

Managing data Lifecycle in the enterprise and the cloud

The British Library, London 8th February 2012



Symantec™



Compellent

Information governance and the management of data is no longer just the concern of the IT department – but has become a key corporate and commercial issue. Never more has it been true to consider information as the lifeblood of the organisation. We have all become only too well aware of the catastrophic operational and financial impact to the business of losing data. Consequently, we have seen a tightening of the corporate compliance obligations now placed upon businesses to align their data strategies to much higher standards.

Information can be your greatest asset and our greatest liability. Managed properly, it can accelerate growth. Managed poorly, it can drive up costs and expose your business to risk. In the face of this, many organisations are now re-evaluating their IT solutions with regard to information management, availability, scalability, efficiency and cost; so it is particularly important, now more than ever, to assess the potential business benefits from managed services in the area of data backup and disaster recovery. The primary objective is to feel 100% confident that your most prized asset – your data – is managed, protected and recoverable, whenever and wherever you need it – and arguably the recovery component is the most important ingredient.

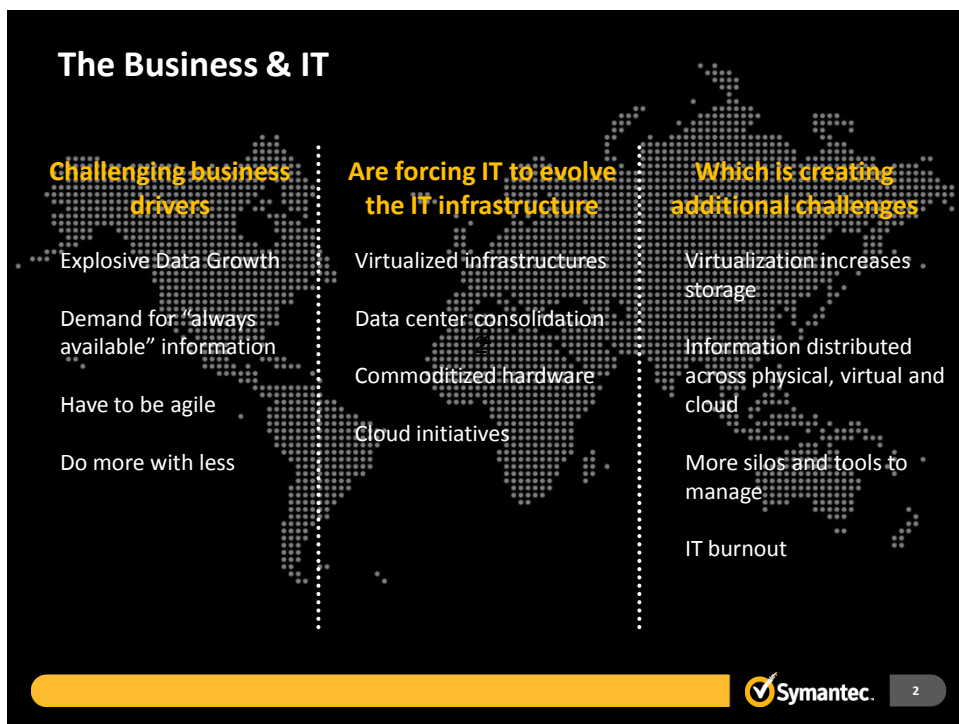
Most organisations are embracing virtualisation – certainly in the back office with servers, and since the turn of the year are beginning to make the shift to virtual storage and desktop service provision... The pulse of the room was typical of NCCs recent cloud and VDI events 60% doing it and 40% still evaluating and realigning their strategic thinking. Certainly it feels more a case of **when** than **if**.... The next 12-24 months will also see a major shift to Cloud services – VDR and managed services in particular playing an increasingly important role as organisations are increasingly losing any appetite for CAPEX.

If you still need convincing – on what, how and where to virtualise your data backup and disaster recovery – hopefully the various formal presentations, not least the user case study outlined a clear routemap of the key drivers. The event focussed on:

- How to manage virtual data backup and disaster recovery
- Understand the real issue around data security and information governance
- Improved disaster recovery solutions
- Improve information reliability and flexibility
- Managed services as part of a cost containment strategy
- Reduce Capital & Operational expenditure
- Overcoming the standards challenge and meeting compliance obligations

Whilst the economic climate remains tough – that doesn't mean doing nothing... there was a general acceptance of the need to do things very differently to meet the demanding business need and the customers' expectations..

