

# The Cloud Brings Flexibility to Data Storage

*By Duane Tursi and Chris Pupillo*

Imagine your company's financial health revolved around you keeping a large jar on the CEO's desk filled with jelly beans. No doubt your first move would be to keep an adequate supply of jelly beans somewhere on the premises.

Sounds pretty simple so far, doesn't it? Now imagine that some weeks the CEO really digs into the jelly beans while other times he/she barely touches them at all. There's no pattern to consumption so you have to plan to have enough around for the peak times, even if you don't need them for weeks.

Finally, imagine that the number of jars you have to keep filled expands each week from one to two to four and so on, and that you need to double the number of jelly beans on hand just in case the first batch gets ruined. Pretty soon you'll have a lot of money invested not only in the jelly beans themselves but the floor space to store them and the equipment to move them from wherever they're being stored to the CEO's office.

That, in a nutshell, is the situation small and medium enterprises (SMEs) find themselves in today with data. What started out to be a manageable task in the days when the sum of an enterprise's data could be measured in gigabytes has become an expensive, unwieldy mess now that petabytes of data have become typical. And it's an issue that shows no signs of stopping.

In 2009, analyst group IDC estimated the total amount of data in the world was nearly 800,000 petabytes. (Just for perspective, a petabyte is one million gigabytes; if you had a petabyte of DVDs the stack would reach to the moon and back.) In 2010, that number increased to 1.2 million petabytes. By 2020, IDC predicts the total data will be 35 zettabytes, which is 44 times the level it was in 2009; that stack of DVDs would now reach halfway to Mars.

With such exponential growth, SMEs are finding that the old model of building Storage Area Networks (SANs) and keeping all their data on servers inside the enterprise - plus a full backup, plus a mirror of it for disaster recovery protection - is no longer practical or affordable. Server virtualization offers temporary relief, but it merely treats the symptom. At the rate of data expansion it won't be long before even a virtualized internal data center is overwhelmed and cost-prohibitive.

The simple truth is there's no reason to deal with all the complexity and cost that comes with storing data on internal, hardware-based systems that are expensive to build, difficult to manage and ultimately prone to failure despite your best efforts. After all, data is simply the raw material for most enterprises in the way polypropylene is a raw material for manufacturers. The value in the data comes from the applications that turn it into information. Those applications are where the organization should be placing its technology focus.

That's why more enterprises are now looking toward a software-based solution using the "cloud" for their data storage requirements. Rather than building complex SANs in massive, power-hungry facilities to house the data, and then hiring costly experts to manage the entire operation, they're making storage the responsibility of a third-party specialist and simply tapping into the data as needed.

One of the biggest reasons is flexibility in terms of cost per byte. In many industries, data needs ebb and flow based on seasonality, the economy or other factors. Think of our jelly bean example. Enterprises cannot afford to have less capacity than required, so they gear up for peak levels; when data needs ebb, however, they're still paying for peak capacity even though they're not using it.

With cloud-based storage, enterprises only pay for what they use. This pay-as-you-go model is in line with how they consume other commodities, such as raw materials, heating/cooling, electricity and so on. In addition, the cost of managing the data, backing it up, having a cold/warm/hot mirror site, etc. is built into the rate structure, further improving the value of going to the cloud instead of trying to build everything yourself.

Another reason to move to the cloud is to provide your employees with easier access to information. With cloud storage, data can be accessed anywhere, anytime, on any device by anyone with the proper authorization. All they need is a connection and they're ready to work - no need to build SANs for each location or copy files and carry them. Given some of the recent news stories about millions of confidential customer records being compromised due to a stolen a laptop that alone is a good reason to consider it.

If the cloud sounds like a good choice, one thing that's important to understand is that there are actually many cloud models - at least five or six, depending on who is describing it. The most well-known, of course, is the Internet, also known as the public cloud.

The Internet cloud is great for handling your personal email or calendar, downloading movies, doing your taxes and other small tasks. But for storing and moving data around an enterprise it's not the best choice. No one is managing the Internet end-to-end, so if there's a breakdown it could take hours before it's rectified. Ask your CFO how much being without access to your data for a few hours would cost. Security is also an issue. The cost of data stolen as it passes through the public Internet - both hard costs and those to your enterprise's reputation - is nearly incalculable.

A better option is a hybrid cloud that offers the scalability and access of the public cloud with the speed and security of an internal SAN. This hybrid cloud can be delivered via infrastructure as a service, a cloud storage appliance or even a virtualization machine appliance.

With these options, the software that drives them is designed to optimize the delivery of data across a distributed enterprise while using low-cost hardware. This model allows you to operate out of four Tier-1 data centers rather than having to rely on a single, expensive Tier-4 data center with all of its inherent loss-of-service risk. Having a multiple-site configuration across different geographies delivers redundancy and failover capabilities that would normally be beyond the budget of an SME.

The hybrid cloud also improves security by taking data off the public grid. Options also exist to generate and keep private your own encryption key, although you may want to think twice about that because if you lose that key it's gone - along with your access to the data.

Another advantage is that automatic hardware refreshes and software maintenance are built into the cloud model. Your cloud storage provider takes care of all of that, freeing your IT department to perform work that adds more value to the enterprise.

If you are considering moving storage to the cloud, be sure you check out the potential provider, the hardware they use and the options they offer thoroughly. There is no one cloud service that fits all, so it's important to be sure you select one that will work with the way your enterprise works. It's also important to make sure you have the network infrastructure to handle the increased load. If it's been a couple of years since your last technology refresh, you may want to check your current capacity versus the expected flow of data to be sure you're ready for the move.

It's highly unlikely the success of your organization depends on the flow of jelly beans to the CEO - although stranger things have happened. But it does depend on the smooth flow of data to the users who need it.

Moving your data from internal storage to the cloud, particularly a hybrid cloud, can help you store and deliver that data far more cost-effectively while focusing your IT resources on the activities that create value instead. Which is the point of having data in the first place.