

White Paper

Hitachi Data Systems' Silver Lining

HDS Enables a Flexible, Fluid Cloud Storage Infrastructure

By Terri McClure

November, 2010

This ESG White Paper was commissioned by Hitachi Data Systems and is distributed under license from ESG.



Contents

Factors Driving Interest in Cloud Storage	3
Cloud Storage and Why It Matters	4
Essential Pieces of the Cloud Storage Puzzle.....	5
The Silver Lining: HDS and Cloud Storage.....	6
The Bigger Truth	8

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.

Factors Driving Interest in Cloud Storage

Data growth is one of the few sure things in life and a major IT cost even in a down economy. In fact, for many businesses, data growth is accelerating as more and more organizations use analytics to find new revenue opportunities or to tweak their business models and drive top and bottom line growth. Therefore, it is no surprise that ESG's research finds that keeping pace with overall data growth is the number one storage challenge faced by IT managers today (see Figure 1).

To intensify the issue, the IT market is evolving: the Internet Era of computing is here and commercial enterprises are going to get dragged in—whether they like it or not. Web 2.0, cloud computing, and SOA content and data will coexist with transactional and distributed content in commercial enterprises, and effective data management in this environment will require a new type of storage platform that can provide tiers of storage with distinct sets of price, performance, and functionality in data management to enable efficient storage consolidation with differentiated performance that can support both types of data.

Traditional storage technologies were not designed for use in the multi-petabyte scale Internet Era. With traditional storage architectures, new arrays need to be added as capacity requirements increase. As the number of arrays under management grows, the storage environment becomes increasingly complex, harder to manage, and more costly to operate. This brings negative consequences to the business including increased time to market, loss of productivity, decreased flexibility, and increased risk. Traditional storage technologies continue to excel in the areas they were designed to address—namely, transactional and distributed computing—but these solutions fall short at Internet scale.

Figure 1. Top 10 Storage Challenges Facing IT Managers Today



Source: Enterprise Strategy Group, 2008

The pressure to reduce operating costs and do more with less has never been higher, especially in the face of accelerated data growth. IT is continuing to face limited capital budgets for buying new systems and is running out of floor space, power, and cooling capacity using traditional underutilized and stovepiped storage models. These factors are driving users to consider new storage deployment options like cloud storage.

There continues to be market confusion about what cloud storage actually is, mostly driven by competing vendor definitions. But there are really two sides of the cloud storage coin that need to be considered. On one hand is the cloud service provider and on the other is the cloud infrastructure provider. Cloud service providers deliver application, platform, and/or infrastructure services to end customers via the Internet. Cloud infrastructure providers sell products and technologies, which allow clouds to be built and cloud services to be delivered efficiently either privately and behind a corporate firewall or by a service provider. Users need to consider both options: buy storage as a service or build a flexible private cloud. There are tradeoffs to be made regarding data performance, availability, and security requirements with each approach.

Few storage vendors offer solutions that cross these borders and provide truly flexible deployment options that can meet a variety of business needs. Hitachi Data Systems (HDS) can. It offers the underlying infrastructure components required to make cloud storage a reality and takes the extra step of providing a *private* cloud platform that *can be consumed as a service*, giving users the comfort of having infrastructure inside the firewall and the efficiency of only paying for storage actually used. But before going into how HDS helps users drive efficiency and build a dynamic cloud storage environment, it is important to clarify what the storage cloud is, including its characteristics and differing deployment models.

Cloud Storage and Why It Matters

ESG defines cloud storage as a combination of hardware, software, and processes designed to efficiently deliver storage services, via subscription, to the business. The relationship of the cloud services provider to its end customers determines the deployment model for a storage cloud: a provider of cloud storage services can be a company's internal IT group (a "private cloud" deployment), a third party company that delivers storage services (a "public cloud"), or a combination of both (creating a "hybrid cloud"). In all three deployment models, cloud storage economics enable both the service provider and the subscriber to benefit though each model has its own characteristics.

