NVMe-IP Introduction for Xilinx

Ver1.8E

Optimal Solution for Data Recording Application!

Agenda

- NVMe SSD Overview
  - SSD Trends
  - Merit of NVMe SSD for embedded system
- NVMe-IP Introduction
  - Summary
  - High Performance and Compact Size
  - Easy User Interface
  - Rich Features
  - Development Environment/Reference Design
- Application

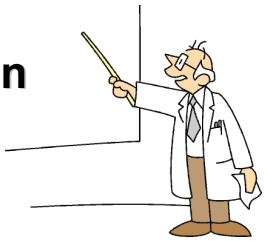

SSD Trends 1

- SATA interface is now performance bottle neck
  - SSD Read/Write speed is limited to 600MB/sec SATA bandwidth
- Move to PCI Express for faster speed
  - PCIe GEN3 x4lane can provide 4GB/sec transfer speed
- M.2 and BGA package suitable for compact application
  - M.2: Wid=22mm, Len=20/42/80/120mm DIMM-like small outline
  - BGA: 20mm x 16mm x 1.5mm, 1gram package

Current 2.5" SATA SSD

Latest M.2 type PCIe SSD

BGA type SSD

SSD Trends 2

- Host Controller Standard moves from AHCI to NVMe
  - Latest standard to extract maximum performance of SSD
  - Extended Queue size, 65536 concurrent command process
  - Most OS provides NVMe driver

FMS2015 "Annual Update on Interfaces" presentation

Merit of NVMe SSD for Embedded System 1

- High Bandwidth: 3.5GB/s for Read, 2.5GB/s for Write
- Cost effective: Cost difference from SATA SSD is small

$197.99 for 500GB

Samsung

Samsung 970 EVO 500GB - NVMe PCIe M.2 2280 SSD (MZ-V7E500BW)

161 customer reviews | 55 answered questions

Amazon's Choice for "mz-v7e500bw"

Price: $197.99 + $21.58 Shipping & Import Fees

Deposit to Japan Details

Capacity: 500 GB

1 TB	2 TB	250 GB	500 GB
$394.99	$794.99	$106.82	$197.99

(Amazon.com 18-Jul-2018)

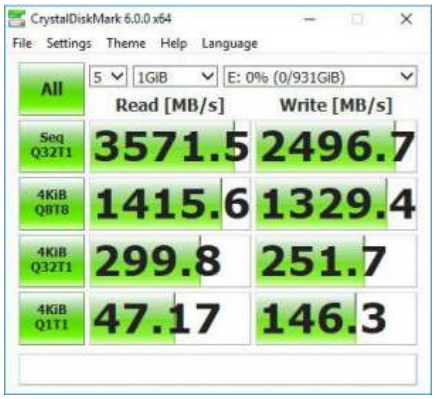

https://www.techadvisor.co.uk/review/ssd-hard-drives/samsung-970-evo-review-3677171/

Cost and Performance of M.2 NVMe SSD (Samsung 970 EVO 500GB)

Merit of PCIe SSD for Embedded System 2

- Various form factor
  - HHHL(Half-Height,Half-Length) general PCIe board
  - M.2 cost saving module
  - U.2 (SFF-8639) of 2.5” drive compatible size
  - BGA package for direct mount on PCB

Merit

Big Capacity

Small, Extractable

Hot Swap

Mount on PCB

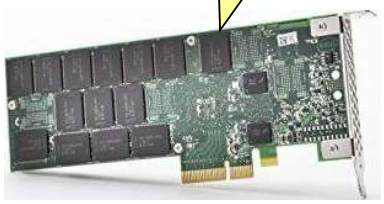

HHHL PCIe board

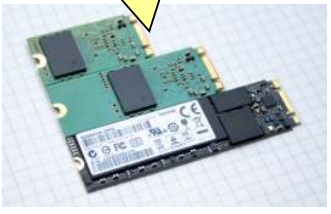

M.2 module (length=42/60/80mm)

U.2 package

BGA pacakge

What's NVMe-IP

- What's NVMe-IP? -> Directly connect NVMe SSD with FPGA
- How to use? -> Just connect with user logic. No need for CPU, its F/W, External Memory
- Application -> Best for ultra high speed data recording system
- User Merit? -> Can develop Storage Application in short period

