

The Death Of Storage As We Know It

Let's be clear, regardless of what you've heard, the onsite storage systems that reside in datacenters today are not going away any time soon. There's no substitute for performance-sensitive, mission-critical applications. Storage systems were taken out of the mainframes of past decades for a reason, to gain performance and redundancy independence from compute systems.

We can expect to see traditional storage systems transition from a general-purpose storage device to a "hot" data storage device. Hot data storage devices will be populated with SSD drives, flash memory, and DRAM. The recent success of all-flash arrays is a testament to this transition.

The Capacity Challenge

One of the biggest white elephants in the storage industry is that eighty percent or more of the data stored on most storage systems is inactive or "cold". With the massive growth of big data over the last decade, the expensive design of traditional storage systems just doesn't make sense for storing data that is not being actively used.

The very design that makes storage systems reliable and mission-critical also makes them unable to cost-effectively scale to support large amounts of cold data. The concepts of RAID (Redundant Array of Independent Disks) that were developed and first deployed in the late 1980s made perfect sense for the last 30 years when data growth rates and disk capacity grew at relatively constant rates. With today's explosive data growth rates and single drive capacities exceeding 10 terabytes, RAID concepts just don't work.

The biggest problem with RAID is that it doesn't work well with low-cost high-capacity disks. RAID system rebuild times after a large capacity disk failure can sometimes take days. Clearly, this is not a workable solution for deploying a low-cost, high-capacity system.

Houston, We Have a Problem

Storage systems that use forward error correction, or erasure coding, provide a means of making multiple copies quickly from the original data source. This concept was used by NASA to communicate with astronauts on the moon. The original message had enough redundancy built into it that the entire transmission could be quickly rebuilt from a message fragment. Storage systems based on this principle are typically called object storage or a cloud.

Object storage systems are well suited for cold data storage. They work extremely well with low-cost, high-density drives simply by making multiple copies. If a drive fails, all the information already exists on other drives that can immediately take over for the failed drive. Object storage systems also scale well on general-purpose servers and can grow into very large storage clusters. The core of Amazon S3 and Google storage is based on object storage.

Unfortunately, object storage deployment outside of large cloud storage service providers is relatively new and has not yet reached a point of being considered a mainstream storage platform for most enterprises. The adoption of the technology has been somewhat slow due to a number of factors:

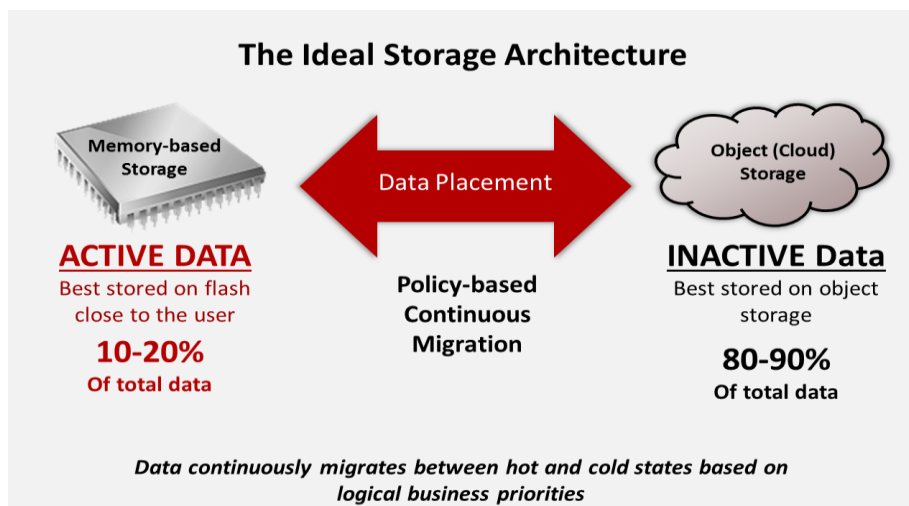
Gateways are needed to convert file or block storage into objects. Think of the gateway as a language translator. They often introduce another layer of complexity by being an abstraction layer between applications and the object store.

Performance is not suited to severing hot data, so in their native form, object stores are best for archiving or applications that are not very transactional.

Installation and migration challenges that come about integrating a different storage technology into existing systems.

The Ideal Architecture is Both Hot and Cold

The idea storage architecture would be similar to a video distribution system. Hot data is kept near the users on very fast media and cold data is stored centrally on a low-cost platform. Building this type of architecture today would combine a flash-based storage array with the intelligence to automatically and continuously migrate data to an object store based on its cold/hot state.



Migration to and from the object store would ideally not be based solely on usage, as most storage and caching systems have done in the past. Since not all data is created equal, migration and movement of data should be based on logical business policies.

Mission-critical data gets the highest performance by not being burdened with coexisting on the same system as cold data. And IT administrators can choose a best-of-breed system to store cold data. In other words, by separating the platforms cold and hot data reside on, an optimal architecture can be achieved.