cds

White Paper | 5 November 2015

Defining the data storage strategy
Is cloud really the answer?

Executive Summary

The visibility and discussions on Cloud strategies have created plenty of buzz and attention throughout the IT and business communities alike over the past few years. With easy access to information, business users are being bombarded daily with promotions extolling the benefits of moving applications to the Cloud, such as reduced operating costs and increased performance and availability. At the same time CIOs have been struggling to reduce cost, deliver five 9s of application availability, and maintain compliance; in the meantime many of their business users have been making independent decisions to adopt or acquire Cloud based applications and services.

This white paper is intended to address the needs of the business user, enabling them to better understand the pros and cons of making Cloud decisions without involvement of experienced IT leaders. It also serves as a resource for IT leaders on how best to approach a Cloud solution that best meets the requirements of the business while protecting their most important corporate asset…..Information.

Contents

Questions remain about hosted Cloud

Cloud services can be broken down into two key categories; public Cloud services use a shared pool of IT resources to deliver applications and infrastructure, and private services that rely on a dedicated hardware platform for each customer. Public Cloud solutions are often cheaper, easier to provision, and much more widely available.

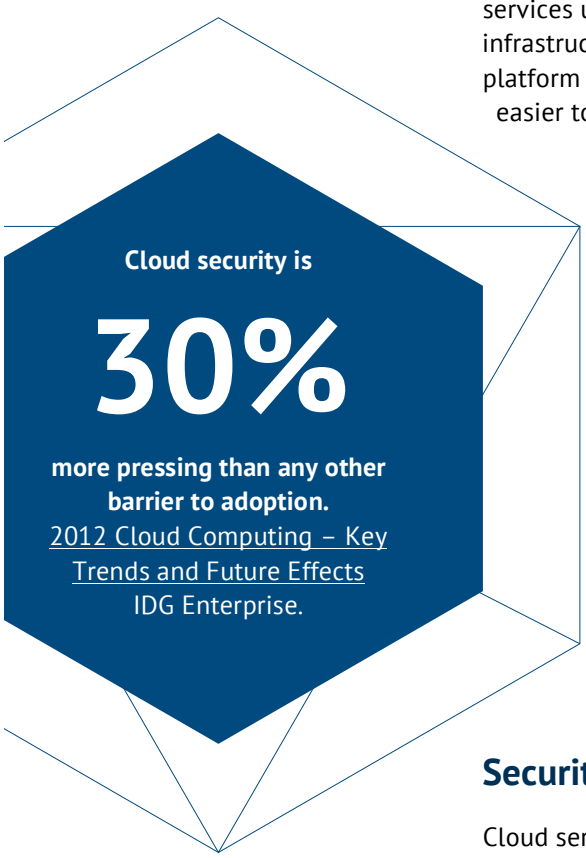

But these headline benefits mask some serious questions that need to be addressed before adopting any public Cloud service.

"Data loss, data breaches, unsecure application programming interfaces (APIs) and shared technology in a multitenant environment are just a few of the concerns expressed by respondents tackling the option of using public cloud. In addition, recent concerns of government snooping in the name of anti-terrorism and general privacy issues contribute to the lack of public cloud adoption."
Laurie Wurster, ***research director at Gartner.***

Security

Cloud services offer a particularly attractive target for cybercriminals as they consolidate vast amounts of data from multiple organisations. By breaching public Cloud defences once, hackers are then able to steal multiple datasets, increasing their reward to risk ratio and jeopardising *every* user of the storage service involved.

"The richer the pot of data, the more cloud service providers need to do to protect it."
David Bradshaw***, IDC research analyst.***

Unsurprisingly, security is the number one inhibitor to Cloud uptake, cited by 70% of IT professionals.

To their credit, Cloud service providers do prioritize security provisioning, but for many CIOs, the very attractiveness of such solutions to hackers presents an unacceptable risk.

Sovereignty

In order to enable maximum resilience and data availability, Cloud technology uses globally-distributed datacenters to prevent localized outages affecting user access. Although this is helpful in terms of meeting internal service level agreements, it presents a massive problem when considering data sovereignty and data export issues.

Corporate information held in US datacenters is protected by US law. But manner of legal and compliance issues:

- Requests by local law enforcement or government agencies in countries where datacenters are located and protection laws are less stringent.
- Litigation and data access requests from third parties in foreign jurisdictions.

Responding to litigation and legal requests in foreign jurisdictions adds significantly to the secondary costs of using and provisioning Cloud services. Because even private Cloud solutions are affected by the physical location of datacenters.

