NVMe-IP Introduction for Xilinx

Ver2.3E

The Very Best Solution for Data Recording Application!

Agenda

- NVMe-IP Introduction
  - Summary, Lineup, Merit
  - High Performance and Compact Size
  - Easy User Interface
  - Rich Features
  - Development Environment/Reference Design
- Optional product (exFAT-IP core)
- Application

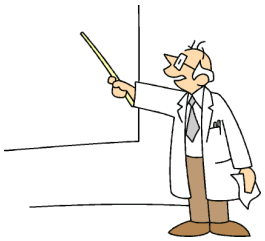

What's NVMe-IP

- What's NVMe-IP? -> Directly connect NVMe SSD with FPGA
- Advantage -> No need for CPU, its F/W, External Memory
  Supports latest PCIE Gen4 protocol
- Application -> Best for ultra high speed data recording system
- User Merit? -> Can develop Storage Application in short period

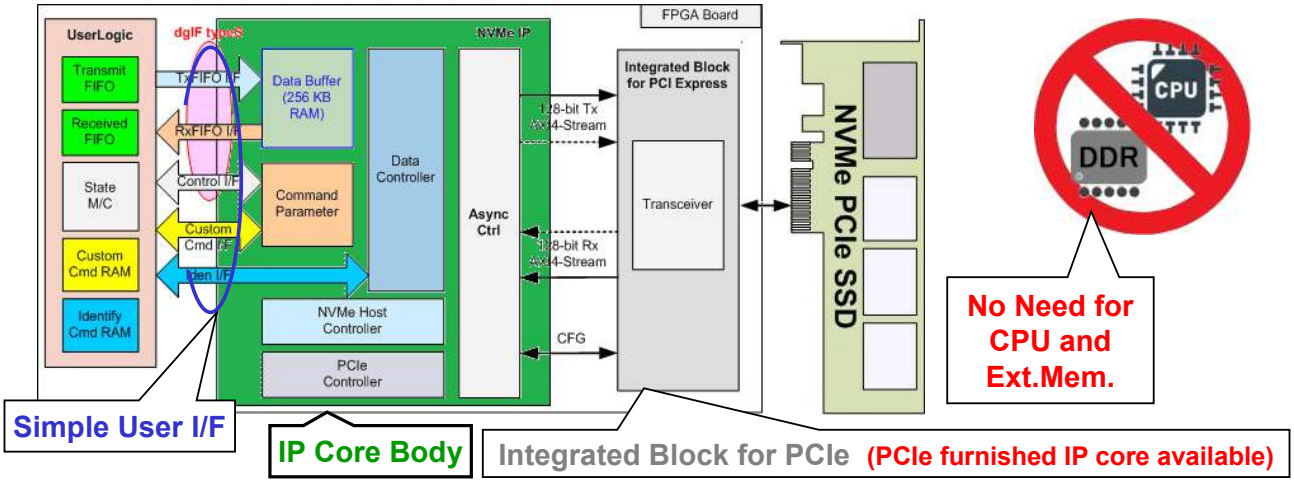

NVMe-IP Lineup

- Multiple lineup for various functions
- PCIe Soft-IP furnished version (Gen4/Gen3) available
- raNVMe-IP suitable for random access application

Core type	Description
Standard NVMe-IP core	Standard core using PCIe Integrated Block in FPGA
NVMeG4-IP core	PCIe Soft-IP furnished, 4-Lane PCIe Gen4
NVMeG3-IP core	PCIe Soft-IP furnished, 4-Lane PCIe Gen3
NVMeSW-IPcore	Multiple SSD connection via external PCIe switch
raNVMe-IP core NEW!!	Supports random read or write access

NVMe-IP core lineup

(Ask DesignGateway for more detail of NVMeSW-IP core.)

PCIe Soft-IP furnished IP core

- **NVMeG4-IP core / NVMeG3-IP core**
  - Can operate without PCIe Integrated Block
  - Includes data link layer and connect with transceiver by PCIe Gen4/3
  - More SSD connection regardless of PCIe Integrated Block count.

