

# NVMe-IP with PCIe Gen4 Soft IP reference design manual

Rev1.04 22-Dec-23

1	Overview.....	2
2	Hardware overview .....	3
2.1	TestGen .....	4
2.2	NVMe.....	8
2.2.1	NVMeG4-IP.....	8
2.2.2	PCIe PHY.....	8
2.2.3	Dual port RAM.....	9
2.3	CPU and Peripherals.....	10
2.3.1	AsyncAxiReg.....	11
2.3.2	UserReg.....	13
3	CPU Firmware .....	16
3.1	Test firmware (nvmeiptest.c).....	16
3.1.1	Identify Command.....	16
3.1.2	Write/Read Command.....	17
3.1.3	SMART Command.....	17
3.1.4	Flush Command.....	18
3.1.5	Secure Erase Command.....	18
3.1.6	Shutdown Command.....	18
3.2	Function list in Test firmware.....	19
4	Example Test Result.....	22
5	Revision History.....	23

# 1 Overview

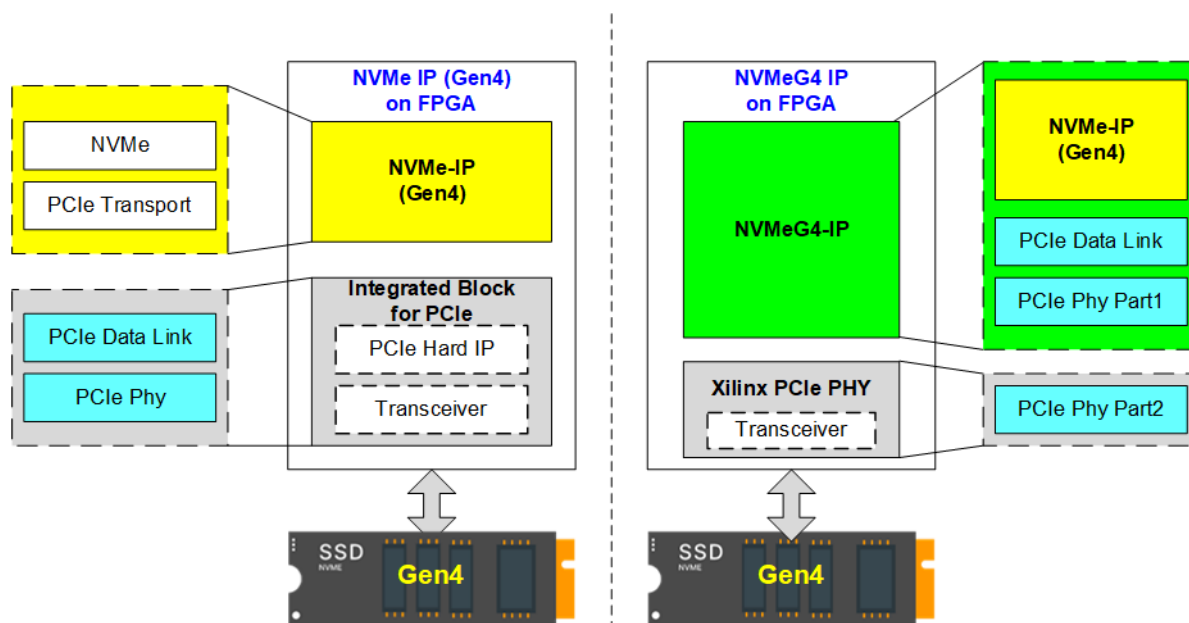


Figure 1-1 NVMe-IP (Gen4) and NVMeG4-IP comparison

The NVMe SSD has become a widely adopted storage solution due to its numerous advantages over SATA storage. Designed specifically to utilize the full potential of SSDs, the NVMe protocol offers an architecture that enables high-performance random access and simultaneous multi-user access. These features are essential for the applications that require exceptional storage performance capabilities.

