

Storage in the Cloud – a look at why the cloud will change how you work with your data

March 2009

Has “the Cloud” knocked on your door recently? If not, it soon will. The IT industry is bursting at the seams with vendors rushing out offerings of innovative, hosted solutions accessible across the internet cloud. Some will acquire your attention through marketing hype, or spark an idea in your thinking with a new approach to an old problem. But some of these vendors will simply demand your attention as a responsible IT practitioner, because they will simplify current processes, drastically lower your total cost of ownership, or deliver better IT than you can yourself, while leaving you in control, and making you the hero responsible for improved IT services.

Sound too good to be true? We see the leading vendors in this space already serving up solutions that because of value and capabilities are rapidly overcoming even the hardest objections.

As one example, Nirvanix this past year came to the market with a caching NAS appliance that acts as a gateway to the cloud and serves up high performance from a local device that unleashes the infinite capacity available in the their Storage Delivery Network in the cloud.

As validation of the Nirvanix approach, we are now seeing other vendors hit the market with local NAS appliances that act as a caching gateway to the cloud, and we expect this to be the year when many such appliances, and even software clients, take the market by storm. For the enterprise, a caching NAS gateway appliance now looks to be the defacto standard and a well-tested approach for the use of cloud storage.

At least one recent announcement – from Iron Mountain – also hints at how innovative capabilities can be wrapped around storage in the cloud. As a complement to their recently announced Iron Mountain Virtual File Storage service, Iron Mountain’s Storage-as-a-Service portfolio has a number of information management offerings that are integrated with storage in their cloud, such as data classification, eDiscovery, tape media ingest, format conversion, and other services. One of the more interesting capabilities is a cloud version of Iron Mountain’s long established experience handling physical data transfer in the form of tape. For the cloud, this new service is called a “data shuttle.” Iron Mountain’s data shuttle further bridges the gap between local and remote storage with a “mini-cloud” that is populated with data and shipped out at the customer’s request. When on-site, the data shuttle sits behind the Virtual File Store local caching NAS gateway and acts just like Iron Mountain’s remote cloud storage repository. Customers can use the data shuttle to either upload or request massive

I N D U S T R Y A R T I C L E

amounts of information, without worrying about over the wire transfer. Since Iron Mountain has years of experience in both receiving and shipping out more traditional physical media, customers will likely find data shuttle allows them to physically transfer massive amounts of information with the same efficiency they have when vaulting or recalling off-site tape.

The market is roiling with other vendors getting ready to bring additional innovative services to the market. We expect to see a number of solutions this year that continue to bring the economics of cloud storage to the SMB and enterprise alike, while innovating in how they bridge the physical and technical gaps between local and remote storage. Each of these solutions are rapidly conquering every disadvantage of remote storage, and wrapping unique capabilities around data that might only be harnessed by service provider environments operating at tremendous scale and with tremendous focus on specific problems.

While we are anticipating the announcement of many types of solutions this year, our conversations with end users make it clear that there are a few select use cases behind the biggest and earliest adoptions. We'll will turn to look at these us cases next. If you are an end user, solutions supporting these use cases will be knocking at your door soon. We recommend you begin thinking about how you might use these services now, and turn your eye toward whether the cloud might deliver these services better and cheaper than you can do it with your own infrastructure.

Stop your data from running amok

We've seen many companies come to market over the past 10 years with claims about "holistic" approaches to data management, but frankly, in businesses inundated by the cycle of replacing aging hardware, upgrading software to never ending newer versions, and constantly pursuing new business technologies, any hope for a single on-going and holistic data management solution seems far fetched. When a vendor says holistic data or information management while knowing they have a software roadmap that predicts yearly upgrades, we often think the longevity and consistency of the solution may be suspect. But caching NAS gateways coupled to the cloud look poised to deliver just such a single, consistent platform for data management, and may well be free from the maintenance and disruption of owned and operated solutions.