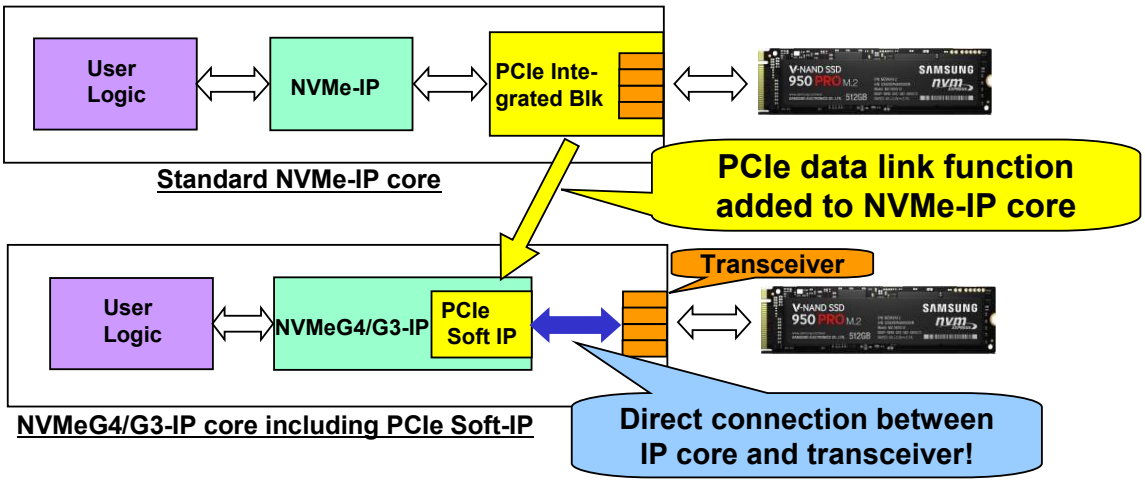

PCIe Soft-IP furnished IP core (Cont'd)

PCIe protocol	Product Number	Target device family	Supported transceiver	Evaluation env.
Gen4 4Lane	NVMeG4-IP-VUP-GTY	Virtex-UltraScale+	GTY	VCU118
	NVMeG4-IP-KUP-GTY	Kintex-UltraScale+	GTY	KCU116
	NVMeG4-IP-ZUP-GTH	Zynq-UltraScale+	GTH	ZCU102/106
Gen3 4Lane	NVMeG3-IP-VUP-GTY	Virtex-UltraScale+	GTY	VCU118
	NVMeG3-IP-KUP-GTY	Kintex-UltraScale+	GTY	KCU116
	NVMeG3-IP-ZUP-GTH	Zynq-UltraScale+	GTH	ZCU102/106
	NVMeG3-IP-KU-GTH	Kintex-UltraScale	GTH	KCU105

NVMeG4-IP/NVMeG3-IP core lineup

Supports all UltraScale+ family and some UltraScale family

Evaluation environment ready for all IP core products

raNVMe-IP for random access

- User can select either Write or Read operation
- Executes 32 commands at maximum concurrently with different (random) address.
- Write or read data per one command is fixed to 4KBytes.

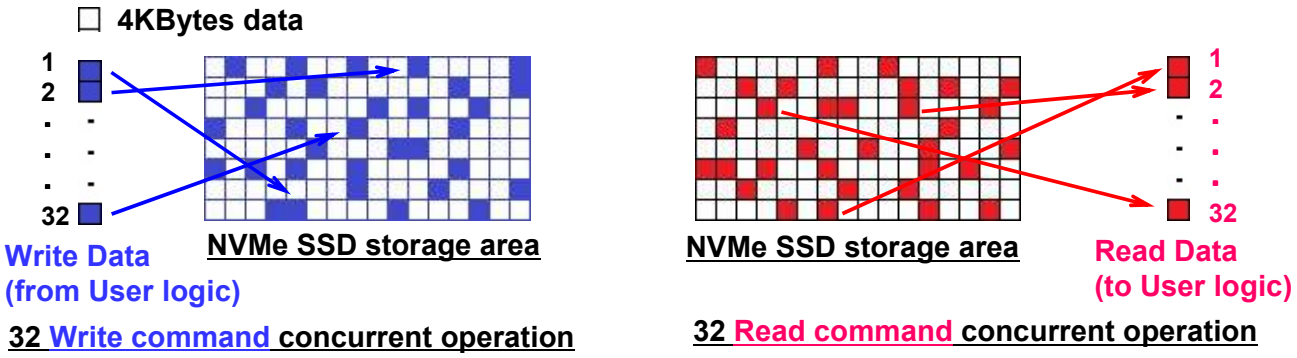

raNVMe-IP concurrent command operation image

NVMe-IP Merit

1. High Performance and Compact size
   - Write=4288MB/s, Read=4670MB/s (measured by VCU118)
   - Support PCIe GEN4 (Operation confirmed on Ultrascale+)
   - IP-Core Size=4170CLBRegs, Memory=59BRAMTile (standard core)
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

Merit1: Performance (PCIe Soft-IP furnished)

- PCIe Gen4 supported core (NVMeG4-IP core) real speed
  - Unprecedented Write/Read performance!

