NVMe-IP Introduction for Intel

Ver2.3EA

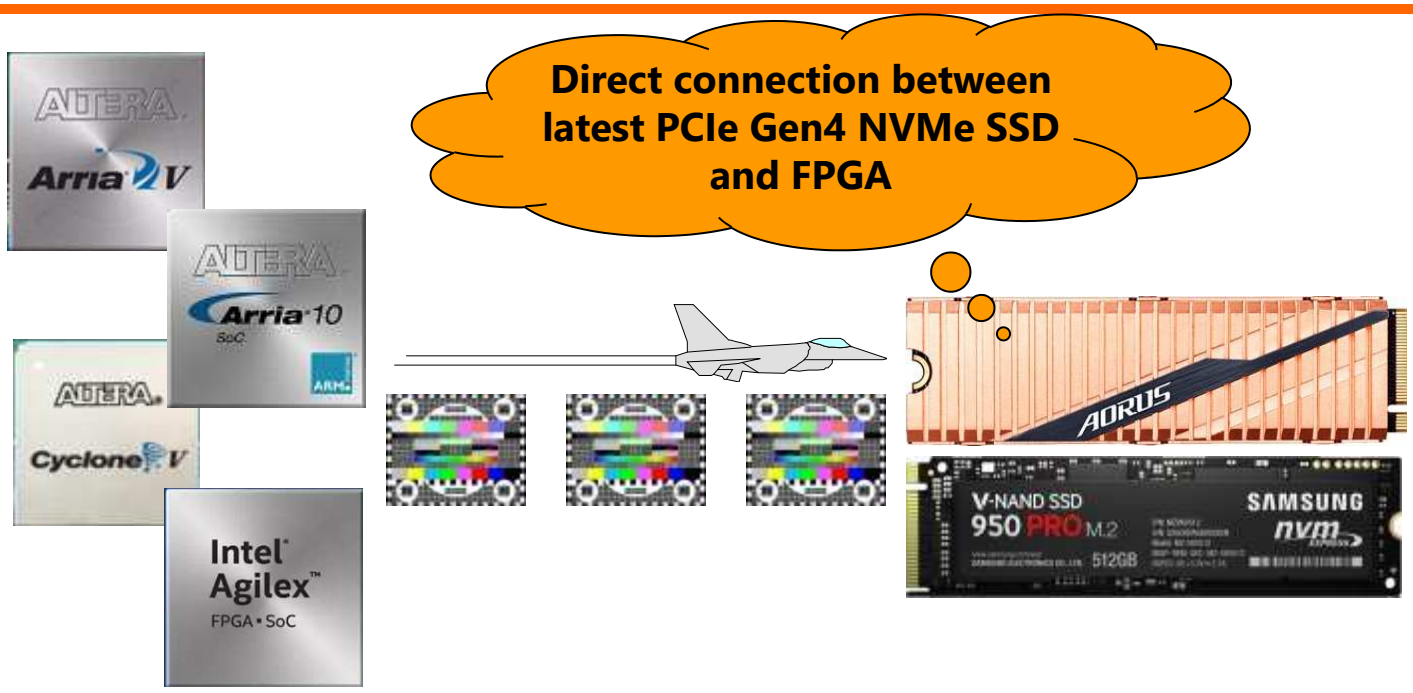

The Very Best Solution for Data Recording Application!

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

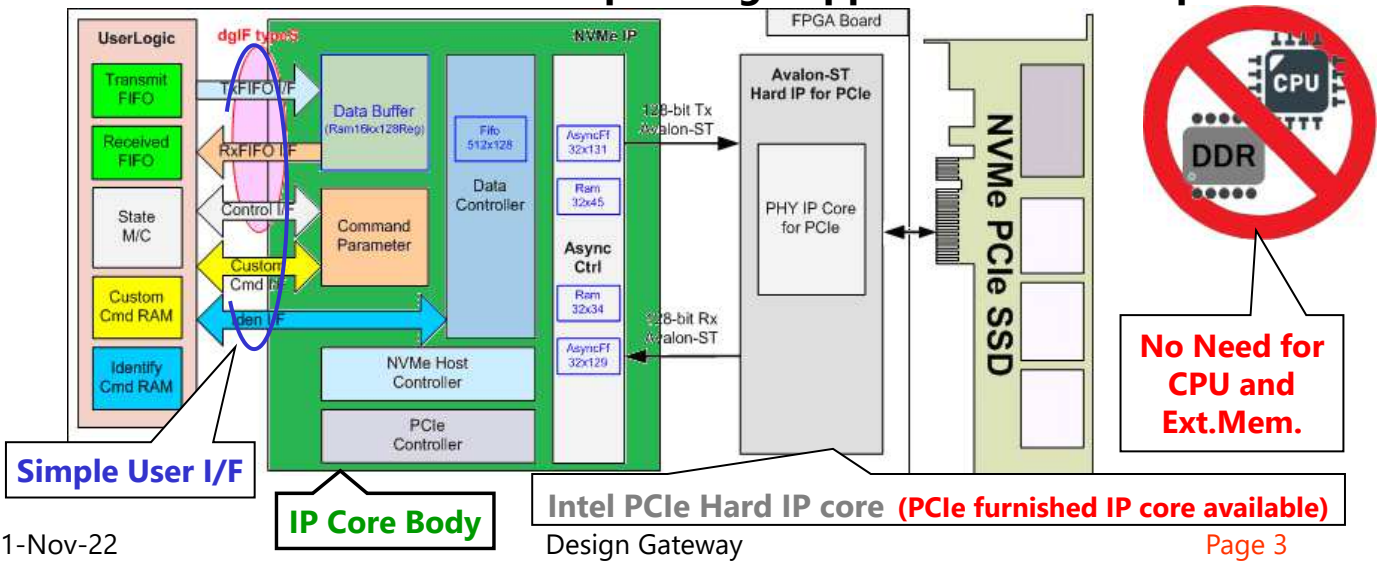

NVMe-IP Lineup

- **Multiple lineup for various functions**
  - NVMe-IP with… PCIe Soft-IP, external switch, multi-user …

