

Advance boosts efficiency of flash storage in data centers

April 3 2019, by Rob Matheson



MIT researchers' modified flash storage drives hold promise to cut in half the energy and physical space required to store and manage user data in power-hungry data centers. Credit: Massachusetts Institute of Technology

MIT researchers have designed a novel flash-storage system that could

cut in half the energy and physical space required for one of the most expensive components of data centers: data storage.

Data centers are [server farms](#) that facilitate communication between users and web services, and are some of the most energy-consuming facilities in the world. In them, thousands of power-hungry servers store user data, and separate servers run app services that access that data. Other servers sometimes facilitate the computation between those two server clusters.

Most [storage](#) servers today use solid-state drives (SSDs), which use [flash](#) storage—electronically programmable and erasable memory microchips with no moving parts—to handle high-throughput data requests at high speeds. In a paper being presented at the ACM International Conference on Architectural Support for Programming Languages and Operating Systems, the researchers describe a new system called LightStore that modifies SSDs to connect directly to a data center's network—without needing any other components—and to support computationally simpler and more efficient [data-storage](#) operations. Further software and hardware innovations seamlessly integrate the system into existing data center infrastructure.

In experiments, the researchers found a cluster of four LightStore units, called storage nodes, ran twice as efficiently as traditional storage servers, measured by the power consumption needed to field data requests. The cluster also required less than half the physical space occupied by existing servers.

The researchers broke down energy savings by individual data storage operations, as a way to better capture the system's full energy savings. In "random writing" data, for instance, which is the most computationally intensive operation in [flash memory](#), LightStore operated nearly eight times more efficiently than traditional servers.

The hope is that, one day, LightStore nodes could replace power-hungry servers in [data centers](#). "We are replacing this architecture with a simpler, cheaper storage solution ... that's going to take half as much space and half the power, yet provide the same throughput capacity performance," says co-author Arvind, the Johnson Professor in Computer Science Engineering and a researcher in the Computer Science and Artificial Intelligence Laboratory. "That will help you in operational expenditure, as it consumes less power, and capital expenditure, because energy savings in data centers translate directly to money savings."

Joining Arvind on the paper are: first author Chanwoo Chung, a [graduate student](#) in the Department of Electrical Engineering and Computer Science; and graduate students Jinhyung Koo and Junsu Im, and Professor Sungjin Lee, all of the Daegu Gyeongbuk Institute of Science and Technology (DGIST).

Adding "value" to flash

A major efficiency issue with today's data centers is that the architecture hasn't changed to accommodate flash storage. Years ago, data-storage servers consisted of relatively slow hard disks, along with lots of dynamic random-access memory circuits (DRAM) and central processing units (CPU) that help quickly process all the data pouring in from the app servers.

Today, however, hard disks have mostly been replaced with much faster flash drives. "People just plugged flash into where the hard disks used to be, without changing anything else," Chung says. "If you can just connect flash drives directly to a network, you won't need these expensive storage servers at all."

For LightStore, the researchers first modified SSDs to be accessed in

terms of "key-value pairs," a very simple and efficient protocol for retrieving data. Basically, user requests appear as keys, like a string of numbers. Keys are sent to a server, which releases the data (value) associated with that key.

The concept is simple, but keys can be extremely large, so computing (searching and inserting) them solely in SSD requires a lot of computation power, which is used up by traditional "flash translation layer." This fairly complex software runs on a separate module on a flash drive to manage and move around data. The researchers used certain data-structuring techniques to run this flash management software using only a fraction of computing power. In doing so, they offloaded the software entirely onto a tiny circuit in the flash drive that runs far more efficiently.

That offloading frees up separate CPUs already on the drive—which are designed to simplify and more quickly execute computation—to run custom LightStore software. This software uses data-structuring techniques to efficiently process key-value pair requests. Essentially, without changing the architecture, the researchers converted a traditional flash drive into a key-value drive. "So, we are adding this new feature for flash—but we are really adding nothing at all," Arvind says.

Adapting and scaling

The challenge was then ensuring app servers could access data in LightStore nodes. In data centers, apps access data through a variety of structural protocols, such as file systems, databases, and other formats. Traditional storage servers run sophisticated software that provides the app servers access via all of these protocols. But this uses a good amount of computation energy and isn't suitable to run on LightStore, which relies on limited computational resources.

The researchers designed very computationally light software, called an "adapter," which translates all user requests from app services into key-value pairs. The adapters use mathematical functions to convert information about the requested data—such as commands from the specific protocols and identification numbers of the app server—into a key. It then sends that key to the appropriate LightStore node, which finds and releases the paired data. Because this software is computationally simpler, it can be installed directly onto app servers.

"Whatever data you access, we do some translation that tells me the key and the value associated with it. In doing so, I'm also taking some complexity away from the storage [servers](#)," Arvind says.

One final innovation is that adding LightStore nodes to a cluster scales linearly with data throughput—the rate at which data can be processed. Traditionally, people stack SSDs in data centers to tackle higher throughput. But, while data storage capacity may grow, the throughput plateaus after only a few additional drives. In experiments, the researchers found that four LightStore nodes surpass throughput levels by the same amount of SSDs.

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Advance boosts efficiency of flash storage in data centers (2019, April 3) retrieved 23 April 2024 from <https://techxplore.com/news/2019-04-advance-boosts-efficiency-storage-centers.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.