Data backup and DR technologies and services have matured – and so too have the strategic models for their adoption. Darren Thomson, CTO at Symantec shared a high level view of the key business drivers across most organisations which have the biggest impact and influence on technology investment.



Data Explosion

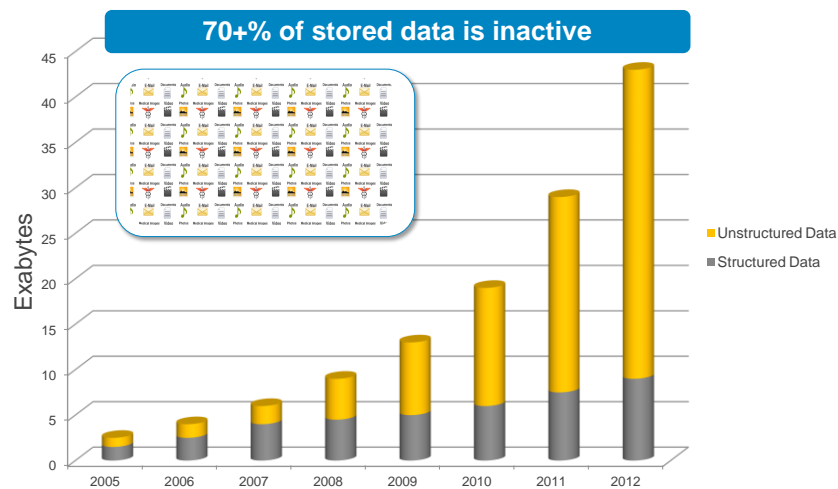
Truth or myth?..... All of the presentations rounded on the impact of data explosion... and whilst there was broad acceptance of the mythical statistics that many of the analyst house throw about there was broad acknowledgement of a year on year growth pattern of around 20%... which presents organisations with huge challenges – not least because of the unstructured nature of most of this data..

Industry statistics show data growth in IT is the single biggest driving Force (Gartner). The proliferation of mobile devices and ability to work from home during off hours yielding 7x24 availability need... Coupled with:

- Industry trend of flat to downturn of IT budgets
- Demand for storage capacity grows relentlessly
- Businesses store richer data – e.g. higher res images and video
- More and more Business Processes are automated – making high-performance online data an increasingly critical resource for all business
- Even emails have become an indispensable business tool for most enterprises
- As mobile devices proliferate, mobile workers generate more and more data as they drive the wheels of commerce ever faster
- And we are centralising more information in enterprise datacentres as we virtualise desktops, servers and move to public and/or private Cloud services...
- Storage has thus become the single fastest growing cost in enterprise datacentres today...

The data deluge continues...

More capacity shipped for enterprise in next 2 years than in prior 20



Source: The Economic Impact of File Virtualization, IStorageNewsletter.com per IDC, May 2010



What this means to IT is revolution to react to changes..

- Expecting >75% virtualization in server market by 2013
- Real estate costs skyrocketing and business combinations yield unprecedented consolidation in IT
- Business can no longer bear cost of custom or specialized hardware

- Ideally make capital hardware expense minimal by “renting” services you need

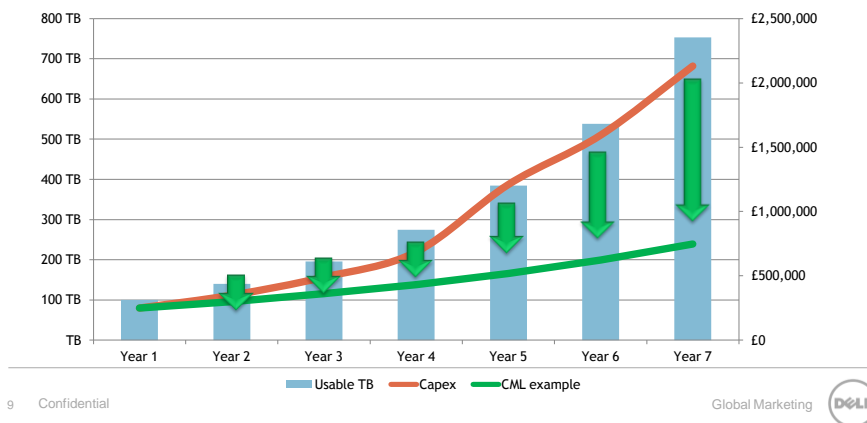
Side Effects

- Virtualisation is not shrinking storage or administrative cost
- Data no longer on traditional tangible owned devices
- New Platforms = New Tools
- Overworked staff burning out faster, increasing turnover