Core type	Description
Standard NVMe-IP core	Standard core using Avalon-ST PCIe Hard IP in FPGA
NVMeG3-IP core	PCIe Soft-IP furnished in NVMe-IP, 4-Lane PCIe Gen3
NVMeSW-IP core	Multiple SSD connection via external PCIe switch
raNVMe-IP core	Supports random read or write access
muNVMe-IP core NEW!	Multiple user (port) with individual access

NVMe-IP core lineup

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

PCIe Soft-IP furnished IP core

- **NVMeG3-IP core**
  - Can operate without Avalon-ST PCIe Hard IP
  - Includes data link layer and connect with transceiver by PCIe Gen3
  - More SSD connection regardless of PCIe Hard IP count in FPGA.

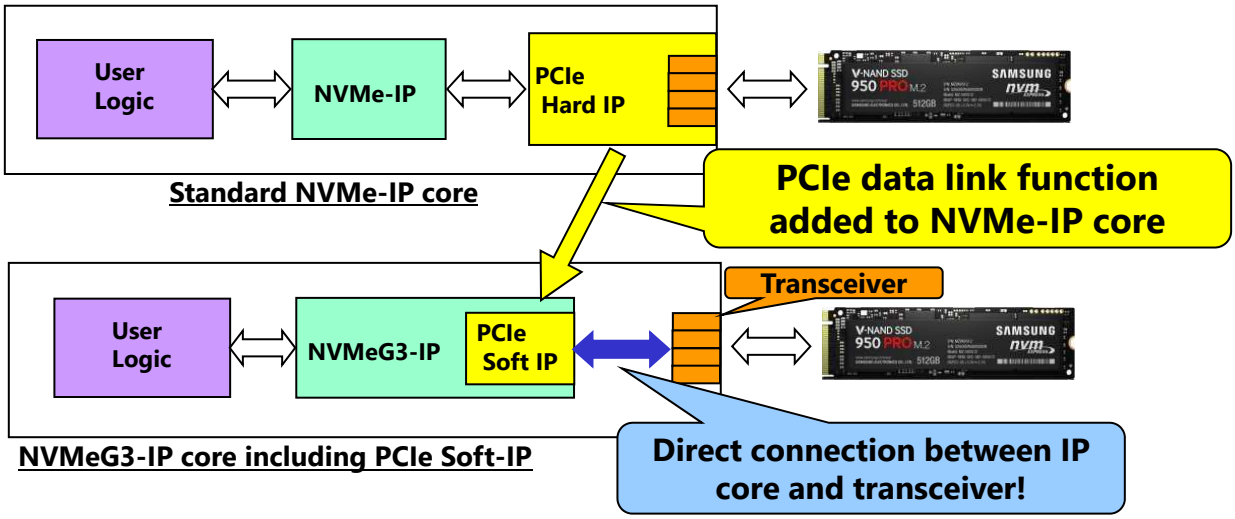

raNVMe-IP for random access

- 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.

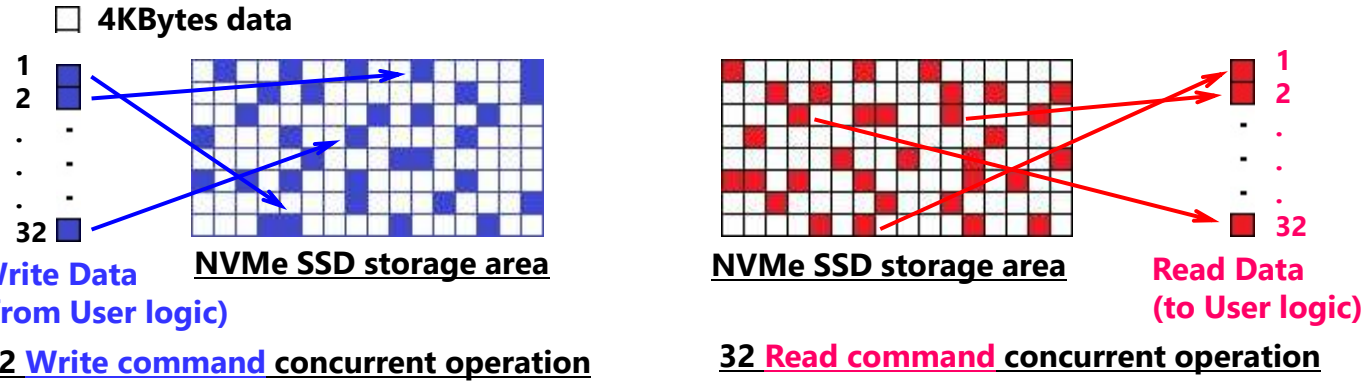

32 Write command concurrent operation

32 Read command concurrent operation

raNVMe-IP concurrent command operation image

Multiple user (multiple port) support

- **muNVMe-IP core**
  - **Multiple user port R/W accessible to single SSD individually**
  - **Suitable for simultaneous recording and reading application**