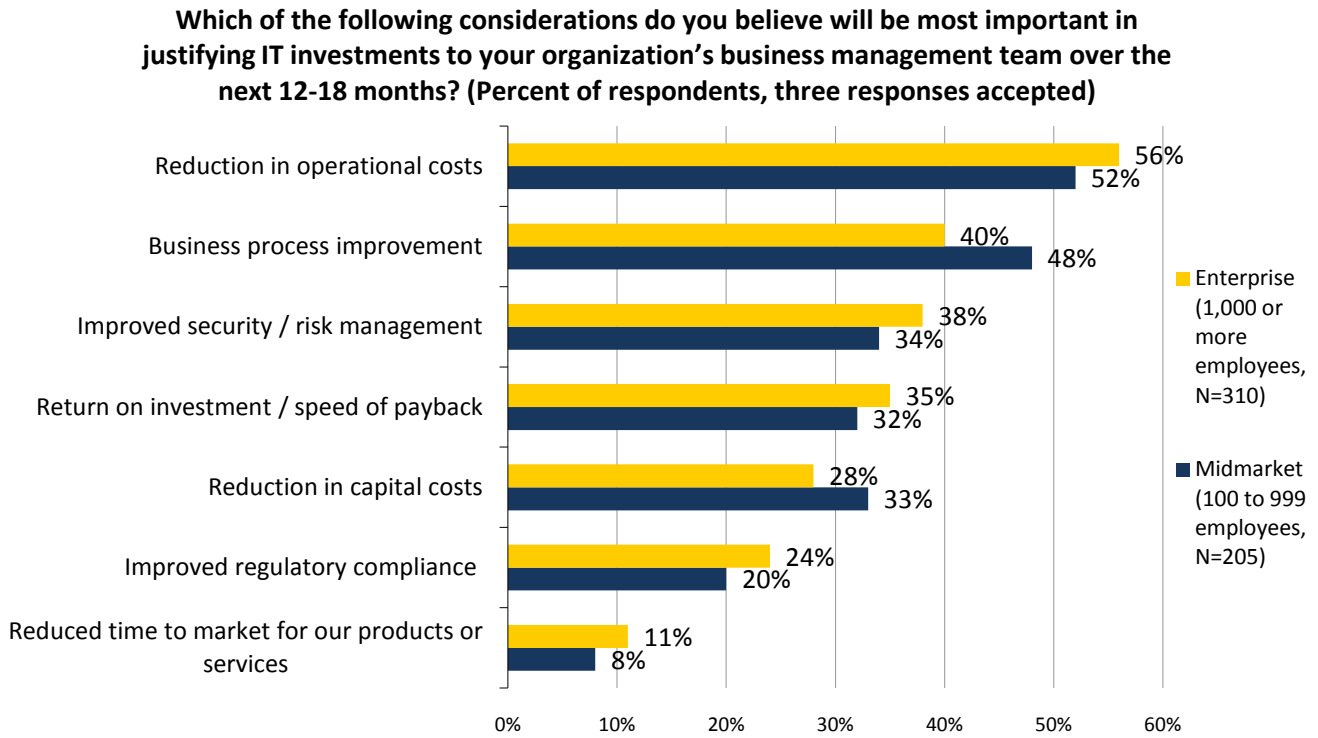
Service providers gain economies of scale via the virtualization of back-end storage infrastructure, consolidation of storage resources on multitenant infrastructure components, and the centralization of management processes to accommodate diverse customer demands. Successful public cloud service providers can anticipate ongoing, predictable revenue streams; it is also possible for private cloud providers to capture revenue from internal customers through chargeback policies.

Cloud subscribers generally will choose public, private, or hybrid cloud models based on the balance their businesses need to achieve in terms of security, control, and cost. While public clouds eliminate the direct costs associated with floor space, power, cooling, and carbon footprint, private clouds can be a better choice for organizations with concerns about their data residing on servers outside their firewall and moving across public networks. For some, a hybrid solution in which sensitive, business-critical data is stored on a private cloud and lower-priority data and files are managed through lower cost public cloud services will strike the right balance.

The benefits to subscribers and consumers of cloud storage are many. The cloud model allows subscribers to consume storage "by the drink," which shifts storage costs to an operating expense. The cloud subscriber no longer incurs expenses to ensure adequate storage headroom that may or may not in fact be used. And because service agreements and SLAs typically cover cloud storage services, operating expenses decrease for the subscriber. Cloud models enable subscribers to balance the value of their data with SLAs and service costs to determine storage tier, security, protection, and performance requirements. On a strategic level, the flexibility of cloud storage services can help a business respond more effectively to rapid shifts in demand for storage capacity and performance and provides a future-proof storage strategy in which storage media and technologies can change behind the cloud curtain without disrupting business.

In the current economic climate, cost reduction is king. ESG recently conducted an in-depth survey of 492 senior IT professionals concerning their organizations' IT spending plans and priorities over the next 12-24 months. For both midmarket and enterprise IT professionals, reducing operation costs far outweighed any other criteria for justifying IT purchases (see Figure 2). The ability to reduce operating costs by shifting to a cloud model is one of the key drivers behind the current buzz around cloud storage.