The only way to be 100% protected against sovereignty issues is to ensure that all data is stored in the same country as the enterprise is registered.

Compliance

The distributed nature of Cloud raises further questions of compliance. Choosing a provider capable of meeting regulatory requirements – such as the protection of personal data again loss or theft – is not enough. Businesses must also be able to *prove* that they meet those requirements.

In the traditional onsite datacenter, proving compliance was onerous but possible, as engineers and auditors had direct access to software, hardware and infrastructure. The hands-off management benefits of the Cloud may be a boon when *using* hosted infrastructure, but it also means that auditing for compliance is a major headache.

51%	**42%**	**47%**
of businesses have difficulty getting documentation of a provider's compliance status to fulfil an audit.	**struggle to get documentation of their own compliance status for workloads they run in the cloud.**	**feel that the information they have relating to compliance is incomplete.**

Is Your Cloud Provider Keeping Secrets?, ***Forrester Research.***

Although not impossible, proving compliance across Cloud services adds significantly to the cost of such services; the expense involved increases exponentially for each service in use. Again, the use of Cloud technologies within company-owned datacenters may make better financial sense in the long term, as well as simplifying the process of proving compliance.

Transparency

One of the biggest criticisms of public Cloud services is transparency; information stored in the Cloud cannot be easily viewed or assessed, particularly when stored in a range of application specific systems. In many ways, hosted software (SaaS) has reintroduced the problem of data silos by encouraging a bottom-up adoption approach, rather than a top-down deployment designed by the CIO and CTO to ensure enterprise-wide information flow.

Obviously it is possible to build dedicated platforms on public Cloud infrastructure, but this requires strong leadership to bring existing service deployments back under central control. Businesses will also need to investigate third party dashboard systems that link disparate services to gain insights into usage and maintain some degree of control over resource allocation and usage.

If simplicity is the ultimate goal, private Cloud allows the CTO to fully assess resource usage and allocation and verify that standards are being upheld.

Capacity requirements and costs

As a fully elastic storage resource, Cloud service users are encouraged to upload unlimited volumes of data. But this in itself present problems in the form of increased file duplication, a problem further exacerbated where multiple Cloud services are in use across an organisation.

There are two serious issues that need to be addressed:
- The issue of data accuracy and identifying which is the definitive copy.
- Paying for additional, unnecessary storage that is used to hold duplicated data.

Without being able to control data sprawl, hosted Cloud makes it a little too easy for organisations to break their budgets. When capacity and processing are allowed to expand without proper governance, businesses will end up paying for services that are not actually needed.

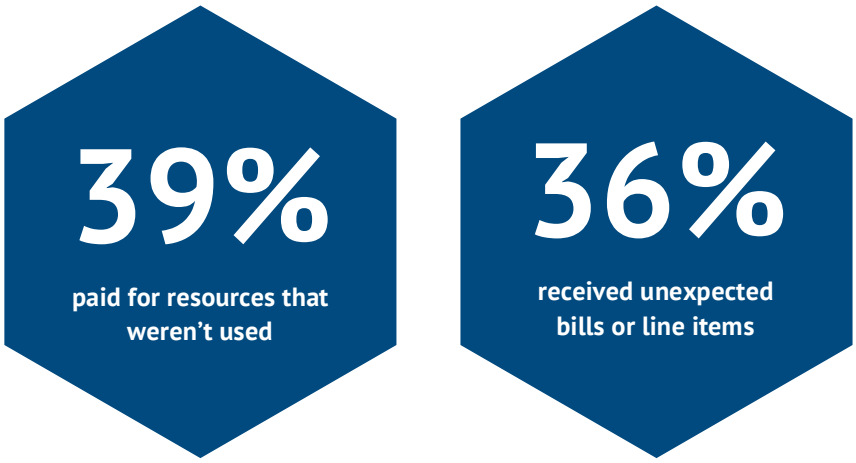

Is Your Cloud Provider Keeping Secrets?, ***Forrester Research.***

Refocusing Cloud strategy and limiting it to a handful of providers for specific use cases, or bringing services back in house with a private Cloud infrastructure, will help regain control of costs and data locations.

Shadow IT

Research suggests that CTOs and CIOs massively underestimate the use of Cloud services throughout their organisations. The average enterprise now uses 1245 Cloud applications – only 14% of which are officially sanctioned.