While traditional NVMe solutions on FPGA devices typically rely on NVMe accelerator engines operating with CPU, Design Gateway presents an inventive alternative – the NVMe-IP. This pure hardwired logic design serves as a stand-alone NVMe host controller, eliminating the need for both CPUs and external memory. The NVMe-IP achieves peak transfer performance, ensuring optimal utilization of SSD capabilities. However, the NVMe-IP requires a compatible FPGA model equipped with PCIe hard IP.

In response to the requirements of users who lack PCIe hard IP in their selected FPGA models, Design Gateway offers a new IP option – the NVMeG4-IP. This solution combines the features of the NVMe-IP with PCIe Gen4 Soft IP to a wider user base. By implementing the Data Link layer and a specific portion of the Physical layer within the PCIe protocol, the NVMeG4-IP delivers optimal performance. To complete the Physical layer implementation, NVMeG4-IP utilizes the PCIe PHY IP provided by Xilinx. Consequently, FPGA lacking PCIe Gen4 hard IP can access NVMe Gen4 SSDs through the NVMeG4-IP. Notably, the NVMe Gen4 SSDs offer twice the data write/read speed compared to NVMe Gen3 SSDs, reaching 7500 MB/s.

## 2 Hardware overview

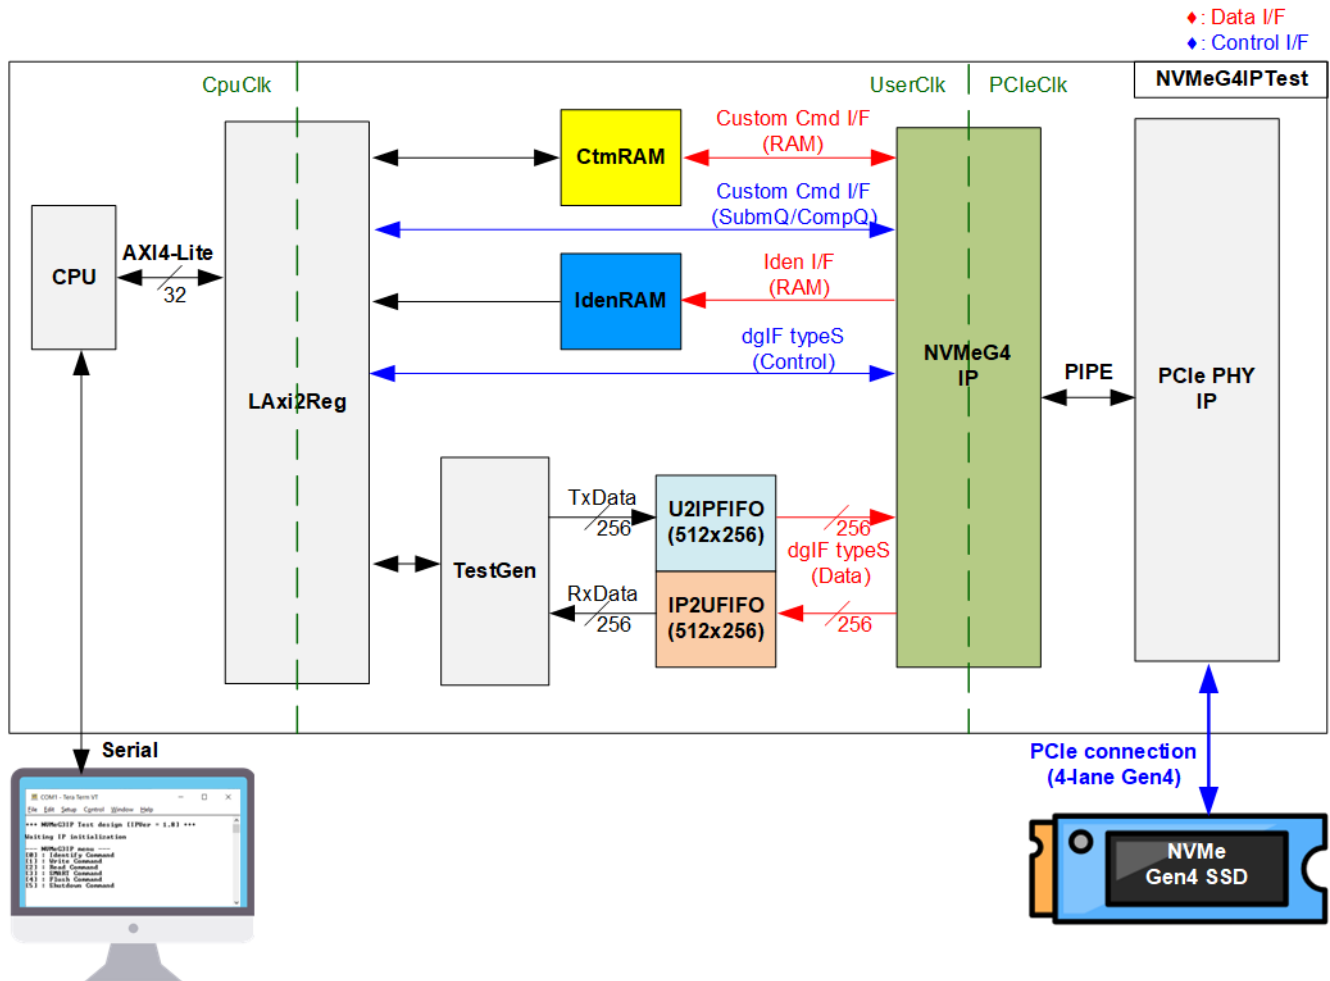


Figure 2-1 NVMeG4-IP demo hardware

The hardware modules in the test system can be categorized into three sections: test function (TestGen), NVMe function (CtmRAM, IdenRAM, U2IPFIFO, IP2UFIFO, NVMeG4-IP, and PCIe PHY), and CPU system (CPU and LAXI2Reg).