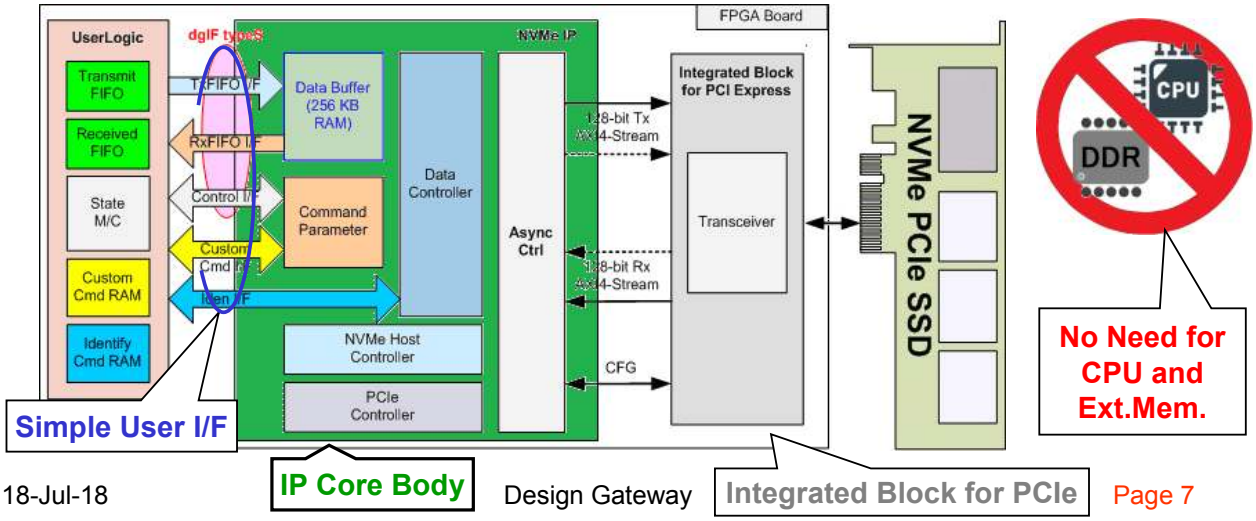

NVMe-IP Merit

1. **High Performance and Compact size**
   - Write=2148MB/s, Read=3252MB/s (measured by KCU105)
   - Support PCIe GEN3 (Operation confirmed on Ultrascale, Ultrascale+)
   - IP-Core Size=4170CLBRegs, Memory=59BRAMTile
2. **Interface: Simple and easy connection**
   - Direct connection to Xilinx Integrated Block for PCIe
   - User I/F control is parameter with pulse, data is simple FIFO
   - Use BRAM for data buffer (external DDR memory not required)
3. **Rich Features: Custom command in addition to Read/Write**
   - Supports SMART/FLUSH/Shutdown custom command
   - Supports both legacy 512byte and 4Kbyte Sector format
4. **Environment: Full reference design project**
   - Full Vivado project with real board operation in the package

NVMe-IP Merit1: Performance

- **Automatic PCIe SSD access by pure hard-wired logic**
  - **Intelligent state machine for complete read/write command execution**
  - **Minimum over head and best performance by synchronized circuit**