At the most basic level, the cloud delivers an unlimited second tier of storage that reduces archive to the simple movement of data from a primary tier to a secondary cloud tier that will last forever, without on-going maintenance and management. Businesses archiving to the cloud will be able to dispense with or reduce complex multi-tiered, multiple-system data movement policies and tools, and set archival storage free from disruptive NAS migrations or complex restructuring exercises. For the naysayer stuck to the idea of existing in-house storage: while low-end storage may seem cost effective up front, archiving on cheap SATA disk incurs significant and often unrealized costs in on-going management, power utilization,

I N D U S T R Y A R T I C L E

security, and data protection. The cloud is built to reap economies of scale that will exceed the best homegrown solutions, and can simultaneously provide much better availability, longevity, performance, security and data integrity.

End the acquisition of bad storage

In typical businesses, both small and large, we see a seemingly never ending list of project implementations that depart from strategic storage plans, or even worse are implemented with poorly chosen storage as a complete afterthought. We've seen such projects range from photo radar ticketing applications, to digital video surveillance storage systems, to enterprise document management, and more. In the worst cases, we see these projects end up with isolated local storage, separate protection practices, and periodic complex archiving and/or capacity management requirements. Such departures from strategically planned storage require enormous unrecognized time and effort for management and are at heightened risk for data loss. Storage in the cloud can be useful in place of these poor storage choices while incurring few of the same consequences. Software clients or local NAS gateways can easily serve the needs of these scattered applications, provide access to unlimited storage capacity, and store data with best in class protection – all while delivering total cost of ownership that will shine the light on what an obvious catastrophe other storage-as-an-afterthought solutions create. If your enterprise is managing local or unplanned storage devices today, or looking at new projects that will force you to implement a non-standard solution, you might well be better off looking to storage in the cloud.

Better storage for important data

While the cloud is compelling for the archiving of general unstructured data, some businesses are faced with unique types of data that simply demand the cloud. While there are numerous examples, one of the more compelling is rich content. The world has rapidly become inundated with all manner of rich content, ranging from digital images to digital video used for purposes crossing the gamut from marketing to training. Businesses that create or manage rich media often face a never ending and ever growing demand for storage that can archive and protect this important content. For the largest users and creators of rich content, the cloud may be the only way they can hope to access enough storage for this content. While past attempts at hosted storage often fell short of serving up the necessary levels of performance and availability, this latest generation of caching local NAS gateways has stepped up to the plate, and serves up an ideal home for rich content or any mission critical data. Moreover, the Storage-as-a-Service vendors delivering these appliances have the infrastructure to secure and protect critical data like rich content better than the best of most enterprises.

Meet the information governance needs you can't afford today

Finally, businesses of all sizes are on a quest for better information governance, ranging from simply fulfilling compliance requirements to executing eDiscovery searches instigated by today's increasingly litigious business environment. For many businesses, the cloud can be a gateway to services that would otherwise be unobtainable.

Over the past couple of years, potential cloud storage providers have been on a steady march to acquire best-in-class classification, search, eDiscovery, workflow, and similar technologies. Today, these vendors are combining such technologies with storage in the cloud. The total product is a suite of data services that pairs storage in the cloud with search and information process tools that are just as, or more, capable than those in the largest enterprises. For the mid-sized business, this levels the playing field by steam rolling over what used to be a distinct competitive advantage of larger businesses with more technology and legal team resources at their disposal. This new playing field may create new winners in all types of litigation, ranging from intellectual property, to class action lawsuits.