The TestGen connects to the user interface of NVMeG4-IP and is responsible for generating test data stream of Write command and verifying test data stream of Read command. The write and read data streams are stored in two FIFOs (U2IPFIFO and IP2UFIFO). The TestGen always writes or reads data when the FIFO is ready, allowing for optimal transfer performance evaluation for the system.

The NVMe consists of the NVMeG4-IP and the PCIe PHY IP, enabling direct access to an NVMe Gen4 SSD without PCIe switch connection. The command request and the parameters of each command (the inputs of NVMeG4-IP) are controlled by the CPU through LAXI2Reg module. While the data interface for both Custom and Identify commands is connected to RAMs that are accessible by the CPU.

The CPU is connected to the LAXi2Reg module, for interface with the NVMe test logics. Integrating the CPU into the test system allows the user to set the test parameters and monitor the test status via the Serial console. Using CPU also facilitates the execution of multiple test cases to verify the functionality of the IP. The default firmware for the CPU includes functions for executing the NVMe commands by using the NVMeG4-IP.

Figure 2-1 displays three clock domains: CpuClk, UserClk, and PCIeClk. CpuClk serves as the clock domain for the CPU and its peripherals, requiring a stable clock that can function independently from other hardware. UserClk is the clock domain utilized for the operation of the NVMeG4-IP, RAM, and TestGen. According to the NVMeG4-IP datasheet, the frequency of UserClk must be greater than or equal to PCIeClk. The reference design utilizes 275 MHz for UserClk at PCIe Gen4 speed. PCIeClk, generated by PCIe PHY IP, is synchronized with the 256-bit data bus on PIPE and operates at 250 MHz for 4-lane PCIe Gen4. Further hardware details are described below.

