

# Tech Field Day VFD3 - Coho Data - Coho-ly Moly this is a cool product!

Firstly, we have to start by congratulating Coho Data on the production launch which was announced today which is very exciting for the team.

If you didn't already know Coho Data, they came from their beginnings as Convergent.io with a specific goal of shifting the way that storage is built and managed. They are proud to say that they are a software company that deploys on top of a storage platform. This is where the good stuff happens.

## **I like my storage like I like my weather - Predictable**

One of the strengths of the Coho Data DataStream product is that they scale a predictable, linear rate when adding storage hardware into their deployment. As you add more disk units to the controller, it grows exactly as you are told so that your 180K IOPS with your first unit are simply multiplied as you grow your storage pool at 180K per unit.



Image courtesy of Coho Data: <http://www.cohodata.com/the-coho-difference>

There are a significant amount of measurements happening inside a storage system, as well as with how it is accessed. The discussion we had in our session talked about the realistic gains we can make in physical components (10GbE connectivity, PCI-e SSD cards, Xeon processors), which can be undone by performance challenges at the hypervisor and application layers above it. There is good, old-fashioned physics that we can't dispute, but the real issue happens with software challenges in utilizing that hardware.

## **SDN and SDS - That's a lot of SDx**

There is some not-so-secret sauce behind this which is in the software layers that are creating the efficiency and predictable IOPS delivery with reliability of the underlying data.

We throw around the term software-defined a lot, but that really is what is the heart of what Coho Data has put together. The controller software is handling the distribution of data, read/write caching on flash, and the protection of data for overall integrity.

Add onto that the SDN (Software Defined Networking) components with the built in [Arista](#) switch running [OpenFlow](#). I'm a big fan of Arista and of OpenFlow with what they are doing in their products so this was particularly cool to see the marriage of an interesting storage platform with a forward-thinking networking environment.



Image courtesy of ThinkAhead:

<http://www.thinkahead.com/coho-data-unveils-hybrid-flash-storage-combined-software-defined-netw>

orking/

There is a 4x10GbE uplink which is very hardware-based of course, so there is no doubt that the access to the storage is not constrained by bandwidth nor latency. I would actually encourage that you head over to the Ahead blog where Chris Wahl ([@ChrisWahl](#)) introduced Coho Data which they have been testing in its beta: <http://www.thinkahead.com/coho-data-unveils-hybrid-flash-storage-combined-software-defined-networking/>

## Where is the right place to cache?

This is the ultimate question facing us as architects of a data center environment. There are numerous places in the data stream to put flash resources to accelerate the data movement, with advantages and penalties at every point depending on a number of factors.

For the storage-specific focus of what Coho Data is doing, the flash layer is designed to handle the bulk of the data storage to provide the most accelerated delivery of hot data to the hypervisor. There is also the ability to do write acknowledgements in the flash layer which makes some folks cringe, but there is obviously a near-immediate write that follows to the persistent storage layer to ensure data consistency.

Host based caching is really slick for a number of reasons, however depending on your host environment topology (number of nodes, distribution of workload), there are some situations where the benefits of host-side caching isn't necessarily the ideal choice. I agree that I would prefer to have my storage environment handle the entire lifecycle of the data below the hypervisor and leverage the flash and high-speed bus to the persistent tiers. It gets even nicer when the SDN layer provides access to it ☐

## Want to learn more?

Head on over to Coho Data (<http://www.cohodata.com>) for more information, and you can download the DataStream Whitepaper [here](#):



Make sure to follow Coho Data on Twitter ([@CohoData](#)) and tell them that DiscoPosse sent you ☐

**DISCLOSURE: Travel and expenses for Tech Field Day - Virtualization Field Day 3 were provided by the Tech Field Day organization. No compensation was received for attending the event. All content provided in my posts is of my own opinion based on independent research and information gathered during the sessions.**