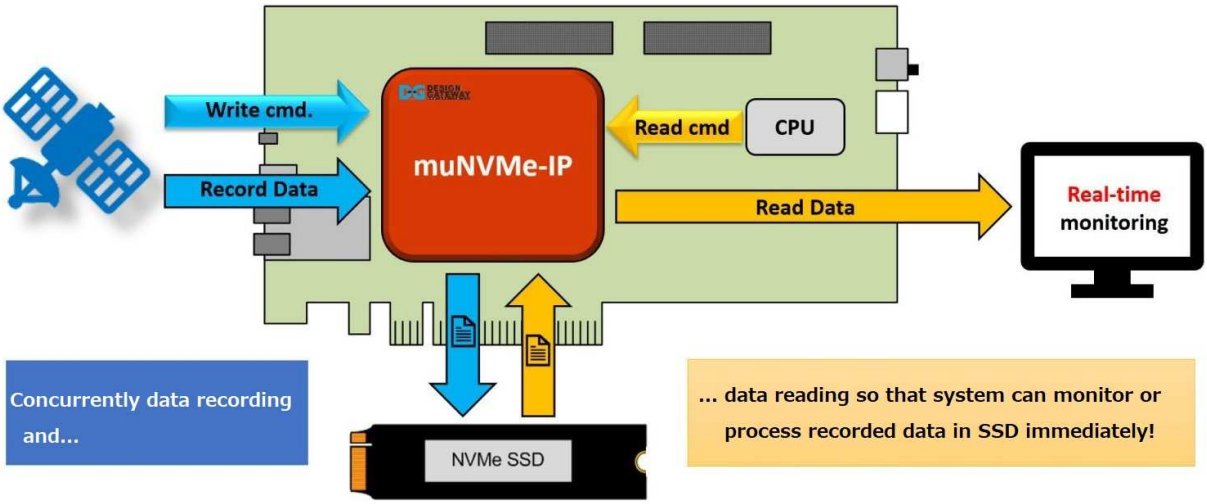

Simultaneous data recording and reading system by muNVMe-IP core

NVMe-IP Merit

1. **High Performance and Compact size**
   - **Write=6294MB/s, Read=6714MB/s (measured by Agilex-F)**
   - **Support PCIe GEN4/3 (Gen4 operation confirmed on Agilex-F board)**
   - **Core size: 1820ALM, 3680DFF (for Arria10SX case)**
2. **Interface: Simple and easy connection**
   - **User I/F control is parameter with pulse, data is simple FIFO**
   - **Use Block Mem. 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 Quartus project with real board operation in the package**

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**

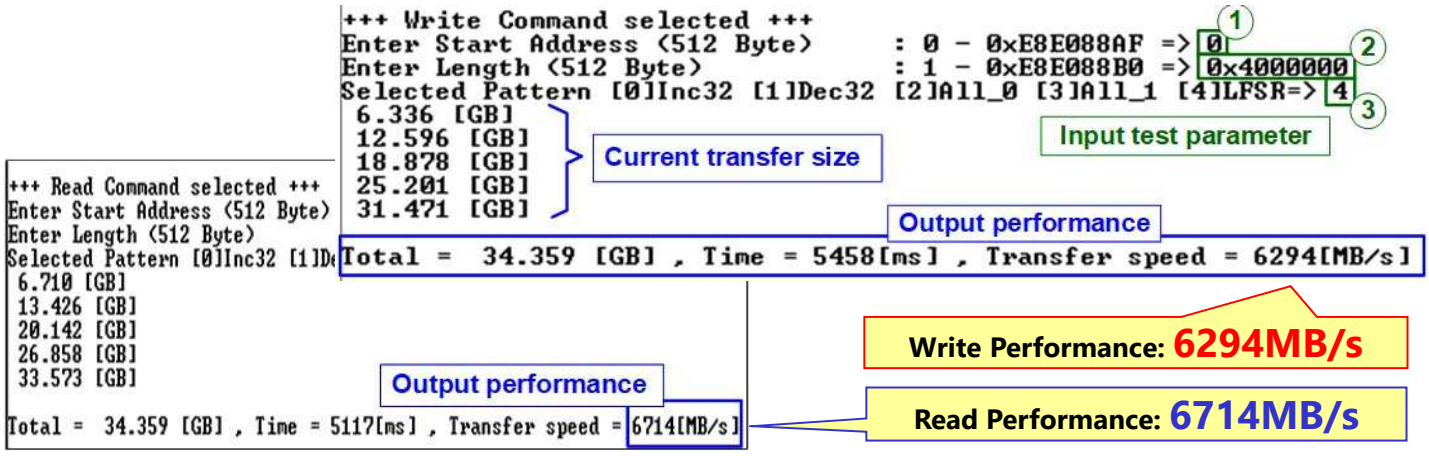

Performance Evaluation Result (Agilex-F)

Notes1: SSD=Addlink S95 2TB Gen4
Notes2: buffer size is set to 1MB for Read performance

Merit1: Compact Size (Standard core)

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

Family	Example Device	Fmax (MHz)	Logic utilization (ALMs)	Registers	Block Memory bit[1]	Design Tools
ArriaV GX	5AGXFB3H4F35C4	212	1835	3637	2,162,688	QuartusII 16.0
Stratix V GX	5SGXMA7N2F45C2	350	1813	3682	2,162,688	QuartusII 16.0
Arria10 SX	10AS066N3F40E2SGE2	300	1820	3680	2,162,688	QuartusII 16.0

Family	Example Device	Buf Mode	Fmax (MHz)	Logic utilization (ALMs)	Registers	Block Memory bit	Design Tools
Agilex F-Series	AGFB014R24A2E3VR0	1 MB	375	4939	10,485	8,496,896	Quartus 22.3
		256 KB	375	3970	9,191	2,205,440	

NVMe-IP (standard) Core standalone resource usage

Merit1: Compact Size (PCIe Soft-IP furnished)

- **Resource usage of PCIe Soft-IP furnished IP core**

Family	Example Device	Fmax (MHz)	Logic utilization (ALMs)	Registers	Block Memory	Design Tools
Arria10 GX	10AX115S2F45I1SG	300	8560	10984	140 M20Ks	QuartusII 18.0