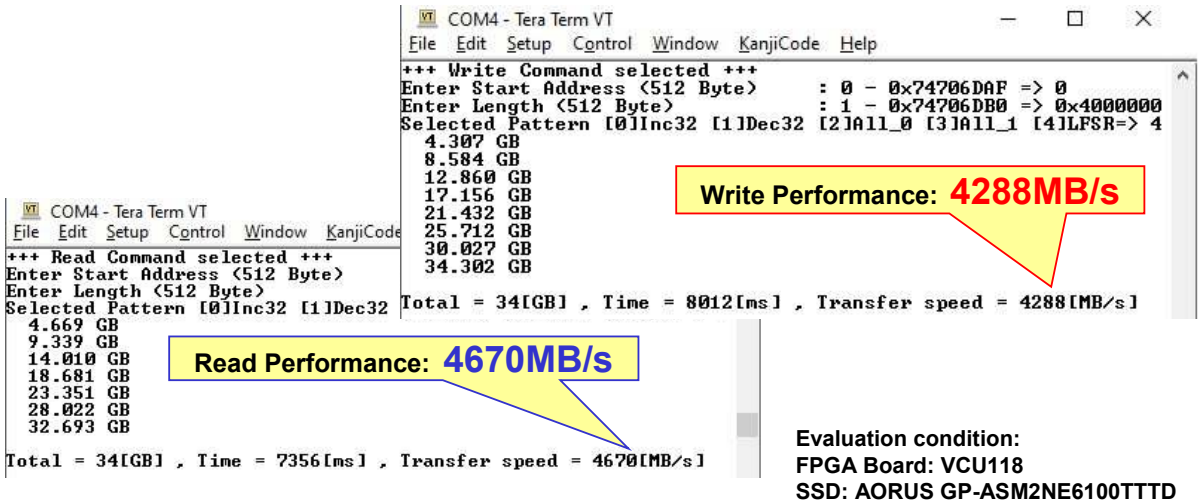

Evaluation condition:
FPGA Board: VCU118
SSD: AORUS GP-ASM2NE6100TTTD
Write test: 256KB data buffer size
Read test: 512KB data buffer size

Performance Evaluation Result of NVMeG4-IP core

Merit1: Performance (Standard core)

- 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

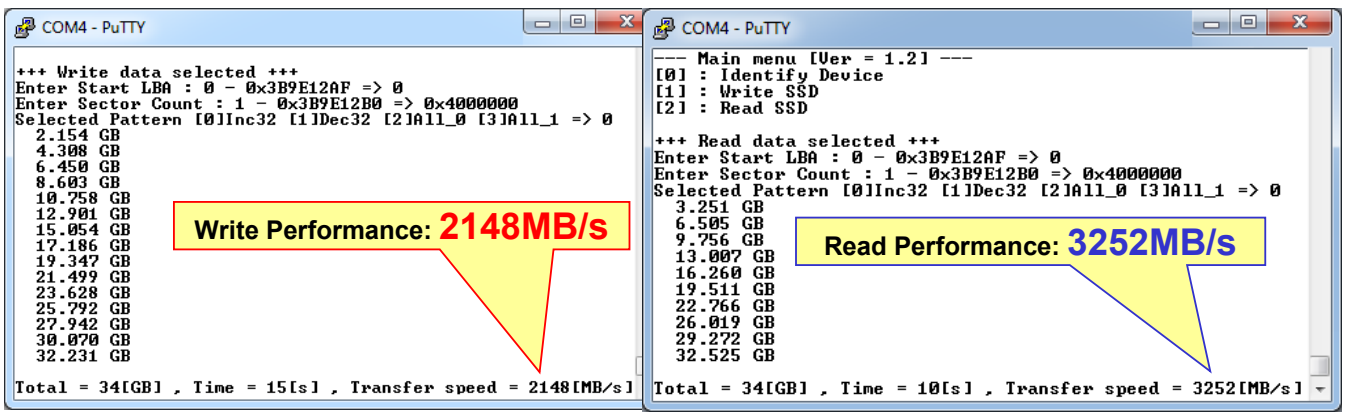

Performance Evaluation Result (KCU105) **(SSD: Samsung MZ-V6P512BW)**

Merit1: Compact Size (PCIe Soft-IP furnished)