"Security vendor CipherCloud analyzed a year's worth of cloud usage data from its enterprise customers and discovered that on average, North American companies used about 1,245 cloud applications. Of that number, an astounding 86 percent were unsanctioned applications that IT groups had little idea were being accessed from inside the enterprise network."

Cloud Adoption and Risk Report, ***CipherCloud.***

Shadow IT greatly increases the risk of data loss, leakage, or compliancy failure as there is no central control of safeguards applied to information stored in these services. Often CTOs discover too late employees have been using consumer-grade services for business purposes, greatly increasing problems with compliance, security and sovereignty.

"We may disclose your information to third parties if we determine that such disclosure is reasonably necessary to (a) comply with the law; (b) protect any person from death or serious bodily injury; (c) prevent fraud or abuse of Dropbox or our users; or (d) protect Dropbox's property rights."

Dropbox Privacy policy

The sensible solution is to provide centralised Cloud solutions that offer the required functionality, but which are fully controlled by the in-house IT business unit.

Availability and performance

The distributed nature of Cloud computing is intended to help improve uptime and availability by using multiple datacenters for resilience. However all of the major Cloud services have experienced varying levels of uptime that could have a significant impact on business operations, particularly when they fall short of the 'five nines' (99.999%) uptime internal IT departments strive for:

2013 downtime statistics:

Microsoft Windows Azure	Amazon AWS	Google
96.8%	**99.23%**	**99.34%**
272.04 hours (11.33 days)	**68.18 hours (2.84 days)**	**58.08 hours (2.42 days)**

Downtime Statistics of Current Cloud Solutions, ***The International Working Group on Cloud Computing Resiliency.***

Realistically downtime is virtually unavoidable, but in the case of hosted storage services, customers are completely at the mercy of the provider in the event of an outage. Financial losses may be covered depending on the SLAs attached to a Cloud storage service, but this recompense does little to repair reputational damage or other secondary effects caused by IT systems going offline.

An onsite private Cloud will still be subject to occasional outages, but:

- The owner employs internal engineers able to address issues, without having to 'share' with other service users.
- Third party support providers offer resources that again, are not shared with thousands of other service users.

Maintaining control of corporate storage using private Cloud technologies *onsite* allows the enterprise to be master of their own destiny.

Vendor lock in and loss of control

Finally, there is the issue of vendor lock-in. In an environment where standards are still not fully defined, let alone implemented, transferring data stores between providers is fraught with hurdles:
• Encryption and storage specifics may require additional time-consuming (and costly) conversion.
• Many providers charge based on bandwidth use, making the export/extraction of stored data incredibly expensive.
• The tools used to manage and manipulate data may be specific to the vendor, requiring additional development or retraining.

This uncertainty makes true computing flexibility difficult to achieve; resources may be scalable with a specific provider, but adopting a new service is extremely complex. Until standards are developed and prices stabilise, businesses will find that data committed to the Cloud may have to stay there.

There is little incentive for Cloud providers to develop standards or improve interoperability. Most are instead focused on developing new services and integrations that tie customers further into a proprietary operating environment for their own enrichment.

"One reason you should not use web applications to do your computing is that you lose control [...] If you use a proprietary program or somebody else's web server, you're defenceless. You're putty in the hands of whoever developed that software."
Richard Stallman***, creator of GNU.***

Such uncertainty with third party services equates to a direct business risk. Far better then to choose an environment over which the enterprise has complete control.

The Big Data factor

Sitting on top of all these issues is the increasing importance of Big Data for analytics and automation. The more data collected, the more each of these potential problems is magnified.

Efficient Big Data needs to be able to access all data stores; if they are spread across dozens of incompatible Cloud services, access itself may be impossible. Even where services can be unified, latency and performance will be a massive issue, slowing down processing and reducing the benefits of real-time analysis.

Cost, efficiency, control compliance, ownership, sovereignty – externally hosted Big Data is subject to them all.

Conclusion

Cloud technologies unarguably provide a useful platform from which to develop and deploy applications quickly and cost-effectively. But for the vast capacity demands of the modern enterprise, particularly those pursuing a Big Data strategy, the case for moving data storage offsite is less clear cut.

The pay-as-you-use nature of Cloud is often cited as a way for businesses to reduce their capital expenditure. Rather than paying for excess storage in advance, says the argument, the provider carries the cost. Although fairly sound, this line of reasoning ignores the fact that businesses often have additional unused capacity in their datacenters. This could be unused SAN storage, hardware that can be further upgraded, or post-warranty hardware that is still perfectly serviceable and which can be used to provide additional private Cloud storage and archiving in-house.

