

# Pure Storage Competitive Brief

Tuesday, April 11, 2017

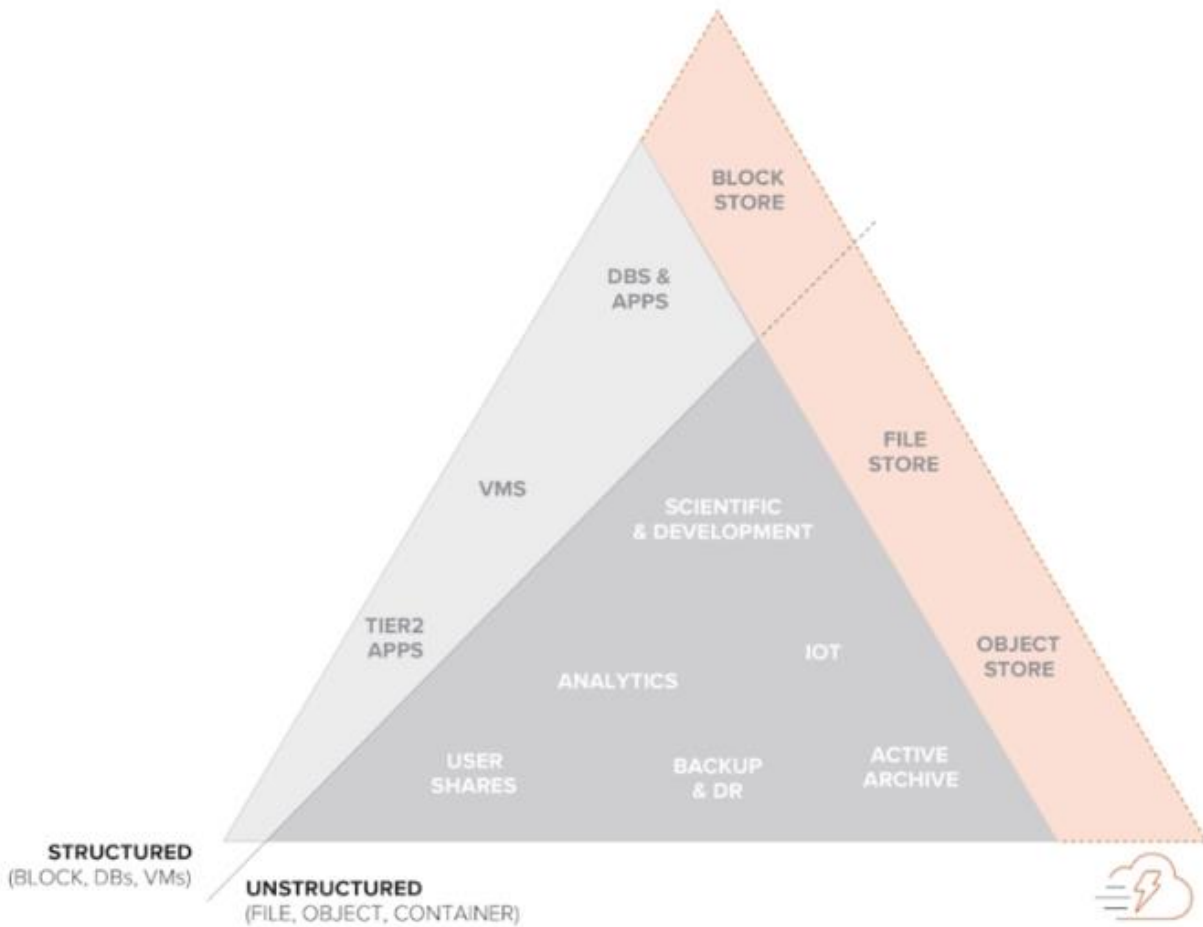
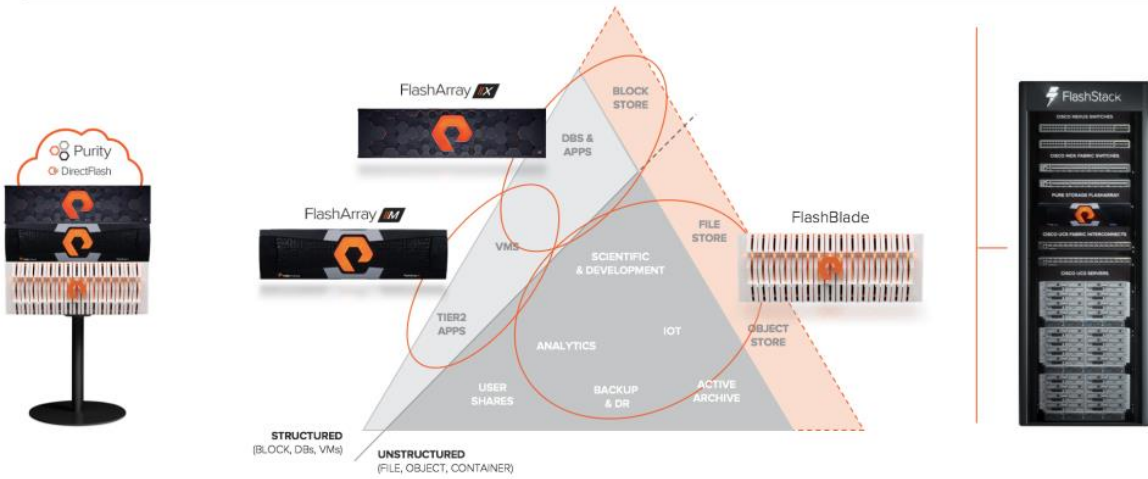
10:14 AM