Consequently there is an unavoidable need to cater for rapid data growth in customers’ datacentres... So if we can’t stem the amount of information we must store, THEN HOW CAN WE CONTROL THESE COSTS? Jon Brooks shared an interesting picture of fluid data storage...

Fluid Data Drives Down Storage Growth Costs

- Add Capacity at cost of cheapest drives available
- Eliminate forced end-of-life replacement / avoid forklift upgrades
- Driving down both Capex & direct data centre Opex / power / cooling costs
- Automate and simplify management



Fluid Data Storage Virtualisation means that as new disk drive and IO technologies come onto the market users are able to use this within the SAN infrastructure that they have already purchased. This reduces the need to move to new models, write-off legacy, or pay again for new hardware and software to replace what they already have; or stay with what they have for a few more years. Either way, the customer will eventually have to face the reality that every few years their system will go end-of life.

Change is the only constant

There are significant changes in the role of IT in enterprises from system-centric to information-centric. To date, IT has brought significant productivity benefits to organisations by automating key business processes and driving efficiencies. This landscape was characterised by various business

applications working centralised databases, supported by physical infrastructure. A lot of what IT departments do focuses on managing systems, including PCs, servers, storage, and networks.

In the context of doing things differently, the role of IT is starting look very different. Organisations are looking for the next level of productivity and business agility by improving collaboration and knowledge sharing. They are looking to better connect their employees, teams, business partners and customers to each other. This is changing the nature of data into highly distributed, largely unstructured information. The infrastructure is moving virtual within the company or turning into an external cloud. Instead of focusing of physical systems management, the role of IT is transforming into more information-centric tasks with governance, policies, risks, and controls.

This evolution is also apparent in how applications are being designed. Historically, the entire 'stack' from database, OS, server, storage was defined per application. Today applications are being written to run on virtually any platform and by design will be agile to move from one system to another based on events (performance, availability, workload, etc).

This balance of 'keeping the lights on' and innovating continues be true as IT organisations look to reduce complexity and increase efficiency so a larger portion of the IT budget can be devoted to embracing the consumerisation of IT and enabling the business.

Flexibility with control....

Not only will organisations need to adapt to an information centric way of delivering applications to the business, it must remain in 'control' while doing so. And, in order to remain in control, there are three key challenges every IT organisation must address. Underlying this is the opportunity to enable the business to collaborate and innovate. IT cannot 'lock down' systems anymore, nor can they take an exclusion mentality. They ultimately need to enable the business while remaining in control.

The first is **Information Risk**. Every organisation will have a different risk tolerance profile but every organization ultimately needs to protect their infrastructure, information, and intellectual property. Make no mistake about it; organised crime and malicious attackers are targeting organisations in new and unprecedented ways. They are not intending to disrupt your systems, but rather breach your security, retrieve valuable information, and remove the stolen intellectual property for profit, not notoriety.

Information Growth is another key challenge. Whether your growth rates are 20, 30, 40, or 80%, the impact of spending more every year on storage, on protecting, managing, and discovering information has often pushed IT departments to their limits. By moving away from a system-centric management model, organizations can dramatically reduce the amount of information that is protected. Your data can be reduced by up to 95% through global deduplication. And, deleting what isn't needed dramatically reduces the amount of storage the IT organization needs to manage, while maintaining compliance and ensuring only what's needed is retained and stored.

The third challenge is **Information Trust**. Virtualisation has delivered many benefits through server consolidation. Now, the challenge is to continue to evolve into an agile infrastructure. In short, private cloud computing. Does the IT organisation have the trust and confidence to secure, store, recover, and discover information in this new computing model? This altered landscape may

dramatically increase the exposure and risk to your business without the proper governance and controls in place. This is particularly true with public cloud computing.