## 2.1 TestGen

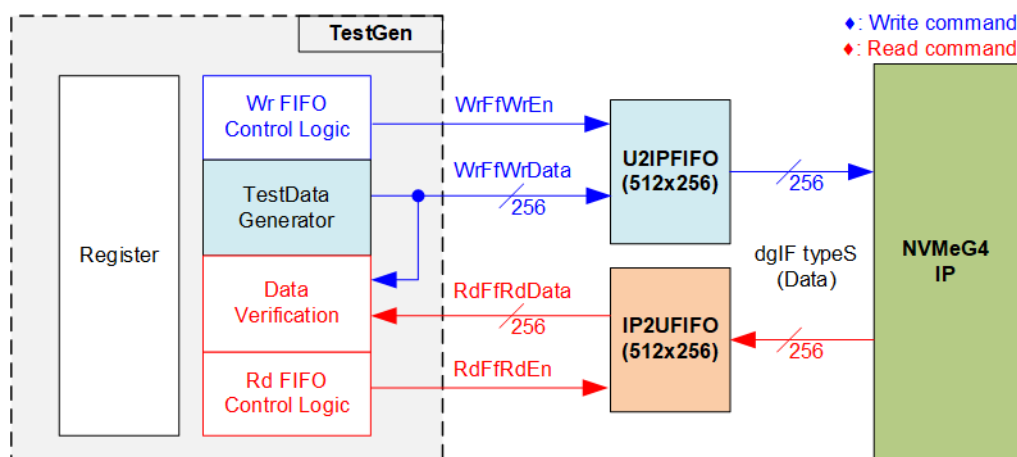


Figure 2-2 TestGen interface

The TestGen module handles the data interface of NVMeG4-IP, facilitating data transfer for both Write and Read commands. In case of a Write command, TestGen sends 256-bit test data to NVMeG4-IP via U2IPFIFO. In contrast, for a Read command, the test data is received from IP2UFIFO for comparison with the expected value, ensuring data accuracy. Data bandwidth of TestGen is set to match that of NVMeG4-IP by running at the same clock and data bus size. The control logic ensures that the Write or Read enable is always asserted to 1b when the FIFO is ready to write or read data, respectively. This ensures that both U2IPFIFO and IP2UFIFP are always ready to transfer data with NVMeG4-IP without delay, providing the best performance for writing and reading data with the SSD through NVMeG4-IP.

To provide a flexible test environment, the user can set some test parameters to control the TestGen module such as total transfer size, transfer direction, and test pattern selector via the console. These test parameters are stored in the Register block. The detailed hardware logic of TestGen is illustrated in Figure 2-3.