PCIe protocol	IP-Core	Example Device	Fmax (MHz)	CLB Regs	CLB LUTs	CLB	BRAM Tile	UR AM
Gen4 4Lane	NVMeG4-IP-VUP-GTY	XCVU9P-FLGA2104-2LE	300	19215	21958	4465	12	8
	NVMeG4-IP-KUP-GTY	XCKU5P-FFVB676-2E	300	19214	21960	4382	12	8
	NVMeG4-IP-ZUP-GTH	XCZU7EV-FFVC1156-2E	300	19213	21961	4374	12	8
Gen3 4Lane	NVMeG3-IP-VUP-GTY	XCVU9P-FLGA2104-2LE	300	11107	12027	2546	70	
	NVMeG3-IP-KUP-GTY	XCKU5P-FFVB676-2E	300	11107	12024	2461	70	
	NVMeG3-IP-ZUP-GTH	XCZU7EV-FFVC1156-2E	300	11147	12029	2779	70	
	NVMeG3-IP-KU-GTH	XCKU040FFVA1156-2E	300	11168	12047	2518	70	

NVMeG4/G3-IP core resource usage example

- **Limitation point of NVMeG4/G3-IP core**
  - **PCIe Gen4/Gen3 only, not support other speed (Gen1/2/4)**
  - **4-Lane only, not support other lane count (1/2/8/16)**
    (Ask other lane count as core customization)

Merit1: Compact Size (Standard core)

