



Using Object Storage to Reverse Cloud Storage



by George Crump, Lead Analyst

A top priority for IT professionals in 2016 is deciding on a cloud storage strategy. Of the available options, most consider using the public cloud for secondary storage purposes, like backup and archiving, to be the most obvious path to pursue. The public cloud also provides a low upfront cost. It is an infinite storage repository that the user never has to upgrade and, in most cases, it is automatically replicated to another public cloud facility. The big payoff for public cloud storage is supposed to be the elimination of soft costs like power, cooling and storage management. Public cloud providers have built custom storage infrastructures that allow them to keep the administration effort per TB to an absolute minimum, but with all of these advantages, on-premises private cloud storage still may hold the upper hand.

Cloud Storage Advantages

The public cloud allows organizations to rent capacity as they consume it. Traditional on-premises storage requires not only the purchase of capacity, but it also requires the purchase of the "system", the components that wrap around the storage media to make it easier to provision and manage. These "system" costs exist in the public cloud storage as well but are barely noticeable when shared across thousands of customers.

Traditional on-premises systems are also traditionally brittle, often hitting either a performance or a capacity wall that forces an upgrade or migration to a new storage system. The ongoing cycle of storage refreshes is costly and time-consuming. Public cloud storage seems to eliminate that concern for the organization. Again, public cloud provider's storage technologies are custom designed to handle this sporadic growth.

Traditional on-premises storage systems also create a hard cost challenge. Virtually all storage is purchased up-front for 3 to 5 years of use, and not pay as you grow. Most of these storage systems use either purpose-built hardware or software that, while technically off the shelf, has had specific attributes added to it.

The Cloud Achilles Heel

IT planners should factor in two weaknesses of public cloud storage when architecting a

strategy. The first weakness is the obvious latency disadvantage of the public cloud when compared to on-premises storage. While bandwidth to the cloud is increasing, it will likely never catch up with the speed of the internal network. The second weakness is its recurring costs. The low upfront cost of the cloud (getting started with a single payment) is appealing to IT professionals, but in the long-term, ongoing repetition of those payments gets expensive. Often the four to five year total cost is higher than an on-premises system. IT planners also need to factor in the cost of storing this cold data in two cloud repositories for protection from disaster. Any organization that needs to store more than a 100TBs of storage in the cloud longer than five years should do the math to understand what the long-term cost of cloud storage will be.

Hybrid Needs to Evolve

Many vendors will point to hybrid cloud storage as the answer. Hybrid cloud storage stores a portion of the organization's data onsite, and the bulk of it in the public cloud. Because a portion of the data is local, hybrid cloud reduces the concern about public cloud latency. However, it does little, if anything, to reduce the concern about continuous periodic billing costs that end up being more than the cost of an on-premises solution.

The Problems with Traditional Storage

Traditional on-premises storage, of course, has its challenges. First, the upfront procurement costs are excessive. Part of the excess costs stem from vendors selling too much equipment, but also because IT planners want to avoid disruptive upgrade cycles. Most traditional systems are difficult to expand and don't handle a high mixture of workloads. All of this leads to storage system sprawl, which is the opposite of the cloud pay-as-you-grow model.

The Private Cloud Storage Answer

The answer for the data center is to learn from the lessons of the public cloud providers and deploy cloud storage technology in their data centers; in other words, private cloud storage architecture. Today, solutions like those from SwiftStack can provide a more off-the-shelf experience to IT professionals operating a traditional data center. Armed with this technology, IT professionals can gain the same pay-as-you-grow capabilities as the public cloud providers but do so on-premises.

Expansion of this type of private cloud storage system is done by adding additional storage nodes (off-the-shelf servers with internal storage capacity). The new node is easily integrated into the existing architecture and solutions like SwiftStack can automatically and seamlessly aggregate in the additional capacity. These solutions also provide the same ease of use and scalable management that public cloud providers enjoy.

Leveraging private cloud storage architecture allows the organization to cost effectively and safely retain data for a long period. For disaster recovery, this cold data can be easily

replicated to a secondary location, which most data centers with more than a 100TBs will likely have. Alternatively, the cold data can be written to tape for off-site storage as there are several tape solutions with native object storage (private cloud) interfaces.

Reversing Hybrid Cloud

With a private cloud storage architecture as the foundation, data centers can evolve the hybrid cloud model by reversing it. Instead of using the cloud as a long-term storage target, which it is not good at, IT professionals can use on-premises private cloud storage for that purpose. They can then use the public cloud for what it *is* good at: the short-term processing of data.

Cloud bursting is becoming a popular way to use cloud resources. When an organization “bursts” to the cloud they are using the cloud as a temporary set of resources to meet a pressing demand. In the cloud storage use case, capacity can be borrowed from a cloud provider until additional capacity is implemented on premises.

More interesting is the cloud compute bursting use case, as tens of thousands of VMs can be spun-up in a short amount of time. Whether Amazon, Google, or Microsoft Azure, gateway solutions like the Avere FXT Gateway can be deployed in the public cloud alongside “bursted” compute instances. Data that these compute instances need can automatically be cached in the cloud but only while it is needed. Once complete the data is removed and only stored locally, which reduces cloud capacity costs and increases organizational security.

A Reverse Hybrid Cloud Strategy also allows the organization to run their private cloud storage system at a high percentage of available capacity. If a sudden spike in capacity demands occurs, the organization can temporarily move some of the older data to the cloud until the current expansion can be ordered and installed.

Conclusion

The public cloud is extremely effective at temporal tasks, such as throwing 1,000 processors at a data set for a quick answer or result. It is not good at more permanent tasks like storing large sets of cold data for a long period. A private cloud allows an organization to leverage the infrastructures that public cloud providers have built while also leveraging their investment in data centers and trained IT personnel to create a more cost-effective cold storage architecture.

Sponsored by [SwiftStack](#)