NVMeG3-IP core (PCIe Soft-IP furnished) standalone resource usage

- **Limitation point of NVMeG3-IP core**
  - **PCIe Gen3 only, not support other speed (Gen1/2)**
  - **4-Lane only, not support other lane count (1/2/8/16)**
    (Ask other lane count as core customization)

Merit2: Command I/F

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

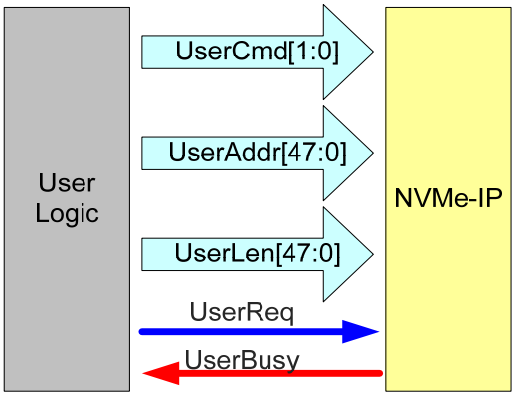

Basic Command I/F Signals

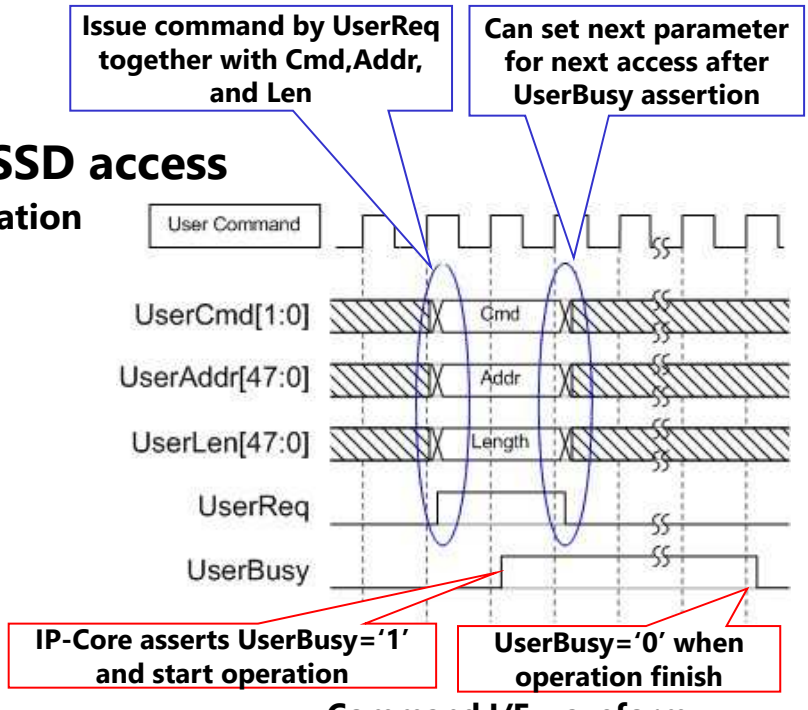

Command I/F waveform

Merit2: Data I/F

- **Simple 128/256bit FIFO for each of read and write**
  - **General FIFO of standard Intel library**
  - **Data buffer using 256KByte Block memory in NVMe-IP**

(*1) NVMe-IP for Gen4 can select 1MB buffer size to increase performance

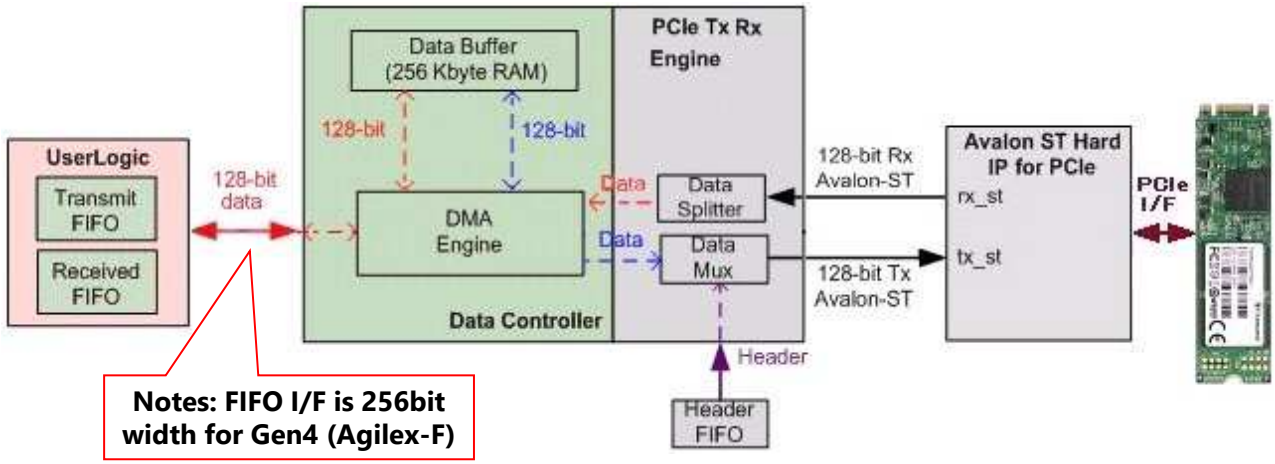

Data path of NVMe-IP

Merit 3: 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 Intel evaluation board**
- **Free sof-file for evaluation before IP-core purchase**