- **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

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

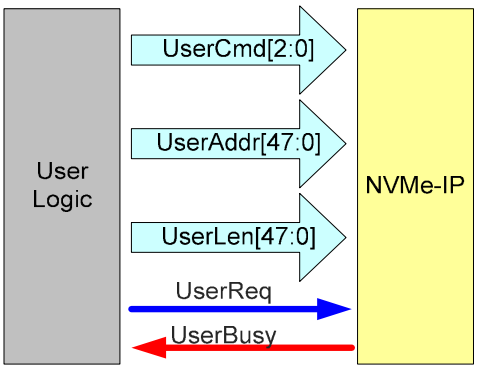

Basic Command I/F Signals

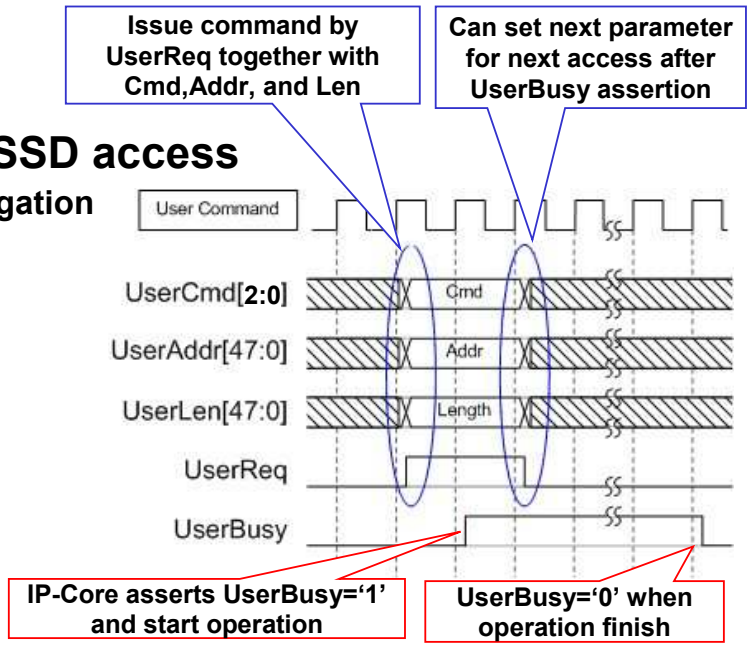

Command I/F waveform

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

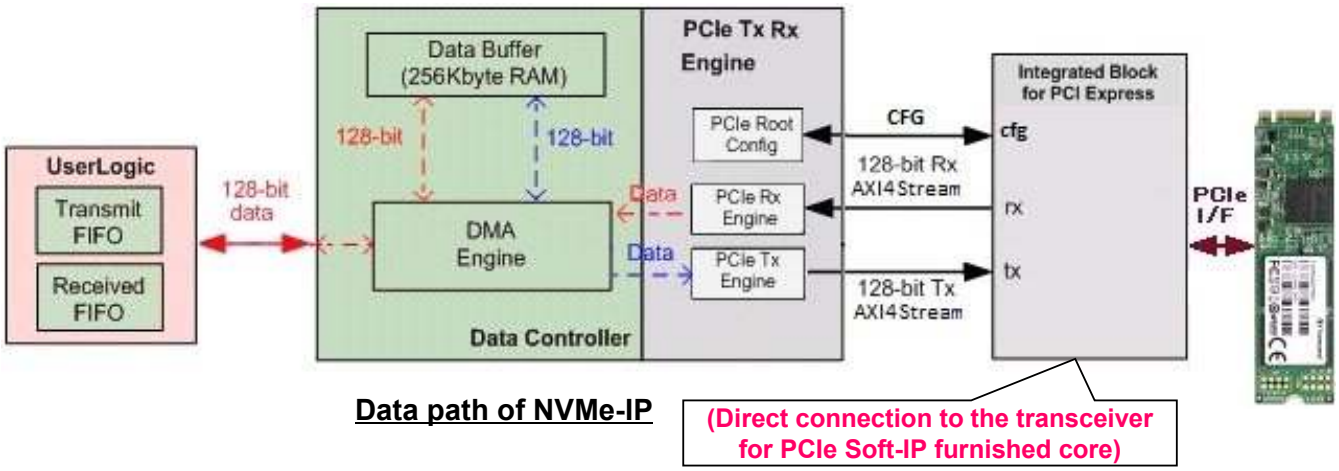

Data path of NVMe-IP

(Direct connection to the transceiver for PCIe Soft-IP furnished core)

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

Merit4: Environment

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

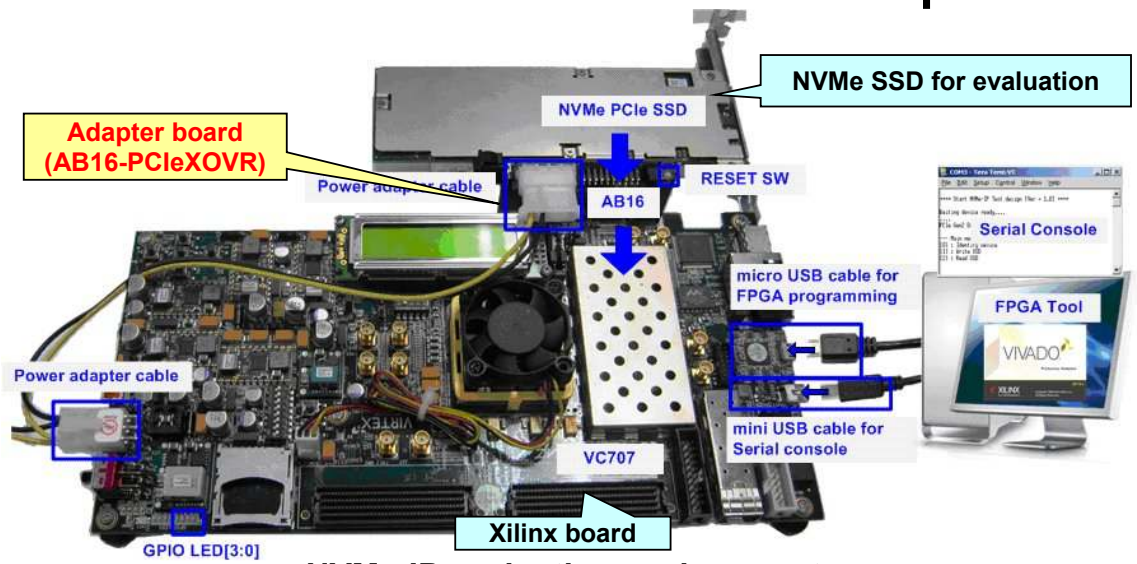

NVMe-IP evaluation environment

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

For 16-lane, use AB18-PCIeX16

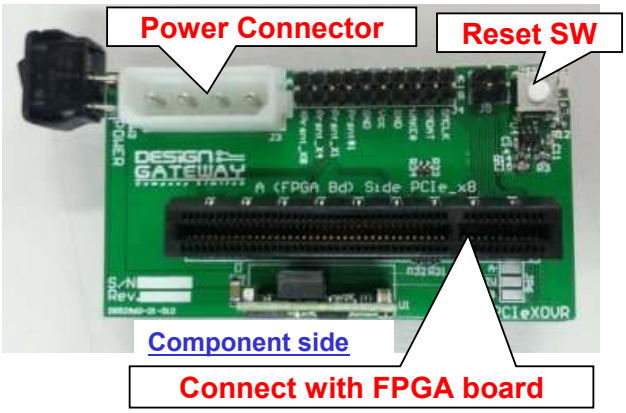

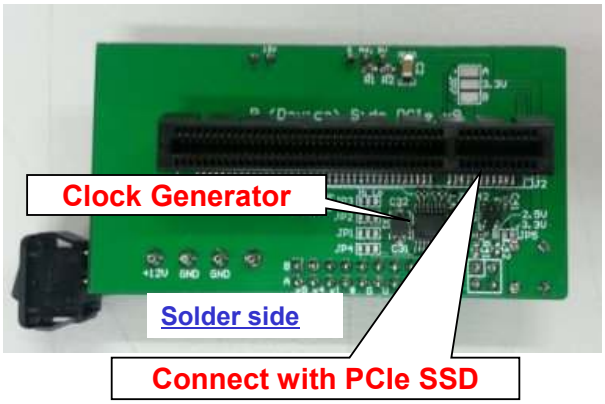

PCIe adapter for NVMe-IP evaluation (AB16-PCIeXOVR)