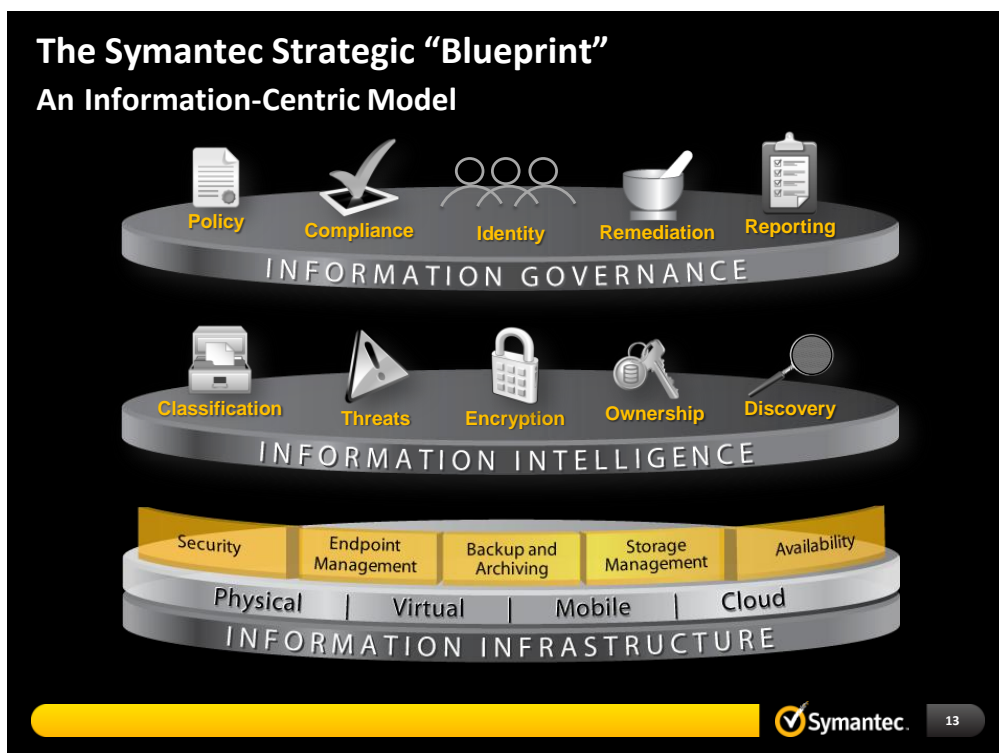
Meeting each of these three challenges is difficult, it must also be balanced against the age-old problem of reducing complexity and increasing efficiency. Typically as much as 70-80% of the IT budget is 'keeping the lights on'. Consequently, if you can reduce complexity and free up existing budget to allocate to embracing information-centric computing and support the consumerisation of IT, that will further enable the business.

This balance of 'keeping the lights on' and innovating continues to be true as IT organisations look to reduce complexity and increase efficiency so a larger portion of the IT budget can be devoted to embracing the consumerisation of IT and enabling the business.

Darren Thomson, presented an information-centric model – that focuses on the journey from physical infrastructure – virtual – mobile and ultimately to the cloud.... most of the room were about half way along that journey some a little further... but most seeing this as a blueprint for future IT service provision.

The Information-Centric Model

The Information-Centric Model is the “vision and strategy” that has been guiding Symantec’s organic product development as well as our acquisitions in response to the need for IT to become more “information-centric” and deliver services independent of any particular system; whether that be servers, storage, or devices; with this model we envision three tenets; information infrastructure, information intelligence, and information governance. This view is very different from others in the industry that are taking the opposite approach and taking a system-centric approach, centered around a database, subsystem, or device.