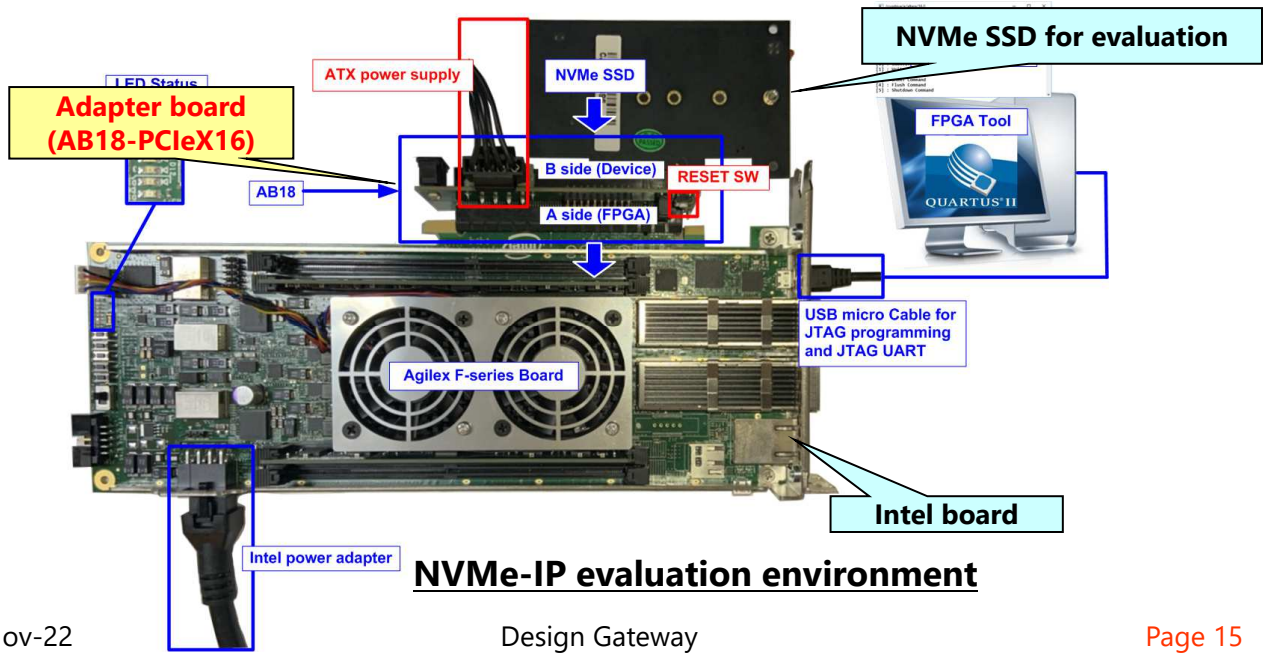

NVMe-IP evaluation environment

Merit4: Development Tool#1

- **PCIe Adapter board for evaluation (Part#: AB18-PCIeX16)**
  - **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**

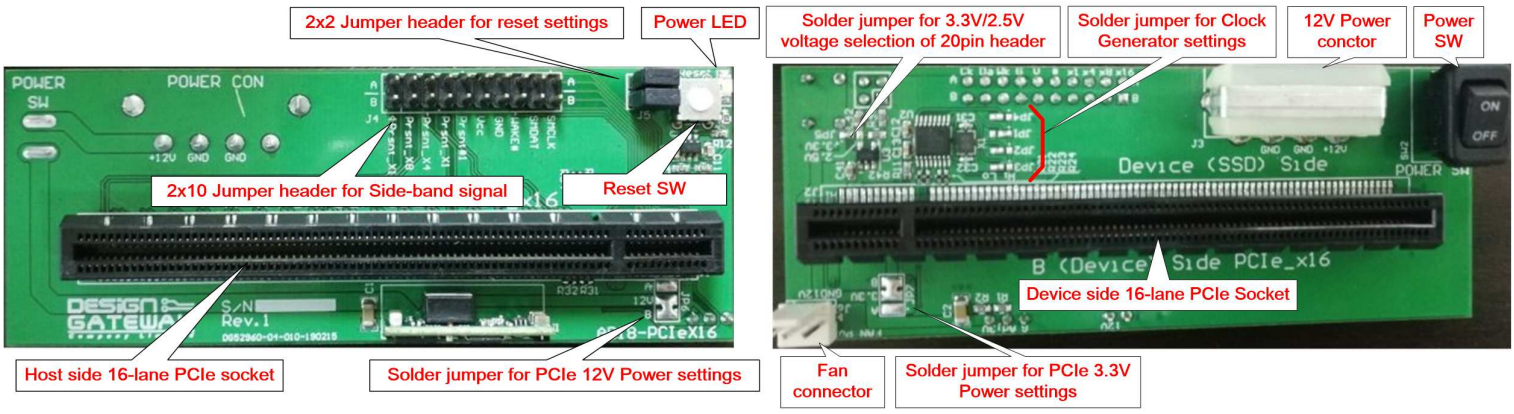

PCIe adapter for NVMe-IP evaluation (AB18-PCIeX16)

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)