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)

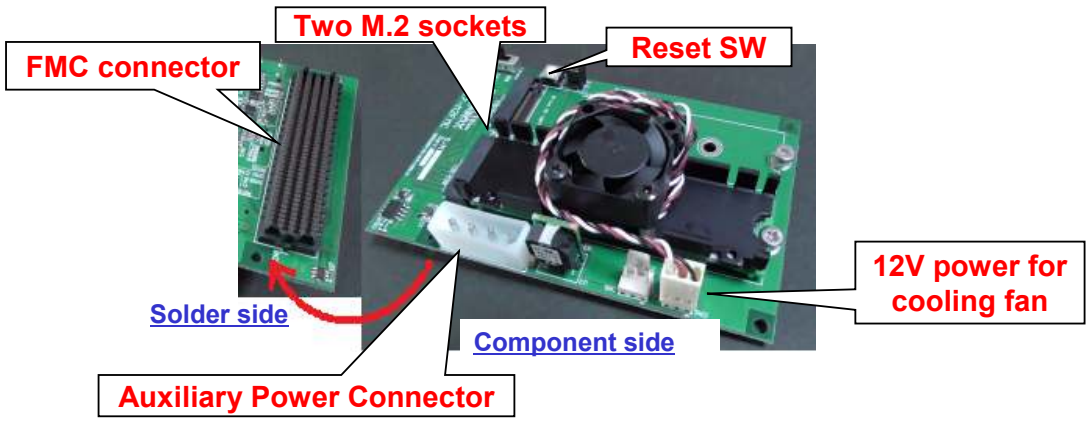

FMC adapter for NVMe-IP evaluation (AB17-M2FMC)

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

Optional product: exFAT-IP Core Introduction

- Optional products for NVMe-IP core
  - Supports data recording with exFAT file format
- PC can directly access to recorded data as a file
  - FPGA writes data to device, reconnect with PC, then PC can read data

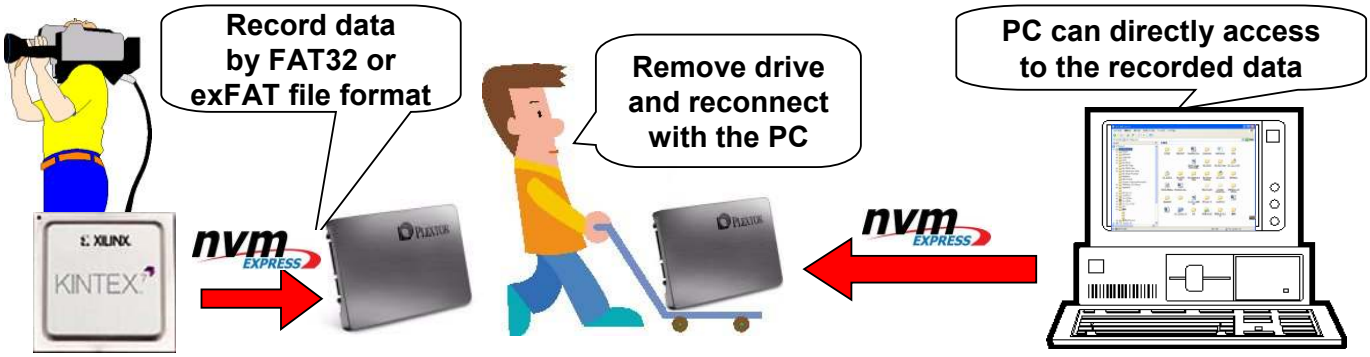

PC can directly read recorded data as a file

Optional product: exFAT-IP (Cont'd 1)

- **Feature description**
  - Executes drive format and data write to file by pure hardwired logic.
  - IP core automatically generates file name.
  - User logic sends file data via FIFO interface.
- **Limitation**
  - Drive must be formatted by the IP core, not by the PC.
  - Files other than those generated by the IP core cannot be written to the drive.
  - File size is determined at format execution and cannot be changed.

Optional product: exFAT-IP (Cont'd 2)

- **Reference design for real operation available**
  - Executes test file generation via serial console.
  - User can confirm file read compatibility by drive re-plug to the PC.