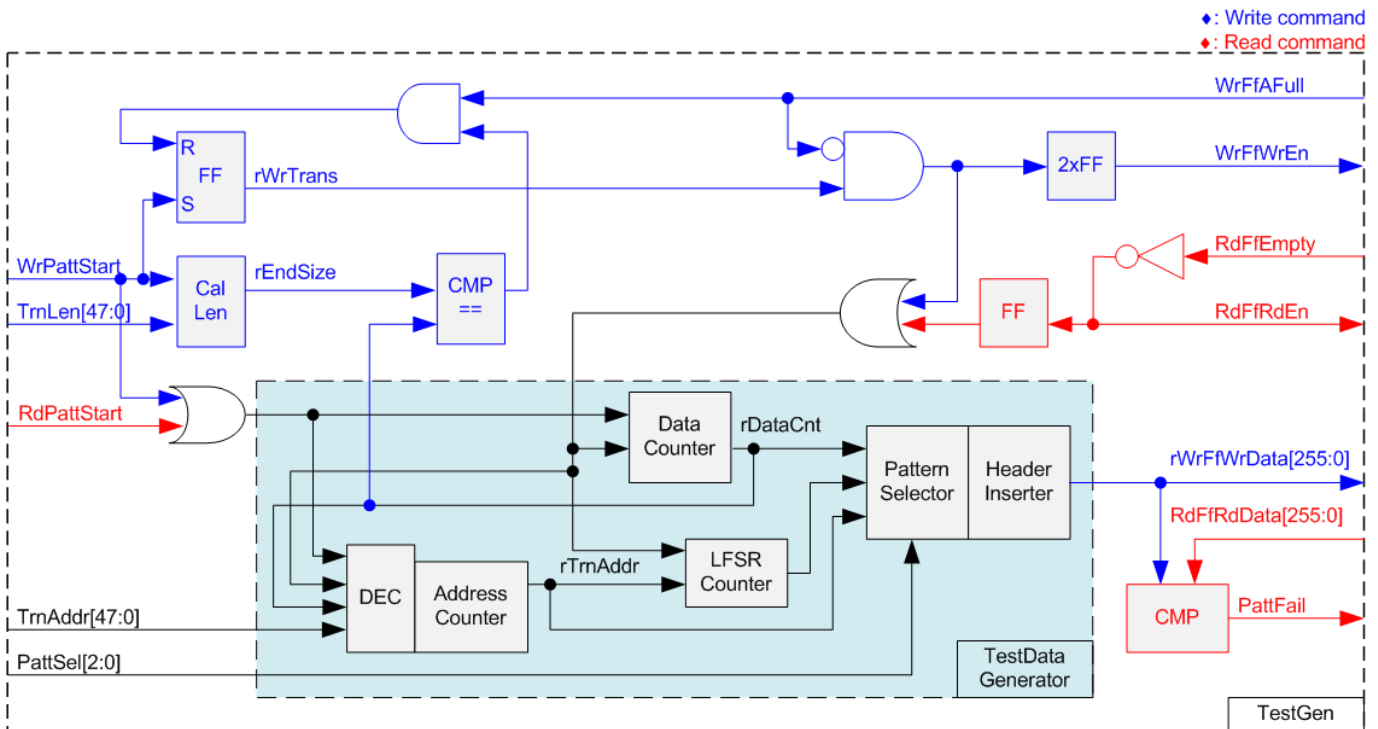


Figure 2-3 TestGen hardware

In Figure 2-3, two key aspects of the system are depicted. The first part illustrates the control of data flow, while the second part details the generation of test data for use with the FIFO interface.

In the upper portion of Figure 2-3, we focus on the control of data flow. Two signals, the `WrFfAFull` and `RdFfEmpty`, are integral to the FIFO interface for flow control. When the FIFO reaches its capacity (indicated by `WrFfAFull=1b`), the `WrFfWrEn` signal is set to `0b`, effectively pausing data transfer into the FIFO. In a read operation, when data is available within the FIFO (indicated by `RdFfEmpty=0b`), the system retrieves this data for comparison by setting the `RdFfRdEn` to `1b`. Furthermore, it is important to note that both write and read operation are completed when the total transferred data matches the user-defined value. Consequently, the counter logic is designed to track the amount of data transferred during this command, and upon command completion, both `WrFfWrEn` and `RdFfRdEn` must be de-asserted.

The lower section of Figure 2-3 outlines the methods for generating test data, either for writing to the FIFO or for data verification. There are five available test patterns: all-zero, all-one, 32-bit incremental data, 32-bit decremental data, and LFSR. These patterns are selected by the Pattern Selector.

For the all-zero or all-one pattern, every bit of the data is set to zero or one, respectively. Conversely, the other test patterns are designed by separating the data into two parts to create unique test data within every 512-byte data, as shown in Figure 2-4.

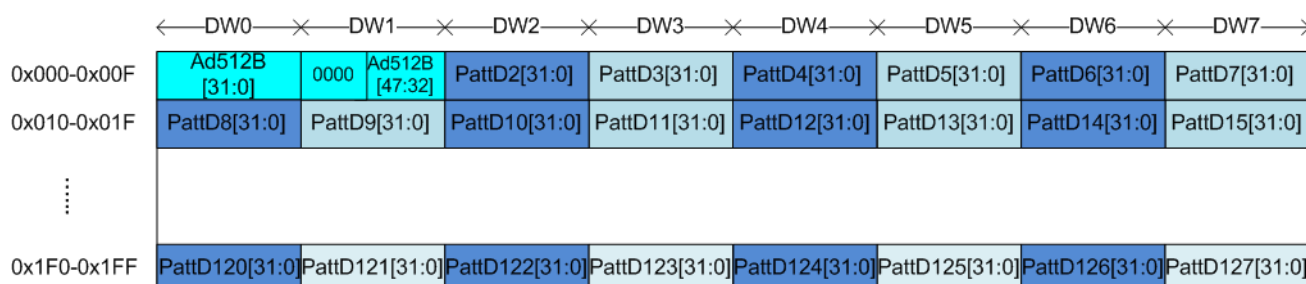


Figure 2-4 Test pattern format in each 512-byte data for Increment/Decrement/LFSR pattern

Each 512-byte data block consists of a 64-bit header in Dword#0 and Dword#1, followed by the test data in the remaining words of the 512-byte data (Dword#2 – Dword#127). The header is created using the Address counter block, which operates in 512-byte units. The initial value of the Address counter is configured by the user and increases after transferring each 512-byte data.