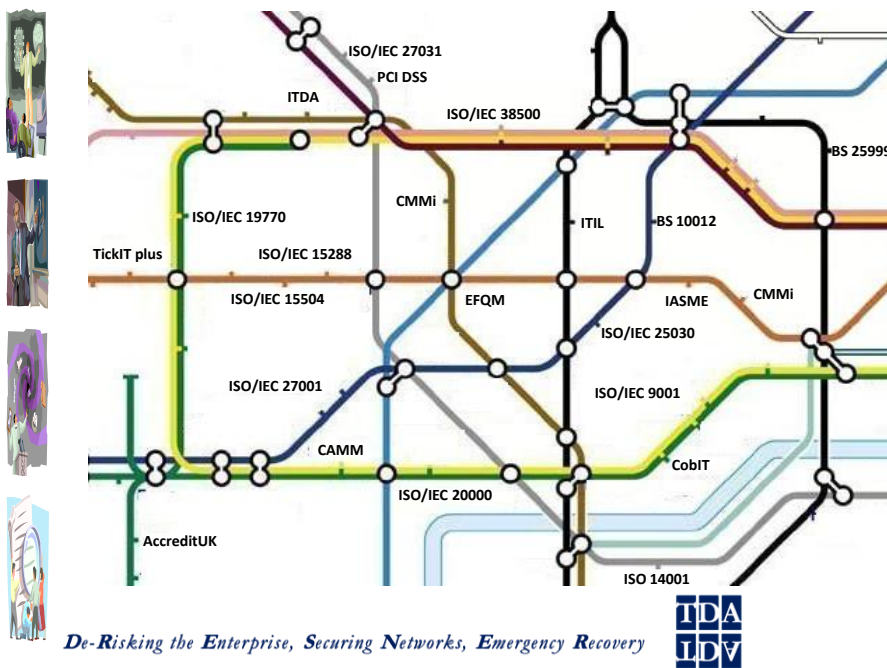


Symantec has a portfolio of heterogeneous solutions from the end-point to the data center. This infrastructure has supported users for 30 years in the system-centric way of delivering applications to their business.

However, it's not enough to provide a sound infrastructure, the value begins to be seen when you can understand the intelligence about your organization's information. Looking at mapping and classifying the information, making a determination as to its relative value to the organization and then beginning to share that intelligence across data loss prevention and infrastructure security is extremely valuable. Additionally, classification can begin to transcend traditional data protection and archiving capabilities into a seamless way to protect information.

Symantec, of course has the ability to protect both the infrastructure and information from sophisticated and very targeted threats.

But by understanding broader data loss prevention policy and compliance standards, actions can be taken based on an event such as trying to copy sensitive information to removable media. That event can trigger a remediation policy to determine one, encrypt the information before copying it or two, and prevent the copying of the information all together.



Not only can ensuring proper access controls be challenging, but answering a seemingly simple question of 'who owns the information' can also be daunting. Storage management and security insight can answer not only the critical ownership question, but also apply governance attributes to the information to protect it. Based on data loss prevention insight, when a policy violation occurs, specific remediation steps can be taken and based on access patterns, ownership information can be inferred and now the remediation process has been dramatically simplified.

An information-centric approach to begin coordinating higher-level governance policies, with the intelligence about the information to then ultimately execute specific infrastructure tasks.

Undoubtedly, there are universal themes to the core drivers... at a strategic level – James Harris shared his thought leadership and drivers centred on data growth, unwillingness to make CapEx investment, complexity and concerns over the business vulnerability over information security.

What keeps you awake at night?

- Real concerns on ability to recover 'really'
- Limited headcount
- See data growth as a factor which will further compress headcount/kit budget
- Reduced TCO and budgets are a key target set by the business
- Eliminate remote site data management complexity
- Technical fatigue and looking to do other things, re-architecture, ongoing issues, media.
- DR is physical, expensive and takes too long to test and recover to
- Offsetting risk and liability
- External audit exposure is also an increasingly important factor
- Data needs to be classified



These factors mapped directly on to those from our user perspective from Jonathan Humphrey from Codemasters Software ... who was challenged by data growth from 3tb to over 200tb in a 10 year period. Challenges of increasing back-up, exacerbated by an ageing infrastructure, exhausted capacity and resource limitations... and a business shying away from CapEx. All in all a foundation for a service requirement to be outsourced or at least controlled through a managed service.

Over the last 10 years...

- Our data has grown from 3tb to over 200tb.
- Increasing challenge of backups.
- Systems are getting easier to manage.
- But we have more systems to manage.
- Customers are more IT savvy and therefore more demanding.
- Networks are getting faster.
- Virtualisation & Cloud are new challenges.
- Downtime is less and less acceptable.