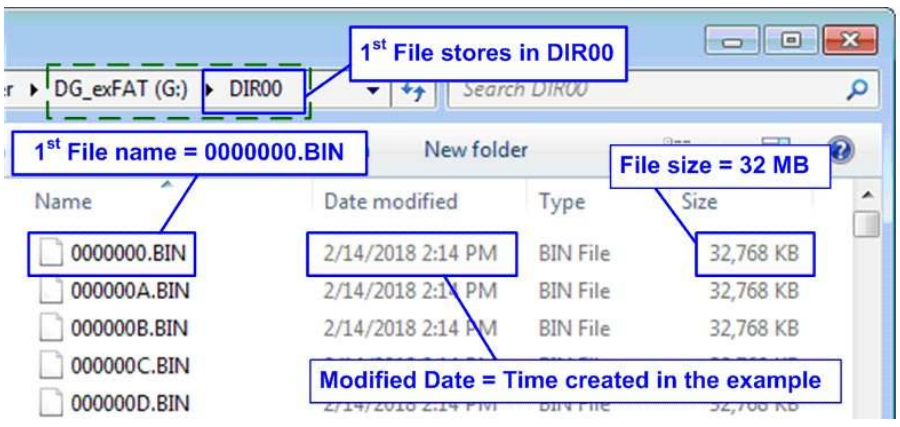

Generate test file, reconnect with PC, and can check file read compatibility

NVMe-IP Application Example 1

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

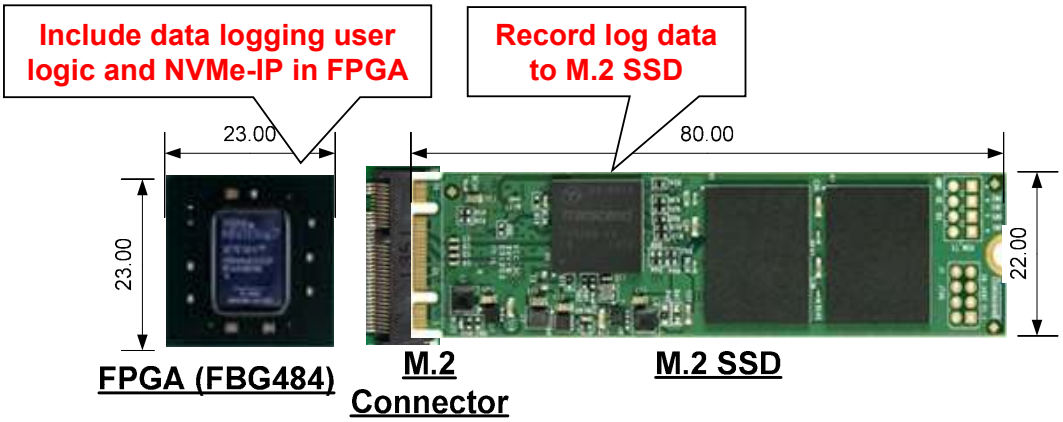

System area image by FBG484 FPGA and M.2 SSD (unit: mm)

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**

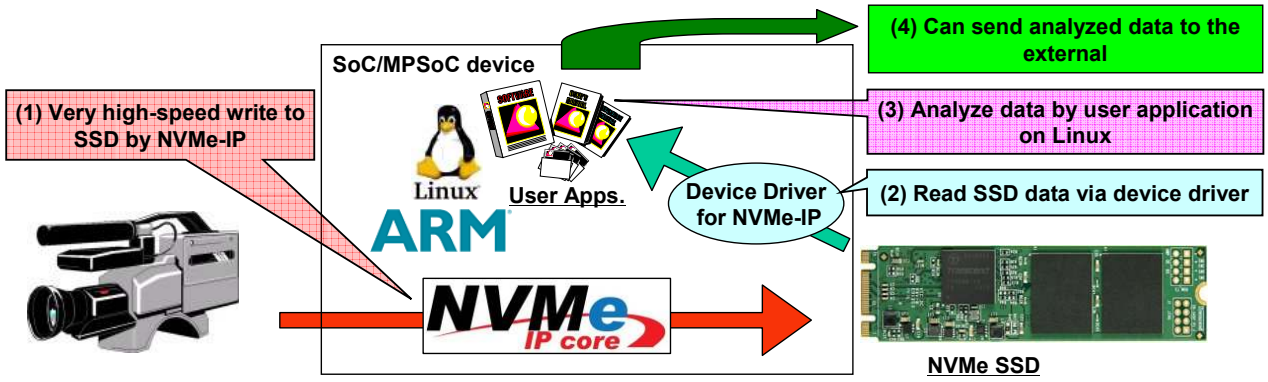

Recording and Analysis sytem on Linux (device driver and reference design available)

