Hardware First - Designing All-Flash Arrays From a Hardware First Perspective

by George Crump, Lead Analyst

Data centers have to meet the increasing performance demands of scale out databases, big data analytics, and dense virtual environments. These data centers need to meet these demands without requiring more data center floor space or consuming more power. All-flash arrays seem to be the "default" answer to today's storage performance problems. But all of these systems are not created equal. How the all-flash hardware and software are designed, and how these two components work together, will impact short-term results and long-term potential of the flash investment.

How Are Flash Arrays Designed?

There are two approaches to developing a flash-based storage system. The first is a software-centric approach that leverages generic off the shelf hardware components. Software first systems tend to leverage form factor SSDs and legacy storage enclosures combined with customized storage software. The other approach is to re-think storage hardware design and customize it to take full advantage of flash-based storage. Hardware first systems tend to offer higher performance in a smaller footprint than the software first alternatives.

The Advantages of a Software First All-Flash Design

The software first all-flash array design has several advantages to both the manufacturer and the purchaser of the system. For the manufacturer, a software-first approach should allow them to get to market faster. While effort behind developing a robust storage software solution can not be understated, it is less expensive and time-consuming than developing purpose-built hardware.

A software first approach, in addition to accelerating time to market, should also allow the manufacturer to develop a robust software feature set. This feature set can include specific optimization to improve overall flash performance and flash durability. A software first approach also allows the manufacturer to use off-the-shelf components, commodity servers and SSDs, which in theory should enable them to reach lower price points.

The key point, when examining software first storage systems is that all of these advantages "should" exist but as discussed below this is not always the case.

The Advantages of a Hardware First All-Flash Design

A hardware first design leverages purpose-built hardware, often custom designed by the all-flash array manufacturer to extract maximum performance from the flash media. Even if the media that is at the core of either a software first or hardware first solution is the same, the package that surrounds the flash media, however, can vary widely. It is in this packaging of flash that the hardware first AFA manufacturer should have an advantage.

The primary advantage that a hardware-first AFA vendor has is that they can use flash in its more native form. Instead of using flash in a drive form factor for compatibility with legacy storage or server

hardware, hardware-first manufacturers typically use flash modules. Drive form-factor SSDs are not as dense as modules, and commodity storage servers have to have larger power supplies to sustain both SSDs and HDD. Using flash in its module form allows for greater density and significantly lower power investments.

Another advantage of an all-flash array is that it has the potential to treat the internal flash storage as thousands of flash die instead of a couple of dozen solid state drives. This means that the hardware-first approach should be able to more efficiently spread workload over a a vast storage pool which will lead to greater resilience and higher performance. The granularity greatly reduces the occurrence of hot spots since the workload can be spread across the wider range of die. A good example is garbage collection, which is easily kept as a background task.

The granular access to the flash die also allows the hardware first flash system to make deduplication an either/or function instead of always on. The more software focused companies, because they count on a low number of flash SSDs, tend to need deduplication turned on to eliminate as many writes as possible. The lack of granularity means that an entire SSD has to be set to read-only or even failed, once a certain percentage of die within that drive have reached their write limits. Deduplication is used eliminate those writes and extend the useful life of the drive but at the potential cost of performance.

The granular access that a hardware-first approach provides is that it sees all of the die individually. A failed die has limited impact on surrounding die, and those die can remain in use. The result is the hardware-first approach can make deduplication a choice to be selected only if it makes sense to for capacity preservation reasons not for drive life reasons.

Finally, an all-flash array has an internal infrastructure that interconnects various components to the flash media. The time it takes that infrastructure to respond to read and write requests determines how well the overall system will perform. This “time”, better described as latency, becomes increasingly apparent as the number of and variety of, workloads increases. The result is that a hardware first AFA should be able to scale further, take up less data center floor space and handle a higher workload mix.

The Misconceptions of a Software First All-Flash Array

Many of the theoretical advantages of a software first AFAs are misconceptions. For example, the reality is that most of the hardware-first AFA on the market today were actually on the market prior to the software-first solutions. The software-first manufacturers used the combination of software and commodity hardware as a means to catch up with the purpose built solutions.

The software first vendors benefited from the reality that initially flash media was so fast that it improved performance regardless of the deployed form factor because the use cases were so isolated. Initially overlooked, the latency caused by combining software with commodity hardware designed for hard disk drives, became apparent as the use of flash evolved in the data center. As flash usage becomes more widely understood, extracting maximum performance from the flash investment is becoming more critical.

AFAs are now expected to support more than one workload type, denser workload usage and developers are designing applications with the expectation that flash storage will be the standard. The result is that “more performance than you need” is becoming an outdated statement, and driving out AFA system latency is a top priority. In fact, some previously software first manufacturers are now offering purpose-built hardware solutions.

Another misconception of software first all-flash solutions is that they are less expensive than purpose-built solutions because they use commodity hardware. Again reality does not necessarily prove this to be true. First, a hard cost comparison shows that software first and hardware first solutions are comparably priced. Second, because hardware first solutions tend to extract more performance from the same flash investment their long-term costs are often superior. A hardware first architecture can scale performance further while requiring less power and data center floor space.

A final misconception is that hardware-first AFA vendors are behind on software feature sets like deduplication and compression. The hardware first AFA vendors in actuality are also benefiting from a software-driven world. Hardware first vendors and their customers have their choice of software solutions from which they can partner, buy or even develop/enhance internally.

The Best of Both Worlds - Hardware Forward with Optimized Software

Hardware-first vendors, with the time investment in hardware complete, actually have a software advantage. They can provide a turnkey solution with their software included, they can integrate with other storage software solutions without requiring the customer to buy storage software twice, or they can offer the AFA hardware with no software for use in extremely high-performance use cases that require 500k+ IOPS.

Conclusion

When IT professionals consider an AFA, they typically focus on resolving a pressing storage performance problem. Which, unless the problem requires millions of IOPS, is usually addressed by either type of AFA approach. By definition, this more standard use case for all-flash leaves the IT professional looking to do more with the investment than just solving the upfront problem. Over time, the IT professional will want to expand the AFA use case, and as they do they will begin to see some of its design weakness exposed.

Expanding the use case typically means more workloads, with increasing variety and developers expecting flash performance all the time, instead of an occasional speed bump. Greater number of workloads and an increase in variety of type will expose latencies in the hardware design of the system. A flash performance expectation will require more consistent performance regardless of the current workload demands.

Purpose built hardware is may be better suited to handle the I/O demands of an expanding use case and an expectation of “always on flash” performance. Flash Arrays may hit a performance wall, not because of the limitation of the flash media itself, but of the limitations of the hardware that surround that media. Purpose-built systems incorporate flash media with hardware designs that are optimized to keep latency low and performance consistent no matter what the load.

Sponsored by Violin Memory Systems

About Violin Memory

Violin Memory transforms the speed of business with high performance, always available, low cost management of critical business information and applications.

Violin’s All-Flash optimized solutions accelerate breakthrough CAPEX and OPEX savings for building the next generation data center. Violin’s Flash Fabric Architecture (FFA) speeds data delivery with chip-to-chassis performance optimization that achieves lower consistent latency and cost per transaction for Cloud, Enterprise and Virtualized mission-critical applications. Violin’s All-Flash Arrays and Appliances, and enterprise data management software solutions enhance agility and mobility while revolutionizing data center economics. Founded in 2005, Violin Memory is headquartered in Santa Clara, California.