Figure 2. Reduction in Operational Costs is Key to Justifying IT Spending



Source: Enterprise Strategy Group, 2010

Essential Pieces of the Cloud Storage Puzzle

A number of characteristics must be a part of the cloud storage offering if both the service provider and subscriber benefits are to be fully realized: cloud storage needs to scale quickly and to tremendous capacities. At a high level, cloud storage must be elastic to quickly adapt the underlying infrastructure to changing subscriber demands and be automated so that policies can be leveraged to make underlying infrastructure changes quickly and without human intervention. Of course, the devil is in the details. To deliver on its promises, the storage cloud must be:

- SLA-driven, automated, and integrated to provide quick response times to user demands.
- Policy-based with deep levels of automation to move data to the appropriate tier based on content value and SLAs.
- Secure, reliable, and scalable to multi-petabyte capacity with a unified management view for both block and file storage.
- Multitenant to support the varying workloads, applications, and SLAs required to serve multiple users or even organizations

The last piece in the cloud storage puzzle is that the storage cloud needs to be federated and flexible. Enterprise users have varying degrees of comfort when it comes to storing corporate data offsite, no matter how secure the storage service provider is or claims to be. The need to meet data security concerns and gain a level of comfort with cloud storage is driving three deployment models: private, public, and hybrid. In a private cloud, all components live within the corporate firewall or in a physically isolated offsite storage infrastructure accessed via VPN. In a

public cloud, all components sit outside the firewall in a shared infrastructure that is logically partitioned (multitenant) for security and accessed over secured IP. In a hybrid cloud, some components sit within the firewall—such as a local gateway, storage cache, or bulk storage capacity—while some components like storage capacity are outside the firewall on multitenant infrastructure

Not all applications are suited for cloud storage. Just as data in the data center needs to have performance, protection, and availability requirements defined to map to the right storage tier, these considerations are doubly important when shifting storage to the cloud. Latency is a key consideration—no matter how good the infrastructure provider is, no one has figured out how to beat the speed of light. Latency is not a big issue in private clouds (unless the internal service provider has geographically dispersed subscribers), but it becomes an issue for consideration in hybrid and public clouds. The latency issue makes the cloud most suited as a storage infrastructure for latency-tolerant applications such as backup, archive, and disaster recovery. It is also suited for meeting cyclical storage demands (peak workloads) for enterprise applications, provided the applications can tolerate some level of additional latency.

For global enterprises and service providers, cloud storage can be leveraged for applications requiring collaboration over geographic distances since policies can be set to place data in the geographic region from which it will be accessed most, significantly reducing latency. One of the greatest opportunities presented to service providers by the cloud is that it is an ideal solution for consumers to extend primary storage and backup capacity. Consumers are accustomed to using the Internet to access data, music, and video—latency is not as much of an issue for consumers as it is for commercial users. But commercial buyers beware: one of the challenges in moving to the storage cloud is the ability for the application to write outside its traditional implementation environment, especially in a hybrid cloud. That is one of the reasons NAS-based applications have seen early adoption—NAS is easily ported to a cloud model. And cloud storage services have largely been marketed and sold to consumers, with service providers only in recent year attempting to commercialize offerings for corporate enterprises. Data availability and protection requirements are drastically different for corporate IT than for consumer markets, so buyers need to examine the underlying infrastructure to ensure it is commercial grade.

The challenges outlined in this paper are strikingly similar to those found in the traditional data center—but in the cloud, they are magnified. The bottom line is that the storage cloud is yet another deployment model for tiered storage that needs to offer a variety of service levels to meet user needs. In today's commercial data center, there is no one-size-fits-all answer to storage. The same is true for both cloud storage providers and IT shops offering storage services (private cloud storage). That is why the foundational technology on which a private cloud or the service provider cloud is based is so important. Evaluating service providers for private, public, or hybrid cloud services requires evaluating the underlying infrastructure on which the cloud storage offering is built; the cloud infrastructure must be secure (encrypted, immutable, access controlled), elastic, automated, and based on mature market-proven technology. Hitachi data center experience and reputation for providing rock-solid solutions are core strengths and will serve it well in the both the private cloud and service provider space.

The Silver Lining: HDS and Cloud Storage

Success in the cloud storage space will require infrastructure solutions that:

- Are flexible enough to accommodate any of the three cloud storage deployment models (public, private, and hybrid)
- Streamline the migration to the cloud for both service providers and subscribers
- Provide robust integration between cloud components from the cloud's edge and through its core.
- Provide edge-to-core, secure multitenancy
- Provide seamless, automated storage tiering that optimizes the data storage environment

HDS has developed a cloud strategy, an architecture, and a portfolio that covers these fundamentals and adds powerful features that stand to drive higher value from cloud-based storage services. HDS is already a trusted infrastructure vendor with a broad portfolio of tightly-integrated data center solutions; ESG believes that with the announcement of its new cloud storage offerings, HDS will continue to play a defining role in the emerging cloud storage space.

The latest components added to the Hitachi portfolio of broad cloud enablement are Hitachi Content Platform (HCP) v4 and Hitachi Data Ingestor (HDI). When combined, these two integrated components provide a flexible, edge-to-core platform that forms the foundation for public, private, or hybrid cloud storage models. HCP's wide interoperability and deep integration with Hitachi Virtual Storage Platform enable cloud infrastructure to scale and evolve as new requirements and technologies emerge without disrupting storage services to end-users, creating a truly fluid and almost infinitely scalable cloud architecture.