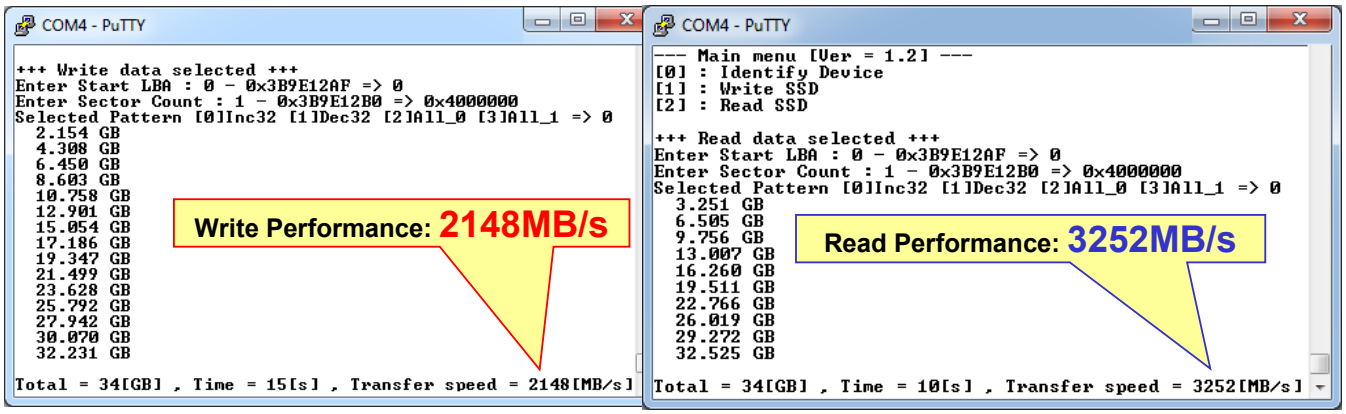

Performance Evaluation Result (KCU105) **(SSD: Samsung MZ-V6P512BW)**

NVMe-IP Merit1: Compact Size

- **Optimized size with minimum resource consumption**
  - **Implements dedicated and optimized logic for NVMe SSD control**
- **Block RAM for data buffer**
  - **Internal block memory can minimize access overhead**

Example Implementation Statics for 7-Series device (PCIe Gen2/PCIe Gen3)

Family	Example Device	Fmax (MHz)	Slice Regs	Slice LUTs	Slices	BRAMTile[1]	Design Tools
Virtex-7	XC7VX690TFFG1761-2	300	4169	2828	1423	59	Vivado2017.4
Virtex-7	XC7VX485TFFG1761-2	300	4159	3465	1446	59	Vivado2017.4
Zynq-7000	XC7Z045FFG900-2	300	4159	3461	1506	59	Vivado2017.4
Kintex-7	XC7K325TFFG900-2	300	4159	3463	1556	59	Vivado2017.4

Example Implementation Statics for Ultrascale device (PCIe Gen3)

Family	Example Device	Fmax (MHz)	CLB Regs	CLB LUTs	CLB	BRAMTile[1]	Design Tools
Kintex-Ultrascale	XCKU040FFVA1156-2E	400	4170	2724	772	59	Vivado2017.4
Zynq-Ultrascale+	XCZU7EV-FFVC1156-2E	400	4170	2670	790	59	Vivado2017.4
Virtex-Ultrascale+	XCVU9P-FLGA2104-2L	400	4170	2675	761	59	Vivado2017.4

NVMe-IP Core standalone resource usage

NVMe-IP Merit2: Command I/F

- **Simple User I/F**
  - Set Command/Address/Length
  - Issue UserReq pulse
- **Full Automatic control for SSD access**
  - User only can wait UserBusy negation