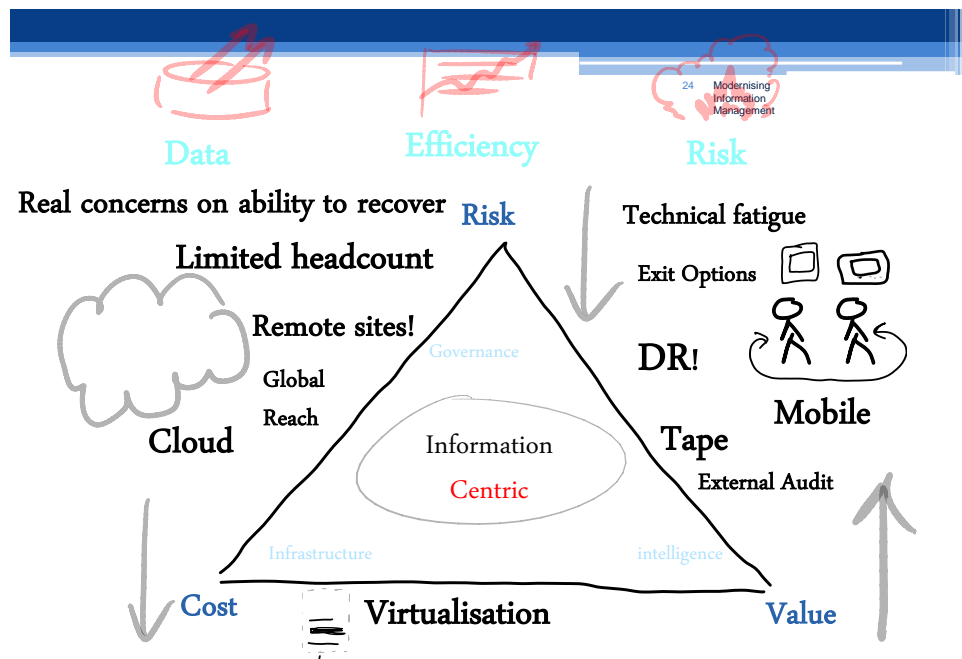


DCG have created a set of framework services – managed services delivering storage, backup and recovery, disaster recovery, security, long term retention and archiving. With a customer base across the globe with multiple petabytes of data under management.

James Harris, DCG Group Business Development Director commented, “the key to any managed service is absolute TRUST.... And this is achieved through proactive management, continual assessment, delivered at lower cost, higher levels of service, with complete peace of mind covering regulatory requirements.

Safety first - A common thread across all four speakers.... was, is my data safe? The short answer is yes – and no. Cloud computing does carry some inherent risks – breaking up and sharing data between servers does create greater vulnerabilities than retaining data on a dedicated server in a known location. These risks, however, are not insurmountable and the key to ensuring data security in a cloud environment is less an issue of what security is in place, but who is providing it. The vital ingredient of trust and confidence featured heavily across all presentations and was apparent in supplier/user relationship between DCG & Symantec, Dell and Codemasters Software.

As well as the integrity and safety of company data, there is also some growing concern of the risks of losing precious intellectual property by placing it in the cloud. Given the variables, transparency is key. Consequently, if your provider is unclear or worse unable to reveal exactly how, and where, the data entrusted to them is handled. This should ring alarm bells. If you, the customer are to be able to properly evaluate the security of your data with a cloud computer service provider, then it is imperative that you are given access to their security architecture as well as being provided with detailed information on data management policies and processes.



In addition to the type and level of encryption used by a provider, some of the important questions to ask revolve around the hardware the supplier uses – do they use storage area network (SAN) solutions, or do they rely on network attached storage (NAS)? What types of V-LAN servers do they have and what switches do they use? There are not necessarily right or wrong answers to these queries, but all can have security implications in certain contexts. Moreover, even with those providers that can demonstrably tick all the right boxes, it is also wise to check what would happen in the event that they are taken over by another provider.

Can they provide both the security and the transparency that your business needs – one key message – go with a dedicated backup provider... an organisation that specialises. The challenge, in what is still an immature market, is to sort the wheat from the chaff.

For the user, the simple truth is that you need to understand the risks and the questions to ask your service provider – here’s some key questions below that came out of the event that will enable you to both benchmark and evaluate potential solutions and the potential suppliers... We are at the start of a huge learning curve with cloud services and you need to undertake the due diligence that your board of directors - and where appropriate, shareholders - are looking for. DCG provided a very useful checklist...