Sticky issues

At the end of the day, even in the face of compelling use cases, our conversations with end users find the average IT practitioner still has some lingering doubts when considering the cloud. But most of these doubts rest on a missing realization of just how far storage in the cloud has come in capabilities. Briefly, these concerns, and what the cloud is doing to address them include:

- **Security.** Security is often the first business and technology area that enterprises try to assess around storing data in the cloud. But many users approach the cloud with higher security requirements than even their own infrastructure can support. Even so, many of the larger and more trusted vendors have stepped up to the plate around these heightened and often unrealistic standards by building or collocating in data centers that are far more secure than most enterprise data centers, and wrapping storage in NAS-like layers of encryption and protection. Today, users more often than not will find that the cloud comes with fewer attack vectors, and far better data security and encryption than in their own data centers can deliver. For the determined objector who thinks that taking data off-site is not a good practice, it is worth noting that enterprise level hosted data storage solutions are coming to the market from trusted off-site and disaster recovery vendors like Iron Mountain, and these companies likely already store your most critical off-site tape data.
- **Portability.** We've seen some users bring portability concerns to the table when they are considering hosted data storage. In some cases, this is warranted, but we are seeing a number of trusted large providers with solutions that promise portability equivalent to the best traditional local storage. Specifically, when storage over the cloud is combined

with the capability to also deliver your entire data set to your doorstep on your choice of media, or even a dedicated storage appliance, then data is open for easy, high speed portability anywhere.

- **Performance and Availability.** Previous attempts at remote storage clearly suffered from latency issues. Similarly, the nature of the internet could easily wreak havoc on the availability of a service. The break through in this latest generation of cloud storage revolves around caching at the client or within a local appliance that effectively masks internet latencies, while keeping copies of the most heavily used data local. By caching locally, these appliances can mitigate the impact of even extended network outages, but just as importantly, can also make remote storage for frequently used data seem just as responsive as local storage. When turning to a local NAS gateway, storage in the cloud can even mimic the availability, performance, and appearance of typical mid-range NAS appliances, while being protected at the remote end with a level of assurance that few businesses can hope to achieve on their own. Trusted vendors can help select caching capacity and appliance performance that match your expectations about data change frequency and volume, with an eye toward making the storage ideally performant and tolerant of outages typical of your specific internet connections. The single remaining caveat around performance is that these solutions may consume too much internet bandwidth when data change rates are too great, and this may cause your cloud storage solution to incur hidden costs. Yet even here, we expect solutions in the market to continue to innovate around capacity optimization and WAN optimization to rapidly reduce data transmission.
- **Access.** Finally, another common objection is concern about sufficient access when unusually large recall or restore demands occur. As we've discussed, some vendors are entering the game who are able to deliver massive sets of data on media of any type, direct to the enterprise doorstep as rapidly as a copy and ship operation can take place. The best of these vendors can model both the copy and the ship times to your doorstep, and easily let you determine if their capabilities match your needs. Moreover, the innovators in the market are able to ship out a set of components that emulate the original cloud destination with a completely local system, and allow the original local NAS gateway appliance to march on without reconfiguration. If and when the biggest vendors build out more regional facilities, data shipment times will only improve. When unexpected demands occur, the right vendor can get the right data to you. More importantly, the right vendor can likely get that data to you even in the face of a catastrophic loss of your local data.

Making the cloud the home for your archive and backup data

The list of cloud storage services is long and getting longer everyday. Some of the biggest vendors in the IT industry have offerings, including IBM, EMC and others. Meanwhile, innovative startups are everywhere, ranging from Nirvanix to Egnyte. Finally, the trusted

I N D U S T R Y A R T I C L E

providers that already service your data needs in some way are rushing into this market, including vendors like Iron Mountain, an organization that in all probability may already manage some of your data in their tape and document storage facilities. As these vendors with established reputations in the marketplace and the business strength to backup the launch of new services enter the market, look to them to innovate and set new expectations for how the cloud can integrate with your data management needs today. The next few years will likely be focused on archival or long term storage needs, but the new solutions you see in the headlines will be innovative new takes on storage in the cloud, and any one may provide you with a chance to significantly change and improve your current long term storage practices. In our view, the impact of those changes and improvements upon core business capabilities, IT service delivery, and cost of ownership stand to be significant.

***NOTICE:** The information and product recommendations made by the TANEJA GROUP are based upon public information and sources and may also include personal opinions both of the TANEJA GROUP and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. The TANEJA GROUP, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors which may appear in this document.*