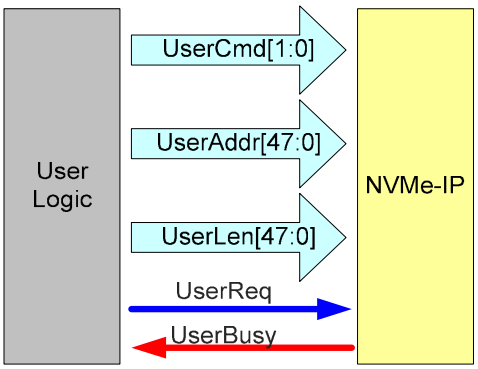

Basic Command I/F Signals

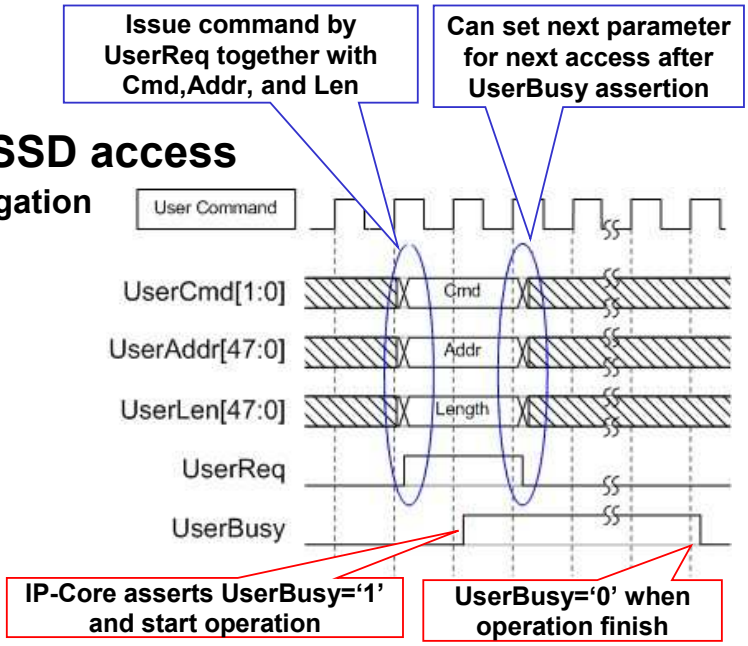

Command I/F waveform

NVMe-IP Merit2: Data I/F

- **Simple 128bit FIFO for each of read and write**
  - General FIFO of standard Xilinx LogiCORE library
  - Data buffer using 256KByte BRAM in NVMe-IP

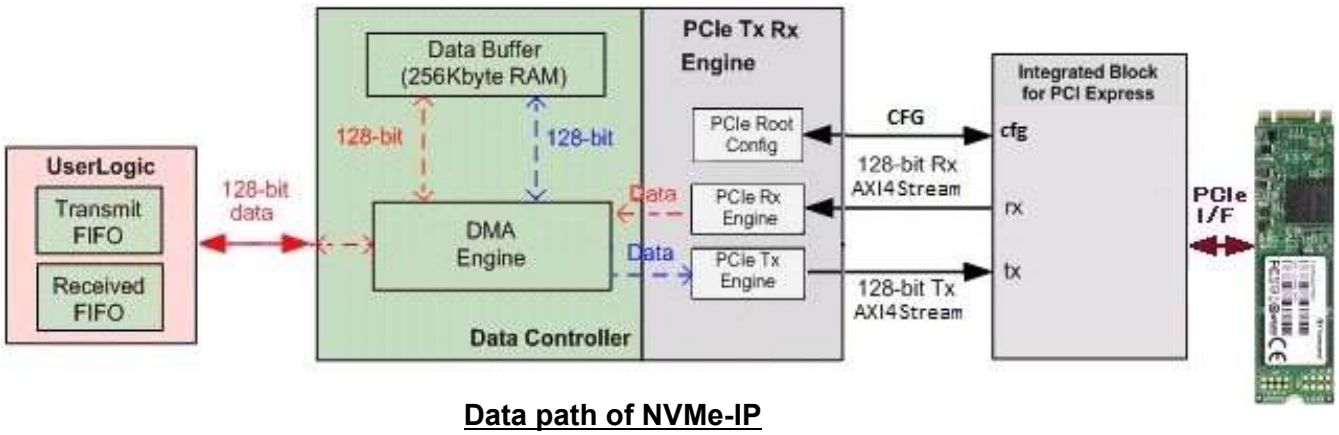

Data path of NVMe-IP

NVMe-IP Merit3: Rich Features

- SMART command for SSD health condition check
  - Can monitor internal temperature, total write size, etc.
- FLUSH command to force cache flush operation
  - User can adjust trade-off between performance and data evacuation
- Safe Shutdown before SSD power down
  - IP-core executes safe shutdown by user request
- Supports both 512bytes and 4Kbytes sector format
  - IP-core automatically selects sector format via Identify command

```
<< SMART Log Information >>
Temperature                    : 32 Degree Celsius
Total Data Read                : 47469 GB
Total Data Written             : 65373 GB
Power On Cycles                : 3991 Times
Power On Hours                 : 79 Hours
Unsafe Shutdowns               : 220 Times
```

SMART command result example

NVMe-IP Merit4: Environment

- Real operation check with Xilinx evaluation board
- Free bit-file for evaluation before IP-core purchase