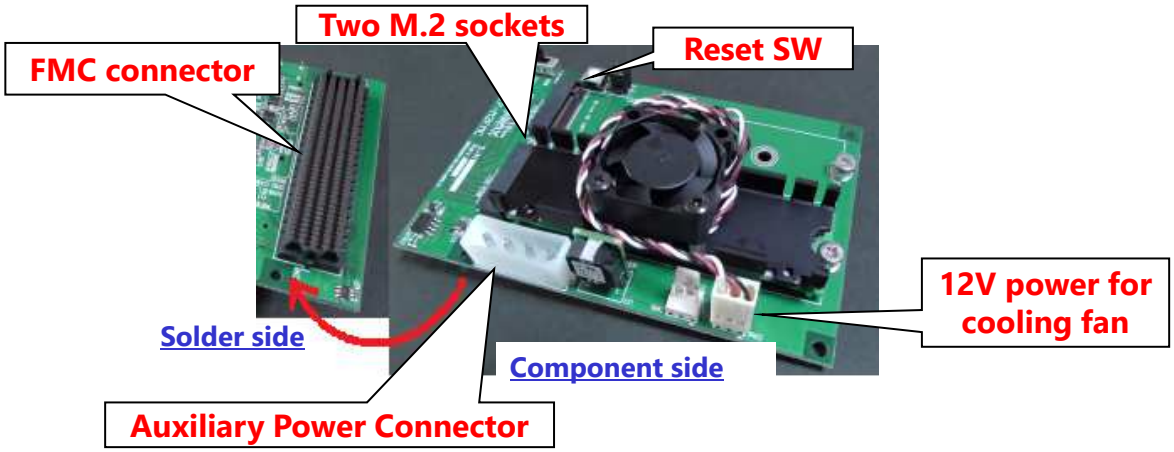

FMC adapter for NVMe-IP evaluation (AB17-M2FMC)

Merit4: Reference Design

- **Quartus/Qsys 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

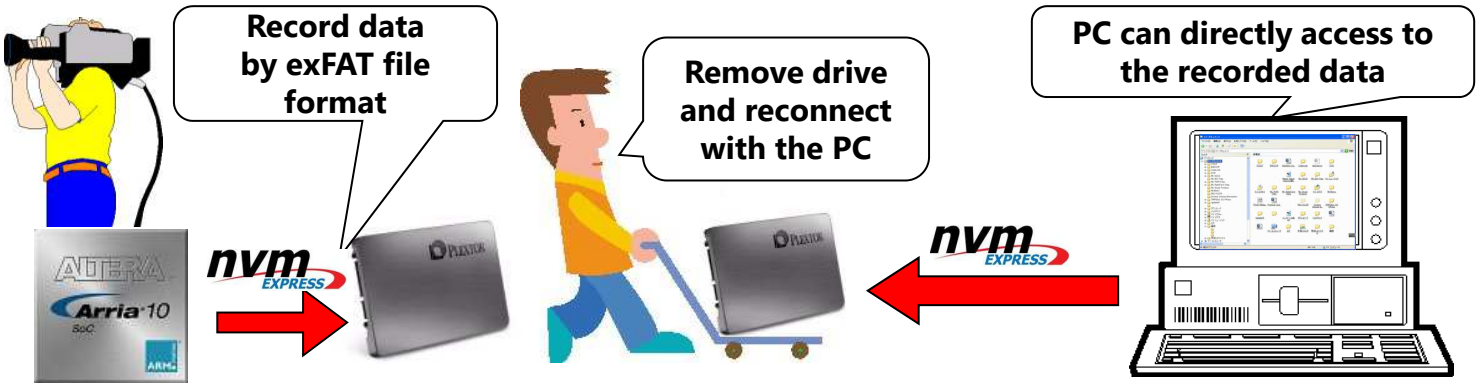

PC can directly read recorded data as a file

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.**

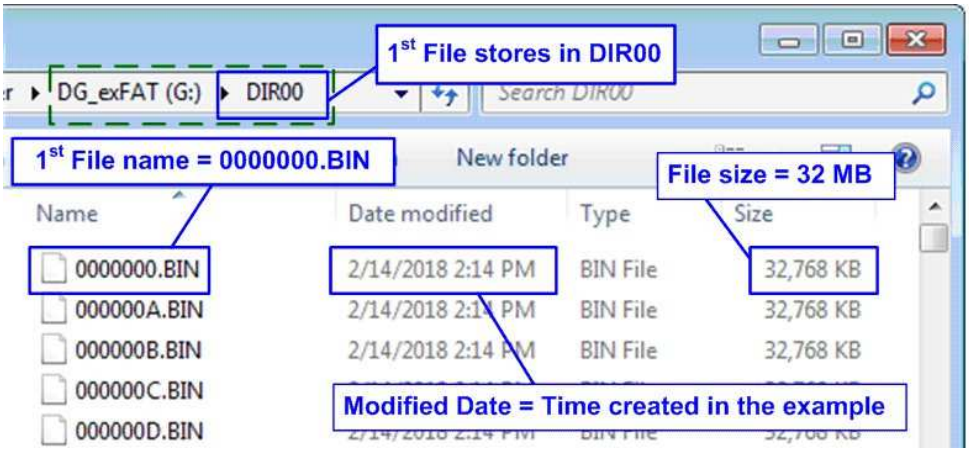

Generate test file, reconnect with PC, and can check file read compatibility

NVMe-IP Application Example 1

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

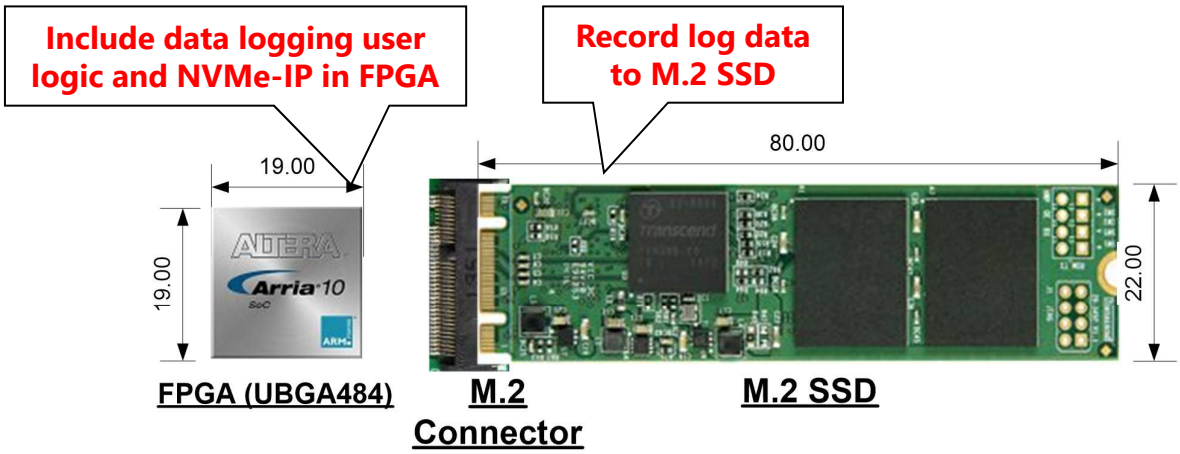

System space image by UBGA 484 FPGA and M.2 SSD (unit: mm)