What did Pure Storage Announce?	The availability of their new FlashArray //X product line.
When will it be available?	Reports on this were mixed initially, but it sounds like it is sampling now in a phase they call directed-availability, and will be fully available in 'early 2H 2017'.
What is it?	FlashArray //X is a shared storage array leveraging NVMe internally, and will be expanded in the future to also support NVMe Over Fabrics to connect to clients.
How is it positioned?	Pure positions //X as storage for the 'cloud-era' where low latency and scalability are required. Sound familiar? It is a block storage device just like the //M line, but they position //M for the legacy enterprise applications. Ultimately, it is a way for them to 'future-proof' the existing product line so that customers feel like they won't have to replace systems once their performance requirements increase, and also allow Pure to try to break into new use cases where shared storage has not typically been used.
What are the quantifiable benefits?	Pure Storage is claiming that these //X storage arrays will offer half the latency of their existing //M product line, and double the bandwidth. However, specific performance specifications were not released. The product line will offer all of the same features as the existing //M product line. They also heavily tout the fact that they are using custom flash modules instead of standard NVMe SSDs, but it is hard to figure out what the customer-visible benefits are that come from this.
Are there other benefits?	Pure is positioning this as a way of future-proofing their existing //M product line, by allowing customers to upgrade the controller module and add new NVMe-capable flash modules to the //M chassis. You would be able to still use the same networking components that exist in the //M chassis. So, customer could upgrade //M to get //X performance, but it is almost a gutting of the system to upgrade it.
Why does Pavilion Care?	The //X product line is positioned as a next-generation NVMe-capable storage array, and we are a leader in that space. They consider the existing NVMe-capable arrays (E8, Aperion, Excelero, Pavilion) as competition.
How does this help Pavilion?	Having companies like Pure entering this space and doing a lot of evangelizing around NVMe-capable storage arrays only helps Pavilion's strategy and direction in this space. In addition, they are also hoping to target similar scale-out database and hyperscale workloads that we are once they add 25/50/100 Gbe support, where direct-attached SSDs are being replaced with top-of-rack flash systems like Pavilion's. This helps validate our direction as well with analysts, investors, etc. And finally, having Pure do this with a product that is inferior to ours means we get the benefit from this evangelization with little cost.

