



Let's face it: cloud storage is here to stay. When the Storage Networking Industry Association (SNIA) formed its Cloud Storage Initiative (CSI) back in 2009, the working group helped usher in the acceptance of the then-new technology.

Now, however, cloud storage is as ubiquitous as the Internet itself. Twitter alone consumes 300 petabytes of data on Google's cloud, and companies like Facebook, Amazon, AliBaba, and Tencent all provide and consume massive amounts of data storage.

That's why SNIA has renamed the group the Cloud Storage Technologies Initiative (CSTI), to signify a newer, more nuanced approach. Adding the T for "Technologies" expands the group's charter to help support the evolving cloud business models and architectures (like OpenStack' software-defined storage, Kubernetes, and more). The CSTI publishes articles and white papers, offers speakers at industry conferences and presents highly-rated webcasts to thousands of viewers.

SNIA member Eric Lakin joined the CSTI group as soon as he realized that it aligned with some internal initiatives at the University of Michigan, where he manages the storage infrastructure for the University's 19 schools and colleges. This storage engineering team provides a commodity storage platform for Windows, Linux Unix, and Mac users across campus to store files using SMB or NFS protocols.

The cloud is an ideal storage platform for Lakin's team, especially for older data. "We're recognizing that older data doesn't need to occupy a premium physical storage space in our data centers when we can just kick it out to the cloud and take advantage of cloud economics," Lakin says. "The business value of data created at the University is the highest right at the time when it's been created. Over time, the value drops. Most of the files we create are saved and stored away on some shared drive and you never go back to them."

Lakin manages around 10 petabytes of storage. "What we're doing is a very rudimentary information lifecycle management within our on-premise storage system," he said. "We are scanning the file shares, identifying files that have not been accessed in over 180 days, and then we are moving those files transparently out into cloud storage and leaving behind a stub file."

These stubs look and act just like the original files, but when a user clicks on it, the system makes an API call out to the cloud provider, which pulls the file for the end user for access. It's almost as quick as local access, too.

This process frees up local storage space. "Ultimately, when we're doing this at a multi-petabyte scale," Lakin said, "each petabyte of old data that I can push out to the cloud will eventually result in four petabytes of on-premises storage that I can decommission or use for something else."

Lakin shares his successes and challenges with other CSTI members via webcasts, like the one he co-hosted titled "Create a Smarter and More Economic Cloud Storage Architecture" for SNIA.

"As we begin to do this at a multi-petabyte scale over the next few years, we fully expect that we will run into some technical and other challenges associated with scale," Lakin said. The CSTI can then function as a user community within SNIA that other technology teams can reach out to, get answers, solve problems, and help each other.

Anyone looking at integrating cloud technology into their storage environment can join the CSTI and participate in these webcasts and benefit from CSTI's membership like Lakin to help them on their way.

Better yet, it's a vendor-neutral discussion. Members of the SNIA CSTI are encouraged to think about the aspects of the technologies that solve end-user problems, rather than specific products or companies. "It's kind of a safe place," said Lakin. "We can talk with and about the technologies without feeling like we are promoting or selling anything. We can just talk openly to get problems solved."

Moving forward, the University of Michigan hopes to put four petabytes of data into the cloud over the next five years. In addition, Lakin said, his team is in the very early stages of integrating cloud storage into its data backup environment as well. His experiences in that realm will also help others who might be working on their own processes as well.

Lakin's webcast co-presenter, Alex McDonald, is the Chair of the CSTI and has been in the IT business for over 40 years, most recently creating standards-based initiatives, educational programs, and working with industry groups. McDonald was part of the CSI before the addition of the "Technologies" to its name. Early on, cloud storage was considered only good for things like backup or archiving data.

It took a few years to convince users that cloud storage could be used for more than just deep, cheap and slow storage, and that it was something to be used and developed in its own right as part of a more significant application offering. “When Amazon and Google came along, they created these quite wonderful services sitting in the cloud,” McDonald said. “And over time, bandwidth & latency improved and became cheaper. That's really the point at which the game started changing and we began to realize that the storage part of the cloud war was won. But there is still an awful lot of education that we have to do to explain that the cloud had changed.”

Now that more advanced technologies exist in the cloud, CSTI can help end users figure out what's happening with OpenStack, Kubernetes, data services, orchestration, and more. “We can already deliver over-the-cloud block and file storage,” said McDonald. “We know how to do that. We can also deliver key-value type object storage. And we can build new and different cloud applications on top of these storage types.”

No one is sure where this move to the cloud will end up, but McDonald is excited. “Nobody's quite defined the whats and hows of computational storage (where there's compute directly on the storage units) yet,” he said. “But I see this kind of technology as something that's very important when applied to the cloud. And that's the kind of thing we want to be pushing and promoting over the next few years.”

Still, McDonald notes that CSTI must also continue its basic educational efforts. New people come into the industry every day, looking to figure out what the current state of technology is and how it works. SNIA and initiatives like CSTI can help them learn what they need to for their real-life business applications. “It's these important “state of the cloud” messages that we will go out and deliver to help make the transition to cloud” said McDonald. “It's an exciting time of change, and there's so much that's new to learn.”

“The important thing for most of end users is that they don't have a great deal of time,” said McDonald. “I think the CSTI education appeals to them because they can get short little bursts: an hour long piece of education and end up better informed at the end of it, having learned something that's of value and applicable to what they're trying to do. And with a technology rather than a vendor focus, and with our goal to the surrounding technologies rather than cloud storage alone, our education program is more relevant and important than ever.”

About the SNIA Cloud Storage Technologies Initiative

The SNIA Cloud Storage Technologies Initiative (CSTI) is committed to the adoption, growth and standardization of storage in cloud infrastructures, including its data services, orchestration and management, and the promotion of portability of data in multi-cloud environments. To learn more about the CSTI's activities and how you can join, visit snia.org/cloud.