NVMe-IP Application Example 2

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

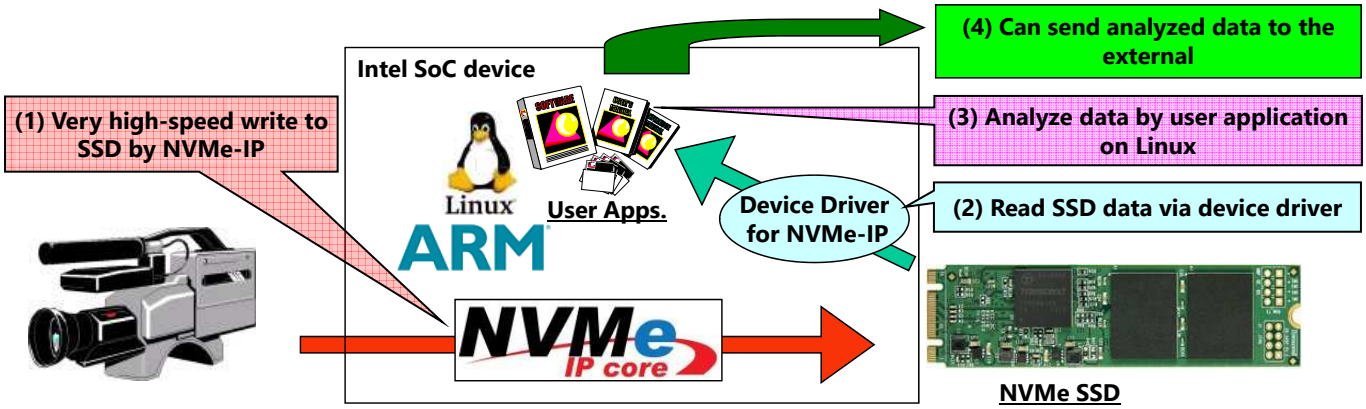

Recording and Analysis sytem on Linux (device driver and reference design available)

NVMe-IP Application Example 3

- **Ultra High-Speed Recorder**
  - Double write speed with multiple SSDs RAID0 configuration
  - Provide RAID0 reference design with 2 or 4 NVMe SSDs

