

Shorten Time to Market for NVM Express based storage solution.

Ravi Thummarukudy, CEO, Mobiveil, Inc
ravi@mobiveil.com

A number of technical and business trends are converging to create a booming market for solid state drives (SSDs), with gigabytes of flash memory capacity along with the related control electronics packaged in the form factor of a 1.8"-, 2.5"- or 3.5". storage device. The first is the emergence of tablets and pervasiveness of smart phones both of which use flash as their main storage. The resulting demand has created innovation in manufacturing and packaging of flash and is driving declining storage cost and increasing the R&D investments to improve flash subsystem performance. The second trend is the continuing growth of "big data", data analytics and the resulting need for faster transaction processing on the ever expanding and proliferating server/storage data center farms. SSDs are being rapidly deployed to increase transactions per second at a fraction of the cost otherwise required to achieve the same result. Lower Power consumption in SSD drives is another important factor that works as a tail wind push to this growth trajectory.

This white paper will examine these trends and in detail the underlying interconnect technology revolution that is enabling SSDs to meet the high performance that servers are demanding. We all have experienced that flash based storage in consumer devices (Tablets, Smart phones and Notebooks) is giving much better user experience. But in the server space a new standard is fast emerging named NVM Express (NVMe), a PCI Express® based scalable host controller interface that uses solid state drives to serve as the data storage element for enterprise, data center and client systems. Defined by an industry group that includes Cisco, Dell, EMC, HGST, Intel, LSI Corporation, Micron, NetApp, Oracle, PMC-Sierra, Samsung, SanDisk and Seagate, NVMe has been architected to deliver unprecedented performance that is demanded by the cloud based servers of today.

The Solid State Drive Market

A Google search of the solid-state drive market lists recent reports from Trendfocus, Leapfrog Research, iSupply, Technologies and Markets, among others. [Trendfocus' finding](#) of 50 percent growth from first quarter 2013 to first quarter 2014 is typical of what most researchers are predicting for this technology. As expected the greatest unit shipment volume is for clients: notebooks, netbooks, and tablets. But, it's in the enterprise where the bigger evolution is occurring: "The enterprise SSD market is gaining more traction in data centers and cloud storage, as the dollar per GB (\$/GB) metric decreases to levels low enough to drive an inflection in demand. Total cost of ownership (TCO) and performance arguments are far more compelling at today's \$/GB levels and with costs expected to continue downward, enterprise SSD adoption will

continue to grow at rates outpacing other storage devices.” In October 2013, the price being quoted for SSD storage was \$0.85/GB.

Two key trends are enabling the adoption of the NVM Express standard quickly. The first one as pointed out by Tiernan Ray writing in *Barrons' Technology Trader Blog*, March 30, 2013, "[Google and Facebook Are Shaking Up the Server Market](#)." He asserts "The 'Net giants are shunning traditional suppliers for cheaper, made-to-order hardware." In other words, these deep pocket denizens of the web are buying their servers from Asia-based original design manufacturers (ODMs) such as Inventec, Quanta or Wistron. Furthermore, because of the rise of the virtual server, these compute engines are added into compute farms to accommodate growth in web traffic. The main driving force behind these web companies' "roll-your-own" server strategies is cost. It provides for an incredibly low-cost infrastructure, Tiernan writes in his blog. NVM Express provides a standard based approach that enables a consistent feature set that allows many companies to develop PCIe SSD solutions for the market. This reduces deployment and acquisition costs

In addition to servers, the roll-your-own server IT departments have an insatiable appetite for storage. A [recent report](#) by Bereford Capital LLC, made the following observation. "Amazon Web Services (AWS) is the largest public cloud Infrastructure-as-a-Service (IaaS) provider. Gartner estimates AWS has 70% market share. AWS is the IaaS provider for Dropbox, Instagram, Pinterest and of course Netflix... AWS just reported (it) stores 2 trillion objects as of May 2013. Last year AWS reported it stored 1 trillion objects. If we assume that an average size of an object is 1 MB (a typical stored email or letter is 20KB and a photo is 100 KB) then we get 1,863 petabytes of storage equal to 1.8 Exabyte's. So that means AWS added about 1 Exabyte of storage in a year." What isn't stated in these figures is the need to access any one of these objects in an instant. To provide faster access to data, NVM Express removes performance inhibitors built into the legacy AHCI and SCSI protocols for traditional hard drives. What remains is an efficient instruction set providing lower latency and increased IOPs."

Outside of the above main trends, with the acceleration of the Internet of Things (IOT) and a world with a Trillion+ sensors, there is going to be even more demand for intermediate sensor data storage between low cost "dumb" sensor devices (alarms, wearables, health monitoring devices) and the high density big data cloud applications on servers. These intermediate Gateway appliance are likely to have low cost, NVMe standard based SSD drives embedded in them. Emerging architecture of IOT will define the scale of the SSD penetration in this high growth market.

The Emergence of Solid State Drives in Data Centers

There are two ways to achieve this fast access. One is to increase the number of servers. The second is to streamline the server's access to these thousands of petabytes of data.

This addresses the demand resulting from server virtualization, virtual desktop infrastructures (VDI) and high performance databases all creating the need for higher storage I/O performance. The solution these data centers are employing is to introduce solid-state drive (SSD) systems between the high speed rotating memory and the servers' main memory (DRAM). [Dell claims](#) that their SSD "solution can reduce latency up to 90%, in 84% less rack space, for 80% less cost/GB than conventional spinning disk solutions."