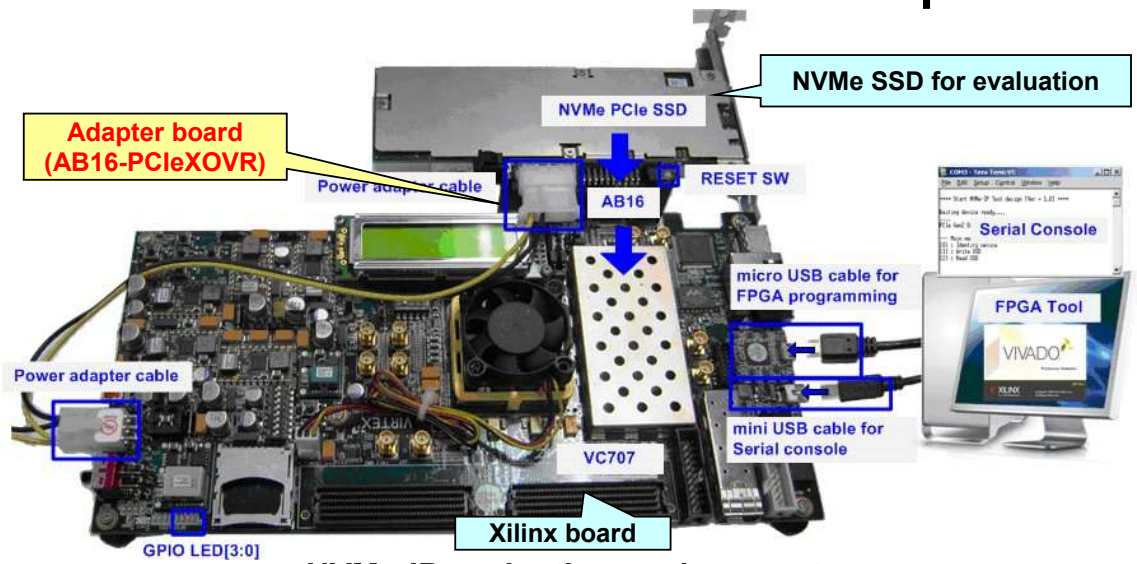

NVMe-IP evaluation environment

NVMe-IP Merit4: Development Tool#1

- PCIe Adapter board for evaluation (Part#: AB16-PCIeXOVR)
  - Connect FPGA board to PCIe socket on component side
  - Connect PCIe SSD to PCIe socket on solder side
  - SSD R/W access via adapter board from NVMe-IP in FPGA

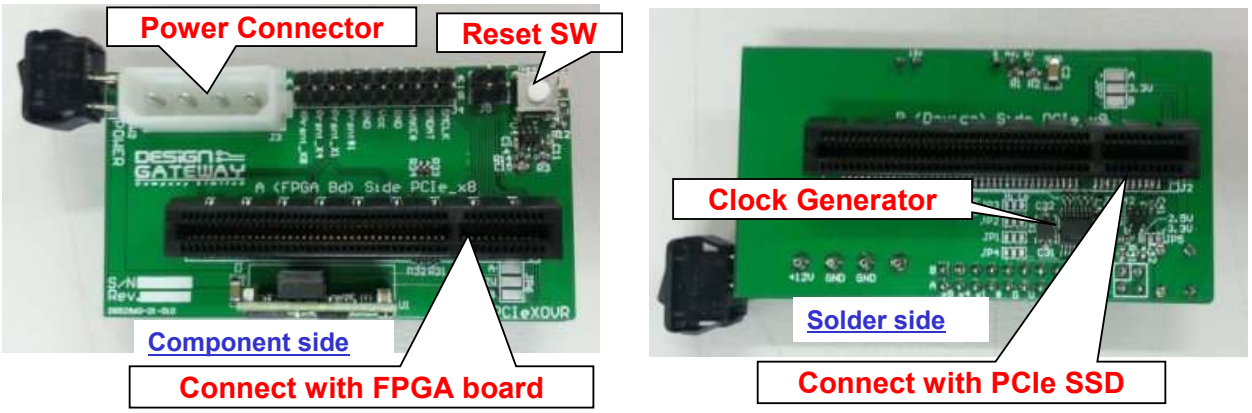

PCIe adapter for NVMe-IP evaluation (AB16-PCIeXOVR)

NVMe-IP Merit4: Development Tool#2

- FMC Adapter board for evaluation (Part#: AB17-M2FMC)
  - Two M.2 sockets on component side
  - FMC HPC connector for FPGA connection on solder side
  - High capacity power supply (max 5A for 3.3V output per one SSD)