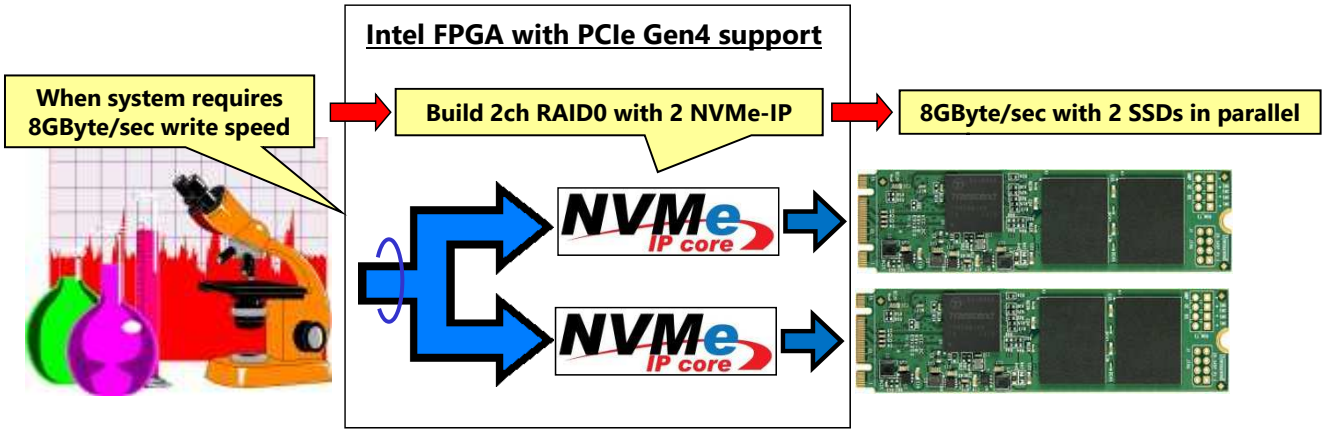

NVMe RAID system supporting 8GByte/sec recording rate

NVMe-IP Application Example 4

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

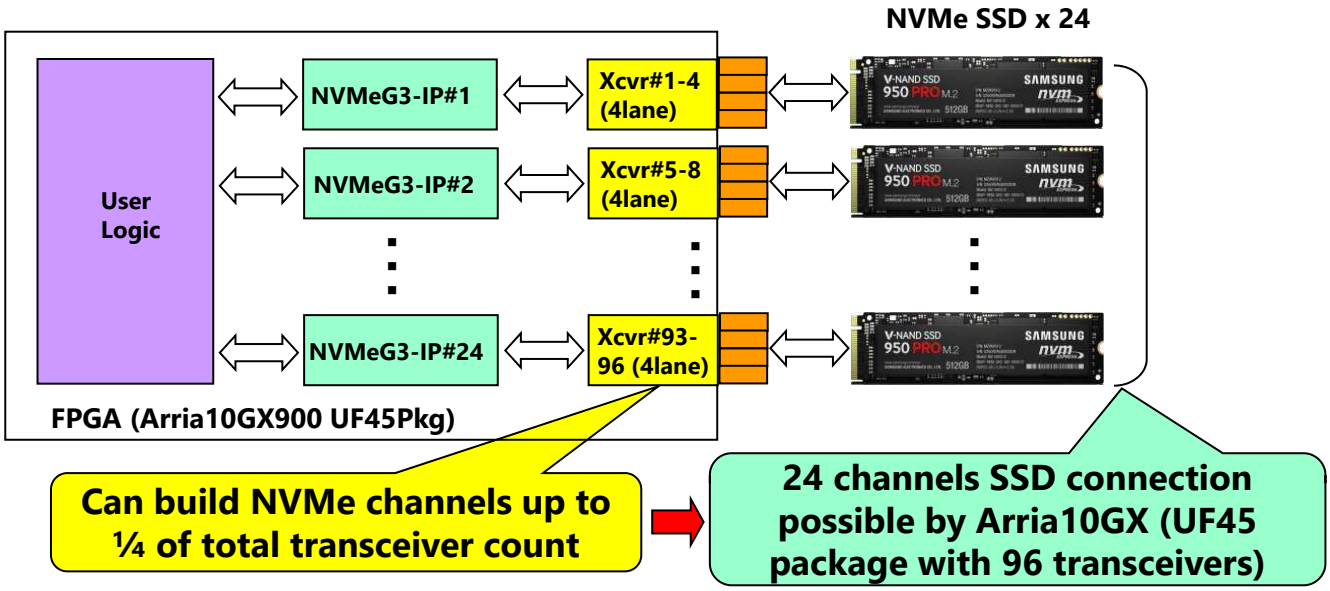

24 channels M.2 SSD Array system using NVMeG3-IP core

NVMe-IP Application Example 5

- **Concurrent data recording/reading system by muNVMe-IP**
  - **SSD access by both sides of User#0(write)/User#1(read)**

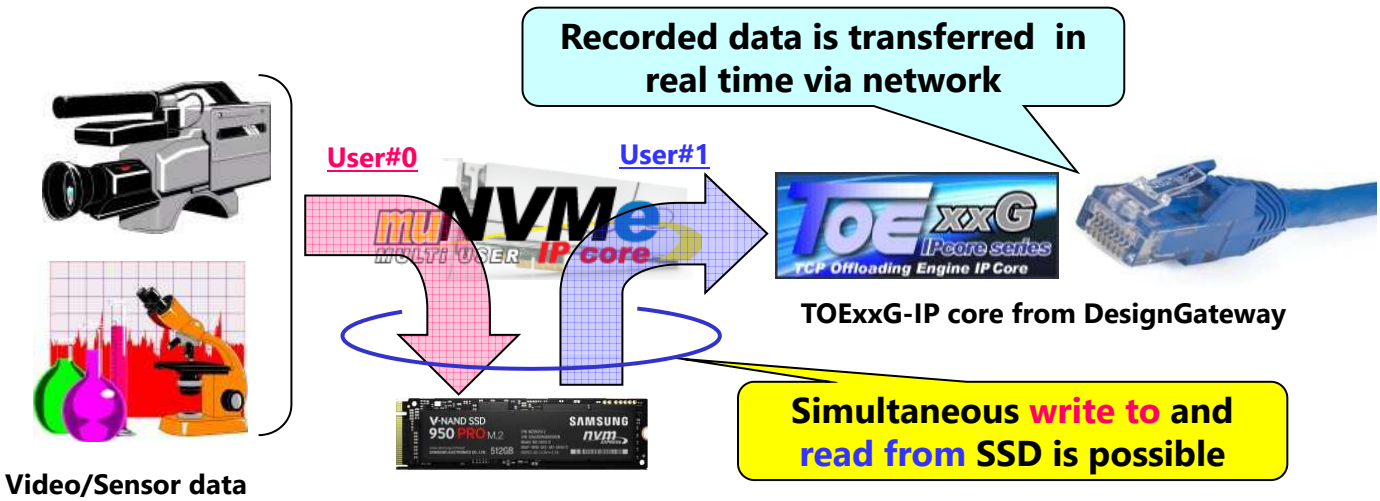

Concurrent recording/reading system by muNVMe-IP core

For more detail

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

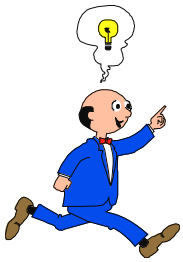

Revision History

Rev.	Date	Description
0.1E	4-Aug-16	English Temporary Version (Ver0.1E)
1.0E	10-Aug-16	First release with resource usage information
1.1E	25-Aug-16	Modify page17 because only one x16 DDR4 device can keep NVMe SSD performance
1.2E	21-Dec-16	NVMe-IP core improvement by removing external DDR chip for data buffer
1.3E	23-May-17	Performance improved by internal PCIe bridge in NVMe-IP core
1.4E	6-Jun-17	Data buffer size fixed to 256KByte
1.5E	2-Nov-17	Added Linux driver application and 2ch RAID0 reference design
1.6E	18-Jul-18	Added 4KB sector format, SMART/FLUSH/Shutdown command support
1.7E	9-Jan-19	Add FAT32-IP/exFAT-IP for NVMe-IP optional products
2.0E	3-May-20	Add new product of NVMeG3-IP that includes PCIe Soft IP core inside
2.1E	27-Aug-20	Add new product of raNVMe-IP for random access application
2.2E	3-Jun-21	Support Agilex-F device with PCIe Gen4 speed
2.3EA	29-Oct-22	Added muNVMe-IP line-up