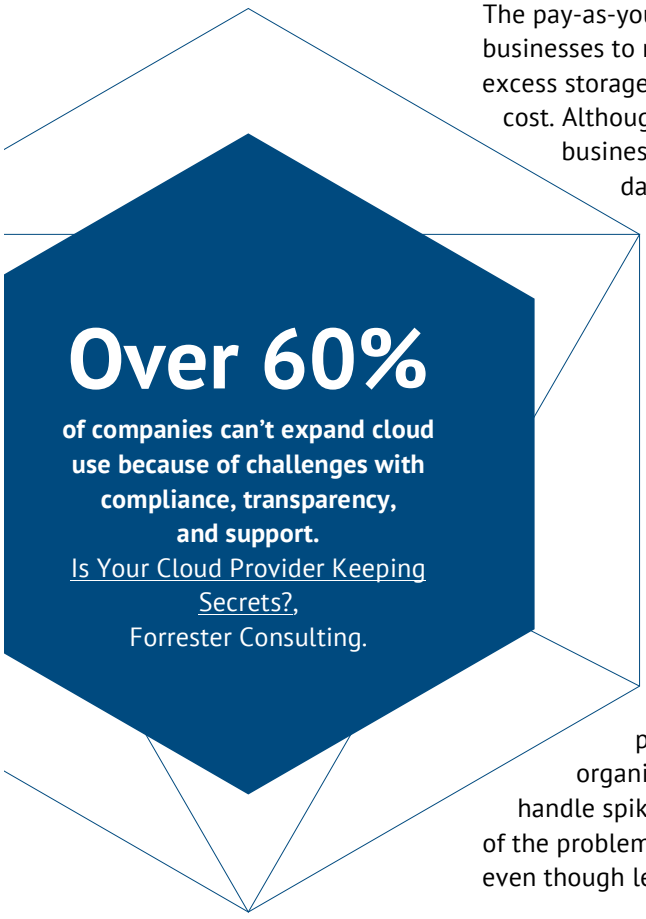

The use of Public Cloud services for mission-critical systems has been almost uniformly rejected by enterprise organisations, with most choosing a Private or hybrid alternative. The reality is that businesses serious about maintaining maximum flexibility and control, keeping data in-house provides the required balance between control and flexibility.

The advent of 'bursting' technologies, offloading processing or data to the Cloud, means that some organisations are creating workable hybrid Cloud solutions to handle spikes in demand. Bursting helps to limit the effects of some of the problems outlined in this whitepaper, but the same risks remain even though less data is being stored offsite.

Cloud technologies and techniques are undoubtedly the way forward for enterprises who need to introduce new flexibility and agility to their infrastructure in order to meet new business challenges. And despite the rapid maturity of hosted Cloud services, the challenges outlined here should give pause for thought to any CTO or CIO considering the use of hosted storage.

Multi-vendor | Multi-Platform | Multi-System | Single Point of Contact

EMC® NetApp® HDS® IBM® StorageTek® Cisco® Sun Microsystems® Dell®

Global Headquarters

Computer Data Source, Inc.
275 Industrial Way West
Eatontown
NJ 07724
USA

+1 732 542 7300
Toll Free: +1 866 237 8008
Fax: +1 732 542 7397
Asset Recovery: +1 732 542 7300
Email: sales@cds.net

DACH

Computer Data Source, Inc.
Deutschland GmbH
Platz der Einheit 1
60327 Frankfurt am Main
Deutschland

+49 69 975 39725
Asset Recovery: +1 732 542 7300
Email: salesemea@cds.net

Canada

Computer Data Source
Canada, Corp.
3780 14th Avenue, Unit 106
Markham, Ontario L3R 9Y5
Canada

+1 905 474 2100
Fax: +1 905 474 2101
Asset Recovery: +1 732 542 7300
Email: salescanada@cds.net

EMEA

Computer Data Source
Europe, Ltd.
Rawdon House, Bond Close
Basingstoke, RG24 8PZ
United Kingdom

+44 1256 362 983
Fax: +44 1256 476 969
Asset Recovery: +1 732 542 7300
Email: salesemea@cds.net

APJ

Computer Data Source Pty Ltd
685 Burke Road, Suite 208
Camberwell, Melbourne
Victoria 3124
Australia

+61-3-9006-1720
Email: jmoshovelis@cds.net