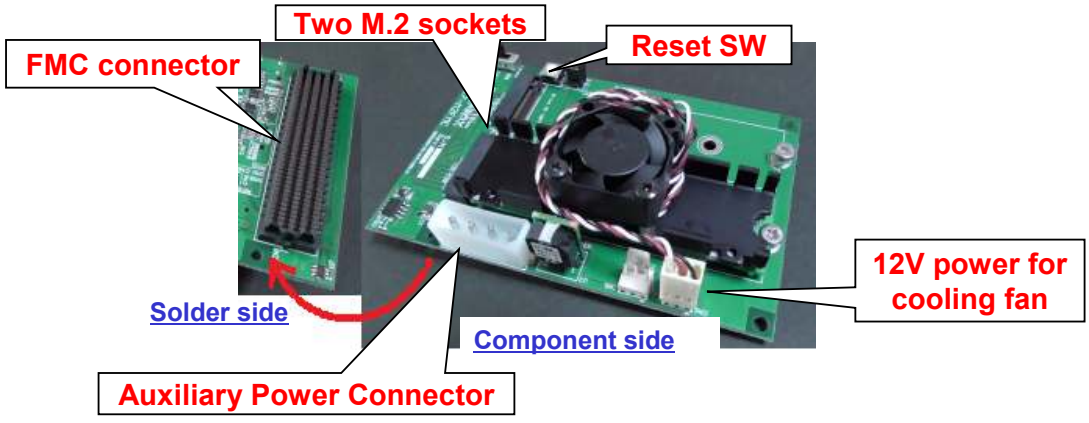

FMC adapter for NVMe-IP evaluation (AB17-M2FMC)

NVMe-IP Merit4: Reference Design

- Vivado project is attached with NVMe-IP deliverables
- Full source code (VHDL) except IP core
- Can save user system development duration
  - Confirm real board operation by original reference design.
  - Then modify a little to approach final user product.
  - Check real operation in each modification step.

Short-term development is possible without big turn back

NVMe-IP Application Example 1

- Space-Saving FPGA data logging system
  - Latest FPGA+M.2 SSD

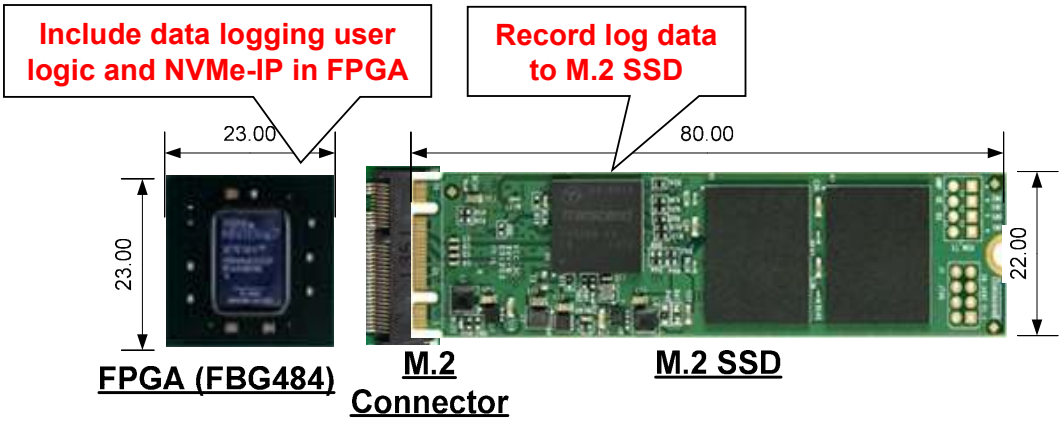

System area image by FBG484 FPGA and M.2 SSD (unit: mm)

NVMe-IP Application Example 2

- **Recording and Analysis system on Linux**
  - Mount Linux and user analysis application on SoC/MPSoC device
  - Very high-speed data recording to SSD via NVMe-IP core
  - Data read from SSD via device driver and analyze by user application