Common storage architectures currently being employed in these SSD storage solutions include Bridged PCIe SSD solutions like PCIe based serial-attached-SCSI (SAS), SCSI Over PCIe (SOP), PCIe based SATA and FPGA-based PCIe SSDs. According to Global Industry Analysts, Inc., "these bridged PCIe SSD solutions currently have inherent architectural limitations that make these solutions less cost effective in the longer run. In addition, these solutions have the ability to handle only medium density IOPS. Higher power consumption is another major drawback associated with these solutions."

To overcome this obstacle, the NVM Express ([NVMe](#)) specification was developed allowing a PCIe SSD to connect directly to the PCIe Root Complex with a simple, efficient command set written specifically for the current and next generation of Non-Volatile Memory. For end-users, this provides a scalable interface that unlocks the potential of PCIe-based SSDs, future-proofing storage investments.

The NVM Express -interface provides the following benefits:

- Performance across multiple cores to quickly access critical data.
- An optimized register interface and command set that reduces CPU utilization resulting in higher performance and lower power.
- Scalability with headroom for current and future NVM performance.
- End-to-end Data protection capabilities and support for standard security protocols, such as Trusted Computing Group.
- Lower power consumption resulting in a lower Total Cost of Ownership and carbon footprint.
- Seamless integration into multiple Operating System environments with standard open driver interfaces

The race is now on to develop controller chips that will implement the full capability of this new interface. OEMs are faced with creating a new chip design from scratch, a 9- to 18-month non-recurring engineering effort. However, there is another alternative that can greatly shorten this time to market: license all the IP needed to replace the existing SSD controller with a fully functional PCIe to NVMe solution.

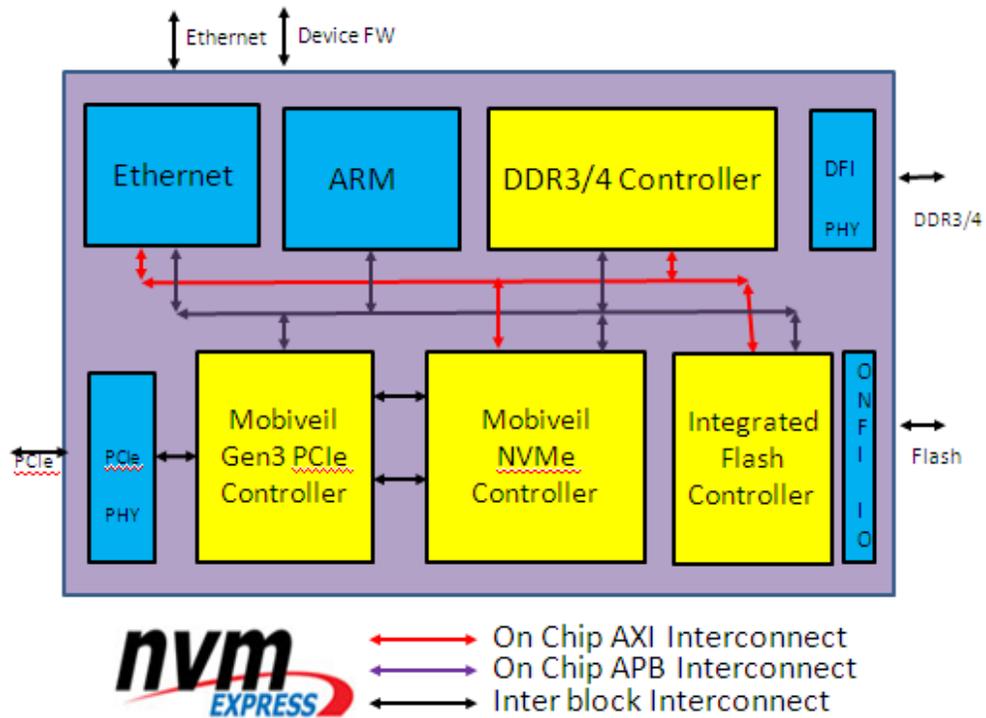
Mobiveil 's Integrated NVM Express based PCIe-SSD solution

Such a solution is available from Mobiveil Inc. the [Universal NVM Express Controller \(UNEX\)](#) shown in the accompanying figure. This complete highly flexible and configurable reference design serves both enterprise and client class solutions. It will unlock the current and future potential of PCIe-based SSDs. The UNEX controller core efficiently supports multi-core architectures thus ensuring thread(s) may run on each core with their own queue and interrupt without any locks required. The controller provides support for end-to-end data protection, security and encryption as well as robust error reporting and management capabilities. The UNEX controller can be used along with Mobiveil's PCI Express controller (GPEX), DDR4/3 and Flash controller (IFC) IPs for a complete NVMe based PCIe SSD implementation.

The IP subsystem is fully compliant to NVM Express 1.1 specification. Thus, it supports configurable number of IO queues and configurable queue depth and the round robin or weighted round robin with urgent priority arbitration mechanism—a powerful feature unique to the NVMe specification. The IP provides a well-defined command interface for local CPU to perform subsystem initialization and to handle all non-hardware accelerated commands. It also provides the following:

- Host memory page size support of 128MB
- Efficient and streamlined command handling
- Supports Fused Operations
- Supports All Optional Admin Commands
- Supports All Optional NVM Commands
- Supports Multi-Path IO and Namespace Sharing capabilities
- Supports Reservations
- Supports multiple name spaces
- Optional AXI interfaces for NVMe implementation in SoC

NVMe IP Subsystem



To explore how Mobiveil’s NVMe IP subsystem can help shorten your time to market, e-mail us at ip@mobiveil.com or visit our website at www.mobiveil.com.