The content of the remaining Dwords (DW#2 – DW#127) depend on the pattern selector, which could be 32-bit incremental data, 32-bit decremental data, or the LFSR pattern. The 32-bit incremental data is designed using the Data counter, while the decremental data can be created by connecting NOT logic to the incremental data. The LFSR pattern is generated using the LFSR counter, using the equation  $x^{31} + x^{21} + x + 1$ . To generate a 256-bit test data using the LFSR pattern, the data is divided into two sets of 128-bit data. Each set uses a different start value, as shown in Figure 2-5.

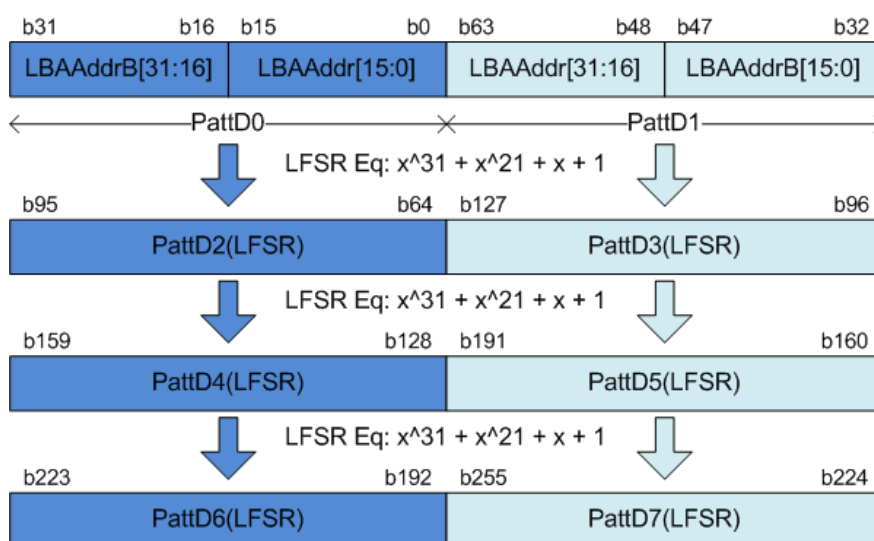



Figure 2-5 256-bit LFSR Pattern in TestGen

Using a look-ahead technique, the test data generation process produces four 32-bit LFSR data or 128-bit data during each clock cycle. These 128-bit data sets are represented by the same color in Figure 2-5. The initial value of each data set is derived by combining certain bits of LBAAddr and LBAAddrB (a NOT logic operation of LBAAddr). This approach ensures that each 128-bit data set uses a unique initial value.

The generated test data is then written to the FIFO as write data or used as expected data for verification with the read data from the FIFO. 'PattFail' flag is set to 1b when the data verification process fails. The timing diagram for writing data to the FIFO is shown below.



**Figure 2-6 Timing diagram of Write command in TestGen**

- 1) The write operation is initiated by setting WrPattStart signal to 1b for a single clock cycle, which is followed by the assertion of rWrTrans to enable the control logic for generating write enable to FIFO.
- 2) If two conditions are satisfied (rWrTrans is asserted to 1b during the write operation and the FIFO is not full, indicated by WrFfAFull=0b), the write enable (rWrFfWrEn) to FIFO is asserted to 1b.
- 3) The write enable is fed back to the counter to count the total amount of data in the write operation.
- 4) If FIFO is almost full (WrFfAFull=1b), the write process is paused by de-asserting rWrFfWrEn to 0b.
- 5) The write operation is finished when the total data count (rDataCnt) is equal to the set value (rEndSize). At this point, both rWrTrans and rWrFfWrEn are de-asserted to 0b.

