The all-flash data center: a reality?

Flash memory has been introduced in data centers few years ago through the use of SAS and SATA Solid-State Drives (SSD). This technology allows very fast read and write IOs compared to spinning disks (HDD). It was first used for applications requiring high performance storage capabilities, as a top tier level only due to the higher cost in term of dollars per gigabyte compared to hard-disk drives. By the way, the total cost of ownership (TCO) of flash-based servers as dramatically decreased, therefore allowing the use of flash-based products for lower tier storage applications. This is thanks to a lower $/GB, software optimization allowing advanced compression algorithms, power consumption reduction, and the effort on higher density form factors. That leads to a huge market opportunity for new flash-based applications. Flash-based products will be designed through a standard server architecture and commercial SSDs, or through a custom all-flash array architecture. In this paper, both design strategies will be described and how the designers will be available to leverage the flash capabilities.

There are 2 main approaches to design a Solid-State Drive (SSD): the first one is to integrate standard components and add a software layer on top of that, which is the main added value of the SSD makers. This is efficient in term of time to market and cost for low/medium volume. The other way is to design a custom architecture based on a FPGA or an ASIC. That allows the implementation of a specific nandflash management at the hardware level, like the number of nandflash channels and how they are managed. In addition to the software layer which is still required, the added value of the SSD makers includes the hardware architecture where the main benefits are the SSD performances, such as a low latency and high IOPS throughput. Time to market is still a key factor, that's why using a first generation of SSDs based on a FPGA is a common strategy, providing design flexibility. In addition, in order to focus design resources on the added-value SSD controller architecture, it is not efficient to design the IPs based on a standard interfaces, such as the nandflash controller or the host interface such a NVM Express (NVMe). NVMe is the host interface defined for the new generation of PCIe SSDs. Both NVMe and nandflash controller are well defined through a specification (e.g. ONFI for the Nandflash interface), then, using interface IPs from third-party partner is highly recommended.

The main advantages of SSD compared to HDDs are the great performances in term of latency and throughput. There is at least a decade or two between the two technologies. There are multiple ways to improve the SSD performances, including advanced and optimized software algorithms for flash management (wear leveling, bad block management, garbage collection, flash translation layer…) and hardware channel management. Only a custom SSD controller architecture will allow to benefit from a performance increase regarding the nandflash management.

For the all-flash array, the design strategy is comparable to the SSDs: using standard components or designing a custom architecture. The first strategy leads to an easiest hardware design, using a server mother board and SSDs with SAS, SATA or PCIe interfaces. But compared to SSDs, there are more possibilities to provide added value through software layers. Most of the all-flash vendors provide a full software suite based on a proprietary

operating system. Software feature often includes deduplication, compression, snapshot…Custom hardware design architecture will allow to benefit from more nandflash performances and higher density form factors. In this case every part could be customized, from the main storage server controller to the memory module. Since it is a fully integrated server, the use of standard interfaces between the main controller and the flash module is not mandatory, and could be based on proprietary protocol. With this specific design, very high density server may be obtained. Only custom flash controller will allow to design flash module with the required nandflash interfaces. As for the SSDs, using third-party IP partner is recommended in order to reduce the time to market and the design effort, focusing on the all-flash server architecture.

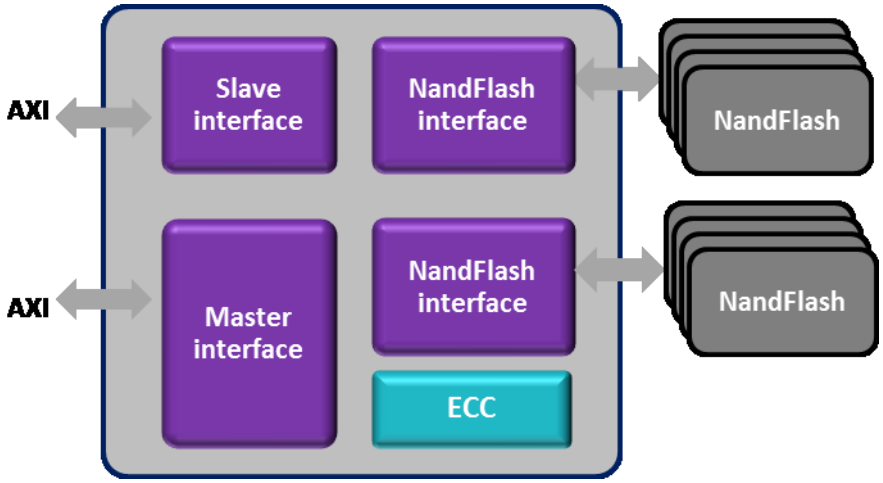

IP-Maker has developed a Universal NandFlash Controller (UNFC) that brings performances and design flexibility to SSD and all-flash server makers. It is ONFI 3.2 compliant. This specification allows throughput up to 400MB/s per channel, which is ideal for flash-based products requiring a high level of performances such as PCIe SSDs. The IP-Maker's UNFC IP core has designed specifically to enable commodity Flash memory to be effectively used in enterprise storage applications requiring high reliability and large interconnect bandwidth. That includes the support of SLC, MLC and TLC NandFlash memories. The UNFC comes with a proprietary interface for full custom designs, or with standard AXI interface for an easy SoC integration. It is a full configurable flash controller, providing flexibility in term of page size, spare size per channel, and internal data bus width, allowing a tradeoff between gate count and performances. The UNFC comes also with a built-in Error Code Correction (ECC) controller, performing error detection and correction from the nandflash data. It is based on BCH algorithm, and it is also fully configurable in term of number of errors to detect/correct and page size. The ECC is fully hardware managed for performance optimization.

The all-flash data center, a reality? Definitely yes. The different design techniques allowed a great TCO reduction, and the gap between flash and spinning disk will be reduced. Of course, that will require a couple of years to migrate to flash only. As tapes are still used for archive, HDDs will still be used for cold data. On the other hands, with the introduction of new smart controllers, the use of flash on the memory bus (DIMM form factor) became a

reality. This is also a great market opportunity for flash-based products manufacturers. The IP-Maker UNFC nandflash controller is the ideal solution for all these flash-based products, providing an optimized, integrated, flexible and validated interface IP. Flash-based products manufacturers will then focus their R&D effort on the top level architecture, providing a higher ROI.