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

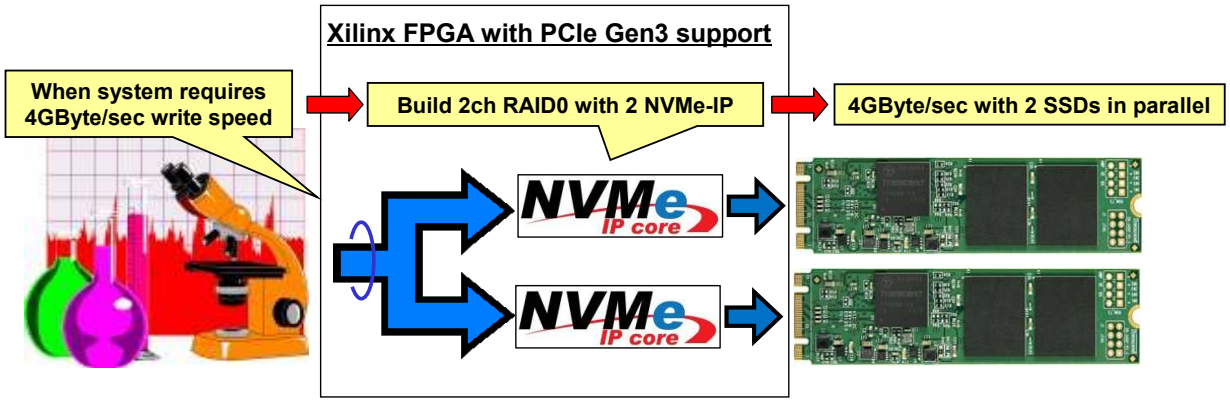

NVMe RAID system supporting 4GByte/sec recording rate

NVMe-IP Application Example 4

- **Super multi-channel SSD Array by NVMeG3-IP**

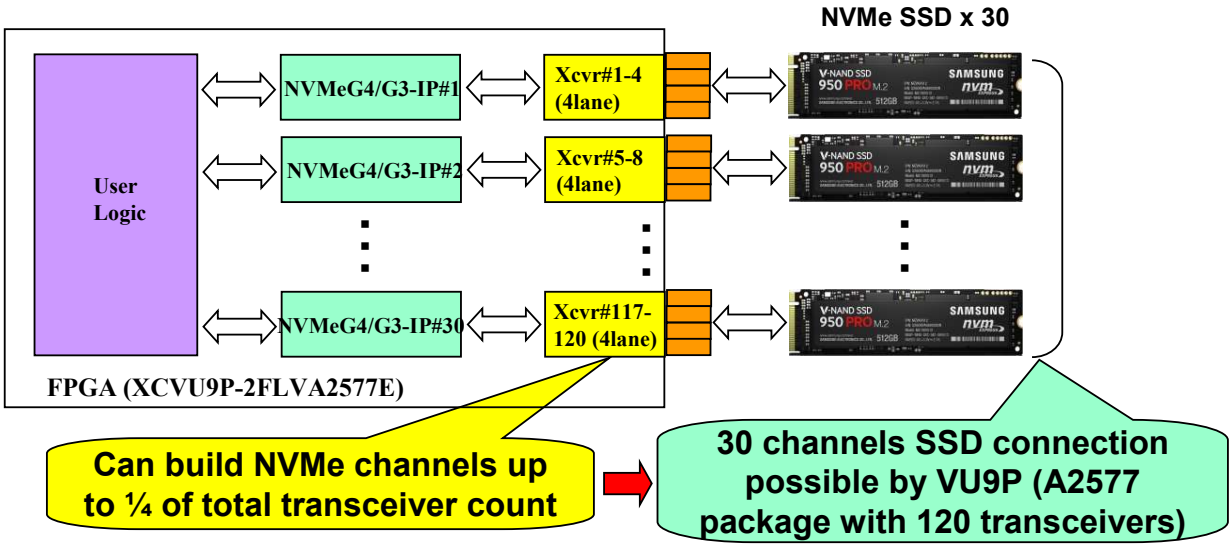

30 channels M.2 SSD Array system using NVMeG4/G3-IP core

For more detail

- **Detailed technical information available on the web site.**
  - https://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
1.9E	9-Jan-19	Add FAT32-IP/exFAT-IP for NVMe-IP optional products
2.0E	24-Sep-19	Add new product of NVMeG3-IP that includes PCIe Soft IP core inside
2.1E	1-Feb-20	Add new product of NVMeG4-IP that includes PCIe Gen4 Soft IP core inside
2.2E	27-Aug-20	Add new product of raNVMe-IP for random access application
2.3E	1-Dec-20	Updated NVMeG4-IP/NVMeG3-IP information