PROPOSED U.2 STORAGE PIN OUT FOR TELEMETRY APPLICATIONS

Chris Budd
SMART High Reliability Solutions
Gilbert, AZ 85233
Chris.Budd@smartH.com

ABSTRACT

The Non-Volatile Memory Express (NVMe) storage interface takes advantage of the internal parallel memory architecture found in many Solid-State Drives (SSDs) to provide high performance bandwidth. While relatively new, NVMe is gaining in popularity, but most vendors do not provide the features needed by telemetry applications. Recognizing that the standard does not define these features, several vendors have collaborated on a standard pin out for the 2.5" U.2 form factor that will provide these features such as write protection, data elimination, encryption, security, and authentication. By following this pin out, both system designers and SSD designers can benefit from this compatibility.

KEYWORDS

NVMe, PCIe, SATA, SCSI, SSD

INTRODUCTION

Industry standards do an excellent job of defining components that are compatible: the designer of each component knows what is needed to interoperate with other components and the consumer of each component is assured of multiple exchangeable sources. However, special markets have special needs; for example, telemetry data storage markets have requirements not found on standard commercial Solid State Drives (SSDs). These requirements include write protection, data elimination, encryption, security, and authentication that are missing from the standard pin-out definitions. Many specifications have unused pins or pins defined for other product families so that several families of products can utilize the same connector; without explicit pin-out definition on these telemetry requirements, each vendor is free to define their own pin-out leading to incompatible solutions. Recently several storage vendors have agreed on a pin-out in the Non-Volatile Memory Express (NVMe) U.2 form factor which includes these features. By following this proposed pin out, designers of telemetry systems can benefit from components that interoperate well and have multiple sources.

1. NVME OVERVIEW

In 2007, several large companies including Intel, who makes SSDs and the NAND inside SSDs, aimed to create a new storage interface to improve the performance of nonvolatile memory-based storage devices. [1] After more than ten years, this NVMe interface has become one of the hottest topics in data storage with many seminars and entire conferences devoted to it.

1.1 NVME OFFERS HIGH PERFORMANCE

One of the main benefits of the new storage standard is higher read/write performance over traditional storage interfaces such as SAS or SATA. An SSD generally has several channels of NAND flash packages with independent address and data lines allowing for parallel access. NVMe takes advantage of this parallel interface by creating a method for a host to issue multiple commands in multiple queues allowing the storage device to determine the best order to complete the commands to minimize the latency and time waiting for data. On top of the parallel command queues, NVMe uses up to 4 PCIe lanes of 8 Gigabits per second (Gb/s) each providing nearly 4 GigaBytes per second (GB/s) of potential bandwidth, rather than one 6 Gb/s interface like SATA providing about 0.5 GB/s.

1.2 NVME GAINS POPULARITY

The first NVMe specification was released in 2011, and then the NVMe organization incorporated in 2014. From 2014 to 2016, the organization nearly doubled in size as it grew from 61 to 115 members. [2] The member companies include several very large consumer storage companies such as Intel, Micron, Samsung, Seagate, Toshiba, and Western Digital. [3] All of these storage companies, and others, offer many different NVMe storage products. Most of these products are focused on the consumer or enterprise markets, and lack the features needed for telemetry applications.

2. PROPOSED U.2 PIN OUT FOR TELEMETRY APPLICATIONS

Recognizing the NVMe specification itself is missing the definitions needed for telemetry applications, several storage vendors that serve the telemetry and industrial markets collaborated in September 2017 on a pin out which incorporates needed features such as write protection, data elimination, encryption, security, and authentication. The pin out was for the U.2 form factor that has the same physical dimensions as existing 2.5” drives, but utilizes a different connector, SFF-8639. [4] Obviously, the exact details of each feature are defined by the SSD vendor, and not every SSD vendor will implement all the features listed here.

2.1 WRITE PROTECT

Write protection (pin S2) allows the user to prevent the drive from accepting writes from the host; this feature can be useful in a couple scenarios: the drive contains host boot information or operating system that never changes during a mission, or after a mission to protect the telemetry data collected from being accidentally overwritten. To help the user identify that the drive is in a write protected state, there is also an LED (pin S3).

2.2 DATA ELIMINATION

The data elimination allows the user to remove data from the drive either to prevent an adversary from acquiring the data or to simply reuse the drive in a different mission. There are three possible ways to eliminate data from the drive. One is a simple cryptographic erase or crypto erase (pin P9) which will erase the cryptographic keys making data retrieval virtually impossible. Another is a standard erase of the NAND flash and possibly could include military erase sequences with multiple overwrites of the NAND flash in an attempt to eliminate any trace of the data; this erase option can use one (pin P2) or two pins (pin P2 and pin S7). A third option is the destruction (pin S4) of the drive such that it will no longer respond as a drive.

2.3 ENCRYPTION, SECURITY, AND AUTHENTICATION

Encryption, security, and authentication come in many forms; four features are proposed. In an attempt to defeat a malicious user from moving an unlocked SSD after is has been authenticated to a different host while keeping power supplied to the SSD, there is malicious unplug detection (pin S1). For those hosts that need to supply their own encryption keys rather than relying on the self-encrypting drive feature of controlling the encryption key, there is the isolated key fill (pins S5 and S6). The proposal includes authentication (pin P1) of the host or the SSD, for example using an Atmel ATSHA204. Finally, a Vbat (pin P7) signal provides low voltage, low current power from the host to the SSD to maintain the encryption key while the main 12V power is off.

2.4 ACTIVITY

NVMe, like previous storage standards, provides an activity signal that the host could use to connect to an LED for a user to see read and write activity to the SSD. This proposal uses the same pin (P11), but also includes the possibility to provide additional activity details on the status of the data elimination, encryption, security, and authentication.

