

# STORAGE SWITZERLAND

## ARMING THE CLOUD WITH STORAGE



George Crump, Senior Analyst

Storage needs for cloud providers has typically been broken into two groups; storage for Cloud Compute and storage for Cloud Storage. Cloud compute data centers often need scalable, high-performance SAN-based storage since most of their actual processing is done locally within the data center. Cloud storage data centers typically need scalable but inexpensive storage, since most of the processing is done outside of the data center and the storage in the data center is just a repository. Many cloud providers actually have both needs and have to use two separate storage systems to meet user demands. What providers are looking for though is a single system that can satisfy both usage scenarios.

### The Cloud Paradox

Cloud providers also have a unique storage software problem. Unlike many traditional data centers they have developed their own storage management and data services applications or have developed their own file systems. They may have also heavily customized an existing storage application or file system. However, there is also a group of providers that have not invested in developing their own storage platforms and do need complete data services, especially in the cloud compute market where robust SAN storage platforms tend to be installed.

Cloud providers are often forced to buy storage systems with a “bundled” set of data services and then don’t activate them. The problem of course, is that those data services cost money and active or not, make the storage software code less efficient. To compound matters they may have their storage software developed for the cloud storage side of the business and not the cloud compute side of the business. Once again the cloud provider is left needing two solutions. In this case cloud providers need a single solution where data services can be applied when they need it, giving them the option to use their own software or the software that came with the storage system.

### Arming the Compute Cloud

The cloud compute section of the provider’s business typically falls into the first category which needs scalable, high performance, SAN-based storage with data services included. As stated earlier, the storage I/O demand is from applications that, while possibly driven from external, cloud connected users, are processed internal to the cloud data center. These data centers still need to provide their services on a competitive per-use basis, so price pressures remain a top challenge, especially when storage is factored into the budget.

Often the cloud compute storage demand is disproportionately focused on performance more so than capacity. At a minimum the demand for performance vs. capacity scales at different intervals. This situation makes scale-out storage systems less than ideal, since in most cases performance and capacity have to be added in parallel. Scale-out systems also typically use SATA-based disk technology, which may be less reliable, doesn't have the RPM speed of SAS/Fibre drives and presents more latency, from an I/O response-time perspective. As a result many cloud compute infrastructures are built with scale-up storage architectures.

Scale-up storage for cloud compute infrastructures needs to have the ability to handle the initial storage I/O demands and the longer term I/O demands as application usage grows. To do that these systems should use the most powerful storage processors available, like Intel's Jasper Forest Nehalem processor as [PROMISE Technology](#) has chosen for its Ex30 Series of storage systems. These processors provide the capability to drive hundreds of drives at their full-rated capacity.

These scale-up storage systems also need the I/O connectivity between the storage shelves and the controller to deliver the data to the controller as fast as it can now process it. Here 6Gb SAS is the answer. Each SAS controller actually has four 6Gb ports for a total of 24Gb of drive-shelf bandwidth per controller. This means that almost any drive technology can be installed in the shelves and there will be enough bandwidth to deliver full performance.

Host connectivity also has to be flexible. Cloud compute environments are among the first to adopt new storage architectures. There should be options for Fibre Channel, SAS, iSCSI (1GbE or 10GbE) and FCoE connectivity. These options allow the cloud compute storage manager the flexibility to choose traditional fibre channel architecture to leverage SAS' affordable networking capabilities or take advantage of a converged infrastructure using either iSCSI or FCoE.

Finally, these cloud compute architectures are mixed on their need for data services software from the storage controller. They may expect the storage vendor to provide capabilities like snapshots and replication or may rather have a virtualization hypervisor or even the application perform those functions. While most vendors can provide those capabilities few can get rid of them. Certainly the storage manager can choose not to use these features, but once they are in the code they're in the code and consuming some level of storage processing resources. Companies like PROMISE Technology have the ability to deliver a "clean" system that just provides the basic storage management functions, like LUN creation. They also can deliver a more robust set of data services when needed, which can be integrated in. This allows the data center manager to use and pay for just the services they need, when they need them.

#### Arming Cloud Storage

Cloud storage providers' storage demands in some ways are the polar opposite of those in the cloud compute space. Here the focus is on scalable, cost effective capacity. There is however a need for performance, because while all the processing requests are done externally through a slow internet connection, enough of those slow connections becoming active at the same time performance could cause an issue. Consistent but moderate performance is what's needed.

Cloud storage often seems to be the classic use case for scale-out storage and perhaps for the multi-petabyte environment, could be the best solution. Many more cloud providers, even storage focused ones, are never going to reach hundreds of TBs in capacity, and certainly a sub-500TB capability would serve them well for the foreseeable future. Scale-up storage systems that can deliver this kind of capacity may be ideal in this environment. They may in fact have a cost advantage, since multiple nodes don't need to be added every time capacity is needed, drive shelves are simply added.

Many of these cloud storage providers have either developed their own storage software that provides just the data services they need or have purchased an off-the-shelf, but 'cloud storage specific' storage application. In these cases the storage provider doesn't need the data services overhead that may come with storage systems that include data services as well. In fact even the overhead caused by managing excessive nodes in the case of scale-out storage may be a challenge. Their software may be able to manage multiple storage pods similar to scale-out storage but, since there would be less pods vs. nodes, costs would be better maintained and overhead decreased.

Once again companies like PROMISE Technology have an advantage here. Their OEM legacy is tailor-made for the need of the cloud storage vendor who often just needs fast, scalable and cost-effective storage hardware, since they are adding their own storage software on top of it. The PROMISE Ex30 for example can scale to support 192 drives supported by a single active-active RAID head for a maximum capacity of 384TB.

### Scale-In Storage For Arming The Cloud

"Scale-in storage" is the term used to describe storage systems that leverage cost effective but high performance storage processors to provide single system performance and capacity that were previously only thought attainable with scale-out storage systems. As a result they can deliver similar or better performance than scale-up systems with the capacity capabilities of most scale-out storage systems making them an ideal solution for both types of cloud providers, especially those that are delivering both types of cloud services.

### **About Storage Switzerland**

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

*Copyright © 2011 Storage Switzerland, Inc. - All rights reserved*