The "engine" at the core of the HDS cloud architecture is HCP, a massively scalable (up to 40 PB on a single physical cluster) intelligent object store that provides robust management capabilities, secure multitenancy, and a host of features to optimize cloud storage operations for both service providers and subscribers. As a content-focused object store, HCP is adept at managing *both* structured and unstructured data, helping eliminate storage infrastructure silos, and providing a single object view across the distributed, multitenant storage environment. Separate policies can be configured for each tenant on HCP to govern and automate performance, protection, capacity, and retention, allowing the platform to accommodate a wide range of subscriber use cases and business models on a single physical cluster. To facilitate provider/subscriber transactions, HCP provides chargeback capabilities and tools that allow third party management software vendors to write to the API and easily integrate with the HDS solution for billing, chargeback, and reporting. HCP also offers a host of advanced storage and data management features, including automated replication, data compression, deduplication, and multiple storage tiers.

Hitachi Data Ingestor (HDI), as a part of the HDS integrated edge-to-core cloud infrastructure approach, resides at the edge of the storage cloud (for instance, at a remote office or subscriber site) and serves as the "on-ramp" for application data, files, and media to enter the cloud infrastructure and HCP. Users and applications at the edge of the cloud interact with Data Ingestor as they would with a local file server or storage system, but perceive limitless storage capacity. HDI acts as a local storage cache while it migrates data into the HCP-based content cloud, retaining links to stored content for retrieval and maintaining local capacity for more active content. To facilitate chargeback and quota management policies, HCP is able to capture and track storage use from each HDI system. A key benefit of this architecture is that it eliminates the need for local tape backup infrastructure and operations—data protection happens automatically through HCP's built-in protection, preservation, and replication processes.

HDS does not just offer a single, infrastructure-based cloud storage solution. It takes it a giant step further and provides a choice of consumption models so users can balance cost and risk tolerance. Of course, users can buy the infrastructure components to build and manage a private cloud; consulting on how to build and deploy a private cloud infrastructure and providing infrastructure components is about as far as many storage vendors take their cloud solutions. Acknowledging that some users may be uncomfortable with a multitenant cloud service provider infrastructure, yet want to receive the benefits of subscribing to a service and only paying for storage as it is used, is where HDS takes it a step further. HDS offers an innovative service in which HDS owns and manages the private cloud infrastructure deployed within a user's firewall and the user pays as capacity is consumed. This allows users to benefit from the perfect economics of cloud storage services, pay only for capacity consumed, and eliminate storage management headaches since HDS runs, manages, and optimizes to meet SLAs—all within the security and safety of the corporate firewall.

The HDS cloud portfolio offers a range of consumption models and capabilities that allows it to provide users differentiated solutions they can feel comfortable with as they explore the cloud, whether it is private, public, or hybrid. It offers a single platform that supports any data type (block, file, or object) and gives users a single, virtual storage infrastructure platform that can grow and adapt seamlessly as requirements evolve or change. It is a combination of platforms, solutions, and services designed for the new demands put on the storage infrastructure by massive data growth, new data types, and unpredictable virtual server environments.

The Bigger Truth

The potential for cloud storage to ease IT operation burdens is huge: no more long planning cycles and having to buy (and install, power, and cool) capacity well ahead of demand. The cloud is a flexible, liquid infrastructure that allows IT to redeploy resources almost instantly in support of the business, gives greater insight into the utilization and audit trail of the infrastructure, and puts control of storage SLAs in the hands of the application owners and administrators (and therefore closer to the business). Cloud storage will also play a significant role in the “greening” of the IT environment—service providers can build more environmentally friendly data centers thanks to efficient, shared, multitenant infrastructures while subscribers reduce power, cooling, and facility costs not just by moving capacity offsite but also by only paying for what is actually used.

While moving from traditional storage models to cloud storage services can be justified in many ways, many IT professionals are rightfully wary. It is only recently that pieces of the cloud storage puzzle have started to fall into place. The latest cloud storage infrastructure additions from Hitachi Data Systems should help accelerate the adoption of cloud-based storage models as they simplify the design and management of cloud infrastructures, provide a secure multitenant storage core that allows each tenant to govern storage and data management policies via discreet SLAs, and present storage consumers with limitless on-demand storage capacity. HDS is giving service subscribers an on-ramp to the cloud that doesn't require them to rewrite applications or commit to forklift overhauls of their existing IT operations; service providers get a single point of control to manage petabytes of data, files, and media resources associated with a wide range of distributed users and business demands. All of this is done via a wide variety of solutions that let users ease into a cloud storage model at the pace they choose.



Enterprise Strategy Group | **Getting to the bigger truth.**