2.5 SUMMARY

Pin Number	**Defense pin name and function when used on U.2**	**Pin IO Direction (With respect to U.2 SSD)**	**Legacy SATA function**
S1	Malicious Unplug Detect (SSD has weak internal pullup): 0= U.2 connector is Engaged. 1= U.2 connector Dis-engaged.	Input (Ground host side)	GND
S2	Write Protect (SSD must have weak Internal pullup): 0 = SSD is hardware write protected. 1 = Not write protected.	Input	SATA RX+ (Input to SSD)
S3	Write Protect LED 0 = Illuminate Write Protect LED (Minimum 5ma sink current) 1 = Write Protect LED off.	Open Drain output Weak internal pullup	SATA RX- (Input to SSD)
S4	Destruct: 0 = No Destruct. 1 = Self-Destruct operation after a SSD vendor defined de-bounce. Signal is ignored if high at power on time and SSD should include an internal pulldown. Actual type of destruct operation is defined by the SSD vendor.	Input	GND
S5	Isolated Key Fill Port: 3.3V RS-232 TX, RS-485+, DS-101+ Protocol is SSD vendor defined.	Output, or BIDIR Differential+	SATA TX- (Output from SSD)
S6	Isolated Key Fill Port: 3.3V RS-232 RX, RS-485-, DS-101- Protocol is SSD vendor defined.	Input or BIDIR Differential-	SATA TX+ (Output from SSD)

S7	Erase Trigger, High True 0 = No secure erase operation. 1 = Trigger a Secure Erase operation, if enabled. If the signal is fully isolated, then pin P2 is the erase return. Signal is ignored if high at power-on and a weak pulldown is suggested. Operating details are SSD vendor defined.	Input	GND
P1	SSD vendor defined. Suggested use: Authentication of Host and/or SSD using an Atmel ATSHA204.	Input/Output	3.3 V
P2	Erase Trigger, Low true or Erase Trigger Return: 0 = Trigger a Secure Erase operation, if enabled. Can also be the return pin for an isolated Erase Trigger using pin S7. 1 = No secure erase operation. Signal is ignored if low at power-on and a weak pullup is suggested. Operating details are SSD vendor defined.	Input	3.3 V
P7	Vbat: Encryption key hold-up voltage used when main 12V is off. Voltage is SSD vendor defined. 1.8 V to 3.3 V. 3V is suggested. Exact operating details are SSD vendor defined.	Input	5 V
P8	SSD vendor reserved	Input	5 V
P9	CryptoErase (SSD must have weak internal pullup): 0 = Trigger a crypto erase 1 = No crypto erase	Input	5V
P11	Activity/Secure Erase Activity (Weak Pullup, Minimum 5ma sink current) Can operate as the standard activity function or by a programmable option, the output can indicate the status of an Erase, Key Purge, or Sanitize operation. Exact operating details are SSD vendor defined.	Open Drain Output	Activity

3. CALL TO ACTION

NVMe has been growing in market share and vendor support for over 10 years; it is not likely to diminish anytime soon. Therefore, as telemetry system designers and SSD designers plan for the next generation of storage using NVMe's U.2 interface, they should consider this pin out proposal. If both system and device designers implement this proposed pin out, then both sides will benefit. System designers will have more choices and more possibilities for second sources, and storage designers will benefit from having compatibility with a wider range of systems and potential customers.

REFERENCES

[1] Intel, "Dell, Intel And Microsoft Join Forces To Increase Adoption Of NAND-Based Flash Memory In PC Platforms," 30 May 2007. [Online]. Available: https://www.intel.com /pressroom/archive/releases/2007/20070530corp.htm. [Accessed 25 May 2018].

[2] A. Huffman, "NVM Express Past, Present, and Future," 9 Aug 2016. [Online]. Available: https://www.flashmemorysummit.com/English/Collaterals/Proceedings/2016/ 20160809_FA11_Huffman.pdf. [Accessed 25 May 2018].

[3] NVM Express, Inc., "About," [Online]. Available: https://nvmexpress.org/about/. [Accessed 25 May 2018].

[4] SSD Form Factor Work Group, "SSD_Form_Factor_Version1_a.pdf," Storage Networking Industry Association (SNIA), 12 December 2012. [Online]. Available: http://www.ssdformfactor.org/docs/SSD_Form_Factor_Version1_a.pdf. [Accessed 8 June 2018].

NOMENCLATURE

ATA or Advanced Technology Attachment: An interface standard from T13 (http://www.t13.org/) for storage devices using a parallel bus (PATA), or serial bus (SATA).

Flash: A non-volatile memory device using an array of transistors each with a floating gate to store a charge.

GB or Gigabyte: 10^9 bytes.

HDD or Hard-disk drive: Traditional mass storage device using a rotating, magnetic platter.

MB or Megabyte: 10^6 bytes.

NAND: A high-density flash device usually with defect blocks marked by the factory; read and write operations must be done at a page level (several kilobytes), and erases must be done at an erase block level consisting of several hundred pages.

NVMe or Non-Volatile Memory Express: A high-speed storage interface commonly using PCIe.

PCIe or Peripheral Component Interconnect Express: A high-speed bus for computer peripherals.

SATA or Serial ATA: A storage bus interface where the data is transferred serially rather than through parallel data wires as in previous ATA devices.

SCSI or Small Computer System Interface: An interface standard from T10 (http://www.t10.org/) for storage devices typically using a parallel bus, or serial bus (SAS).

SSD or Solid-State Drive: A mass storage device typically using the same form factors as traditional hard-disk drives, but without the moving parts. An SSD typically stores data in SLC or MLC NAND flash.