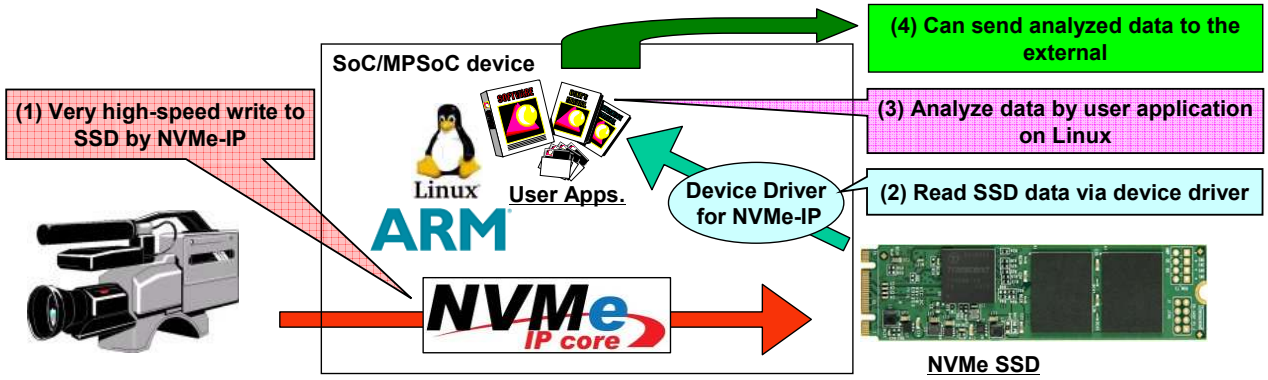

Recording and Analysis sytem on Linux (device driver and reference design available)

NVMe-IP Application Example 3

- **Ultra High-Speed Recorder**
  - Double write speed with multiple SSDs RAID0 configuration
  - Provide RAID0 reference design with 2 NVMe SSDs

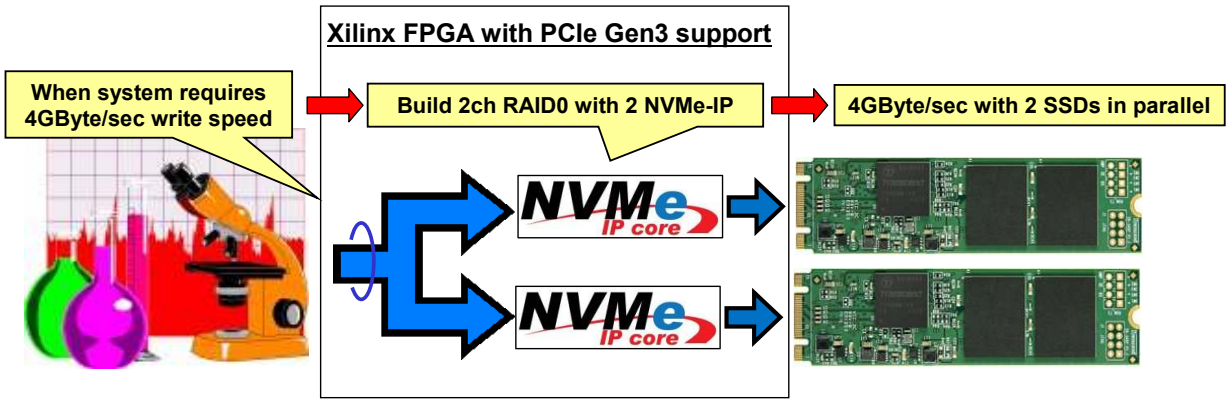

NVMe RAID system supporting 4GByte/sec recording rate

For more detail

- **Detailed technical information available on the web site.**
  - http://www.dgway.com/NVMe-IP_X_E.html
- **Contact**
  - **Design Gateway Co,. Ltd.**
  - **sales@design-gateway.com**
  - **FAX: +66-2-261-2290**

Revision History

Rev.	Date	Description
1.0E	10-Jun-16	English Version first release
1.1E	21-Jun-16	Support Kintex-Ultrascale
1.2E	25-Aug-16	Modify page17 because only one x16 DDR4 device can keep NVMe SSD performance
1.3E	12-Sep-16	Support Zynq-7000 and Kintex-7
1.4E	8-Nov-16	Support PCIe GEN3 on Virtex-7
1.5E	21-Dec-16	NVMe-IP core improvement by removing external DDR chip for data buffer
1.6E	6-Jun-17	Performance improved by internal PCIe bridge in NVMe-IP core
1.7E	2-Nov-17	Added Linux driver application and 2ch RAID0 reference design
1.8E	18-Jul-18	Added 4KB sector format, SMART/FLUSH/Shutdown command support