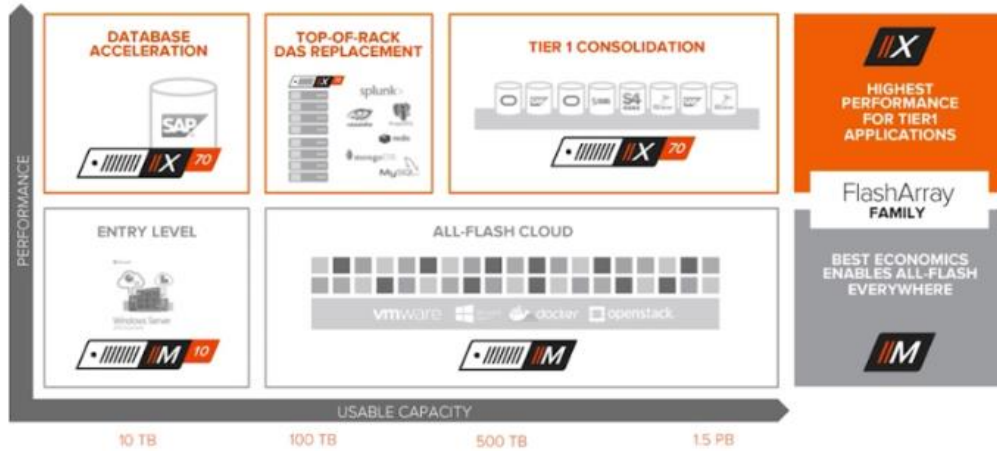
<p>Are there any new features on the //X series arrays?</p>	<p>It doesn't sound like there are any new features outside of performance improvement, since the array uses a software stack that comes from the //M product line.</p>
<p>Does this product compete with Pavilion?</p>	<p>Yes. They are targeting some of the same types of customers and workloads that Pavilion is in hyperscale and cloud environments, but they also are targeting other standard enterprise workloads that they currently satisfy with their //M products. In their favor, they support legacy networking and protocols, as well as a full data management software stack that exists in their //M product line as well as ecosystem interoperability in many application environments, and thus expect the system to be used with traditional enterprise applications in addition to cloud-scale/hyper scale environments. However, again, their penetration into the hyper scale/cloud environment should be fairly limited due to performance.</p>
<p>How much capacity does the system have?</p>	<p>You can put up to 183 TB in a single 3U chassis, using 10 of their removable flash modules (assumes 18 TB each). The smallest version uses It doesn't support expansion shelves like the //M line does. With data reduction technology, they claim you could store 1 PB of user data in a single chassis, but if the workload is latency-sensitive customers may not be willing to pay the latency price of deduplication, particularly if the data is not that reduce-able.</p>
<p>Why does //X use custom flash modules instead of SSDs?</p>	<p>Pure claims that this gives them the ability to 'directly control flash' but they are hard-pressed to identify a specific customer-visible benefit outside of possibly improved endurance. However, that would potentially just improve cost for them, but it is unclear how much and whether it is an actual advantage over vendors who use SSDs. Several identifiable disadvantages exist however, the biggest being that their innovation cycle will actually be slowed down since they have to revise their flash controllers every time they want to incorporate new NAND, whereas vendors like Pavilion can just do a quick qualification of a new SSD to improve performance and/or density. They may also claim a cost advantage just from the fact that they are purchasing NAND directly rather than paying 'markup' to an SSD vendor, but the volumes they would purchase at would dwarf an SSD vendors' volume, so they will likely be paying substantially more for raw NAND than an SSD vendor. In addition, many SSD vendors, including the ones Pavilion uses, are vertically-integrated and have their own NAND, so pricing is pretty compressed already. In the end, I think this is a bad direction for Pure storage in the long run, and their engineering team might have sold them down the river a bit here.</p>

They talked about how NVMeOF front end access coming 'next year' will help drive further up the pyramid displayed below, where high performance block storage is required (hyperscale, SAAS, etc). The orange-colored sections on the right are the 'new' requirements coming from cloud-scale environments, including very low-latency block storage at the top of the pyramid.

# PURE DELIVERS EVERYTHING YOU NEED TO BUILD YOUR CLOUD




# FLASHARRAY MEETS ALL NEEDS



## ALL-FLASH ARRAYS: THE SEQUEL - VIRTUAL LAUNCH

Evergreen™ Storage

EVERY //M CAN UPGRADE TO //X WITHOUT DOWNTIME




PURE STORAGE

# SMALL & SPEEDY, OR BIG & BRAWNY

**JUST 10  
DIRECTFLASH  
MODULES**  
DELIVER MAXIMUM  
PERFORMANCE!



**TOP-OF-RACK**  
CAN REPLACE 100s OF  
UNRELIABLE FLASH CARDS



**1PB USABLE IN 3U!**  
WITH NEW 18.3TB DFM



**ULTIMATE  
CONSOLIDATION**  
FOR TIER 1  
APPLICATIONS



**15 PB**  
/ RACK  
**>100 GB/s**  
BANDWIDTH