What are the final checks you need to look for:

DCG GROUP
Enhancing the power of information

Security, risk mitigation	Data ownership	Lock in
SLA fit for purpose	Migration pain	Is it cost effective
Trustworthy partner	Internal resistance	Reference

The slide features a 3x3 grid of purple boxes, each containing a checklist item. The items are: Security, risk mitigation; Data ownership; Lock in; SLA fit for purpose; Migration pain; Is it cost effective; Trustworthy partner; Internal resistance; and Reference. The DCG GROUP logo is located in the top right corner, and three small images are at the bottom of the slide.

Questions to ask your cloud service provider

Q1 - Who can see my information?

Data loss is now a reality and a sizeable chunk of all data loss incidents are either down internal or human vulnerabilities or to third party providers. As a result, you need to know whether the service provider, who is the administrator of the system, can see your data. Most have this ability. Therefore, do they have the controls in place to avoid sending, copying, emailing etc etc your data? Remember they will be responsible for your data but, the liability remains with you...

Q2 - What happens if the service provider lost some of your data?

You need to ask your cloud service provider what their data protection policy is and what their audit procedures are. And then you should perform due diligence on those procedures.

Q3 – Where is my data located?

What does the third party organisation do to separate information and systems? You need to stipulate your requirements around co-location and or separation.... Could your competitors - who may also be using the service - get their hands on your data? Remember that, in the cloud, you cannot tell whether your data is copied. So you really need to get this one answered!

Q4 - What happens in the event of data corruption?

How many copies of your data does the third party have? Do they use incremental backups and can they reconstruct an image of your data at a given point in the past from these partial backups. How far back to their backups go in calendar terms?

Q5 - How easy is it to migrate to another cloud service provider?

This all too often a question few companies ask - until it's too late. Porting data between cloud service providers is a relatively new capability and only a small number of service providers have implemented what will become a very necessary service.

Q6 - Are you relying too much on service level agreements?

A service level agreement (SLA) is the contract between you and the cloud service provider that enables them to provide the service... Figures are usually central to most SLAs, but remember things can - and do, sometimes, go wrong, so it is important to agree the remediation process upfront. Above all make sure the SLA you agree is fit for purpose and measuring the things you are trying to achieve.

Also check out the service providers customer base, are these companies comparable to your business, ask for referrals? What percentage of contract renewals does the service provider experience? And don't forget, the success of any partnership will ultimately be down to the level of TRUST that has been developed as the relationship grows – and ultimately this is the cornerstone of an effective business centred service.

Conclusion

The ability to harness technology to deliver new customer services, to operate more efficiently, and to facilitate the collaboration that leads to innovation, characterises successful organisations. Whilst server virtualisation has transformed the IT industry through its ability to both reduce hardware and operating costs and increase IT agility and responsiveness. Most organisations today have virtualised some of their application workloads and are planning to virtualise more. As they do so, they typically progress from virtualising less important applications to increasingly critical applications that are the heartbeat of their organisations.

Virtualising business-critical applications can deliver significant strategic benefits, and, along with consolidation, is a key stepping-stone on the way to a cloud model of IT service delivery.

The session provided a very balanced and insightful routemap – we hope you found this of value... if you'd like to discuss any aspect of the event in more detail or would like to engage directly with any of the speakers please don't hesitate to come back to me... This summary includes only a handful of the slides used, the full slide deck for the session can be accessed at:

<http://www.ncc.co.uk/events/events-archive/>

Other useful references highlighted throughout the morning:

1. The practical compliance journey for business on-line: www.eradardar.eu
2. The Institute of Information Security Professionals www.instisp.org for professional development
3. www.cesg.gov.uk for Information Assurance Standard No. 1 and Business Impact Tables
4. SANS Institute for security education www.sans.org
5. Register of ISO/IEC 27001 certificates <http://www.iso27001certificates.com/>
6. If you use IT, you need IT Governance and that means ISO/IEC 38500

IAN JONES

Head of Research & Content

The National Computing Centre

Email: ian.jones@ncc.co.uk

Direct Dial: 01565 631320

Mobile: 07880 788985

All presentations available at:

<http://www.ncc.co.uk/events/events-archive/>