



▶ *E-Guide*

Data Center Storage Architecture Gets Smarter with AI

In this E-Guide:

“We all want smarter, more automated, self-optimizing infrastructure -- especially with storage -- but I don't see storage infrastructure components engaging in a human conversation with people about anything anytime soon. Storage is definitely getting smarter in more practical ways, however, and these changes are being seen in places such as data center storage architecture.”

--Mike Matchett, Senior Analyst at SearchStorage.com

Storage vendors are now inserting intelligent algorithms into the storage layer directly to produce near real-time feedback to the business. To that end, there are a few trends converging on storage that are enabling this evolution. Read this e-guide to view the 5 trends evolving intelligent storage and AI.

Data center storage
architecture gets
smarter with AI

Data center storage architecture gets smarter with AI

Mike Matchett, Senior Analyst

Infrastructure is getting smarter by the day. It's reached the point where I'm afraid artificially intelligent IT will soon turn the tables and start telling me how to manage my own personal "lifecycle." Well, I would be afraid if I believed all those AI vendors suddenly claiming they offer AI-powered infrastructure.

Now, we all want smarter, more automated, self-optimizing infrastructure -- especially with storage -- but I don't see storage infrastructure components engaging in a human conversation with people about anything anytime soon. Storage is definitely getting smarter in more practical ways, however, and these changes are being seen in places such as data center storage architecture.

I'm excited by the hot storage trend toward embedding machine learning algorithms aimed at key optimization, categorization, search and pattern detection tasks. Corporate data assets are growing, and so is the potential value that comes from gathering and analyzing big data. It's difficult to manually find those nuggets of data gold, though. And with the coming onslaught of the internet of things (IoT), data prospecting challenges will add mining huge amounts of fast streaming, real-time machine-generated and operational transactional data to the mix.

To help us take advantage of these potential information riches, storage vendors have started inserting intelligent algorithms into the storage layer directly. By converging

Data center storage
architecture gets
smarter with AI

analytical-type processing into the data storage layer, we can now readily tackle the huge scales of information available today and produce near real-time feedback to the business side of our organizations.

There are a few trends converging on data center storage architecture that are enabling this intelligence evolution to happen quite rapidly.

Lambda architecture

In general, serverless, event-triggered computing (e.g., Amazon Web Services Lambda) is growing in popularity as a means to process increasingly streaming, pipelined and event-oriented data. The main idea here is much like the venerable stored procedures or user-defined functions long supported in structured databases.

You can now store and execute event-driven bits of functional programming directly in new, more universal data storage. Custom compute functions can be triggered at a low, intimate level right in the storage layer as data is persisted or accessed -- or in another data lifecycle event, for example, as data ages or migrates to a colder tier.

Also, view this convergence of application and analytical computing into the storage layer on the same evolutionary axis as big data, where storage gets scaled out (e.g., Hadoop Distributed File System) and computation is mapped across cluster nodes local to each chunk of data. The increase in use of in-memory data grids and the arrival of newer big data "universal" databases, such as combined structured and unstructured data sets, are also helping converge computing and data persistence into the same layer.

Data center storage
architecture gets
smarter with AI

Containerized storage OS

Vendors are writing new storage OS architectures as containerized applications, at least internally. This is part of the trend toward a so-called software-defined world, but is also driven by the desire to eventually make computing completely agile over any kind or mix of underlying hardware resources, such as heterogeneous clusters and dynamic hybrid clouds.

Properly containerized storage services can readily integrate and support end-user or third-party functionality inside and intimate to the host storage functionality. Container architectures for storage can then execute microservices on demand to rapidly respond and dynamically scale functionality as-needed -- perfect for the aforementioned Lambda Architectures.

Big data, advanced analytics

In the last century, machine learning usually took place on isolated, historical (i.e., offline) data sets leveraging algorithms designed for scale-up environments. Today, through a decade of big data development, we now have libraries of easy-to-use machine learning algorithms ready and optimized for distributed (i.e., parallel) scale-out applications on increasingly wider volumes and varieties of data.

Streaming data solutions

With the advent of IoT, we are seeing new sources of rich data showing up in the data center storage architecture as nonstop streams of information that require continuous processing in a pipeline manner. The need to process data in real-time, in parallel and with

advanced contextually informed analytics -- not just with traditional transactional business operations -- motivated much of the aforementioned development.

Accelerating technologies

Despite progress in freeing compute and storage from hardware dependencies, it is continuing advances in data center storage architecture that are helping to superpower smarter storage. Every month, we hear of denser memory, the increasing deployment and transition to flash and new nonvolatile memory express-based architectures, more capable CPUs, the leverage of GPUs and even custom field-programmable gate arrays (FPGAs) for big data. And let's not forget persistent memory on the near horizon.

Consider these developments in total, and it becomes easy to recognize how the storage market is quickly growing vast new intelligence capabilities. Of course, when all of this converges, some will argue it's no longer just storage, but that's a different conversation. For now, it seems storage is once again the most interesting place to be in the data center. Certainly, new intelligence can augment traditional data management tasks that implement policies for data lifecycle management. But it can also power other interesting and valuable services over all stored corporate data, such as the following:

- social recommendations
- native storage search
- advanced data security
- data transformation (e.g., transcoding, translation)
- scoring or categorizing data on ingest
- automatic business intelligence analysis

Data center storage
architecture gets
smarter with AI

Machine learning has arrived in the data center at many levels, applications, augmented management and even embedded in devices. IT Infrastructure is getting smarter, at scales and speeds we have only just begun to appreciate. And storage, where most of our data resides, has become machine learning goldmine. While storage is not going to grow a thinking digital mind with equivalent, or vastly superior, human cognition anytime soon, it is going to start to acting in much smarter ways. IT folks today would do well to start looking for infrastructure that learns.