For read transfer, the read enable of FIFO is controlled by the empty flag of FIFO. Unlike the write enable, the read enable signal is not stopped by total data count and not started by start flag. When the read enable is asserted to 1b, the data counter and the address counter are increased for counting the total amount of data and generating the header of expect value, respectively.



## 2.2 NVMe

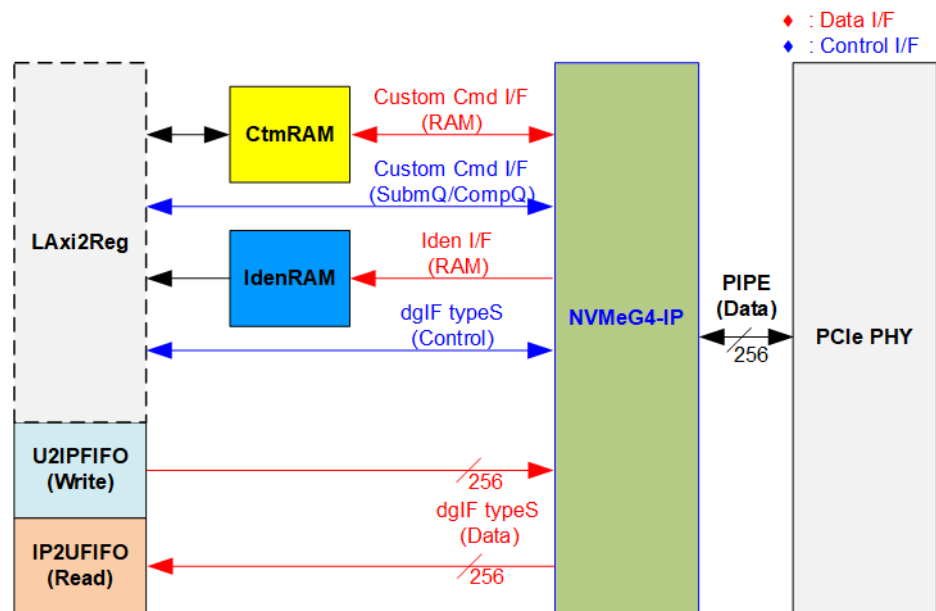


Figure 2-7 NVMe hardware

In the reference design, the NVMeG4-IP’s user interface consists of a control interface and a data interface. The control interface receives commands and parameters from either the Custom command interface or dgIF typeS, depending on the type of command. For instance, Custom command interface is used when operating SMART, Secure Erase, or Flush command.

On the other hand, the data interface of NVMeG4-IP has four different interfaces with a data bus width of 256-bit. These interfaces include Custom command RAM interface, Identify interface, FIFO input interface (dgIF typeS), and FIFO output interface (dgIF typeS). While the Custom command RAM interface is a bi-directional interface, the other interfaces are unidirectional interface. In the reference design, the Custom command RAM interface is used for one-directional data transfer when NVMeG4-IP sends SMART data to LAXI2Reg.

### 2.2.1 NVMeG4-IP

The NVMeG4-IP implements NVMe protocol of the host side to direct access an NVMe Gen4 SSD without PCIe switch connection. Besides, it includes PCIe soft IP operating with PCIe PHY to support the same feature as PCIe hard IP. It supports seven commands, i.e., Write, Read, Identify, Shutdown, SMART, Secure Erase, and Flush. More details of NVMeG4-IP are described in datasheet.

[https://dgway.com/products/IP/NVMe-IP/dg\\_nvme4\\_ip\\_data\\_sheet\\_xilinx\\_en/](https://dgway.com/products/IP/NVMe-IP/dg_nvme4_ip_data_sheet_xilinx_en/)

### 2.2.2 PCIe PHY

The PCIe PHY is provided by Xilinx to enable the utilization of Soft IP instead of Hard IP to construct a PCIe MAC. The PCIe PHY utilizes the PHY Interface for PCIe Express (PIPE) as its user interface. To operate with the NVMeG4-IP, it is configured to 4-lane, with each lane operating at 16.0 GT/s. For more detailed information of the PCIe PHY, please refer to the “PG239: PCI Express PHY” document, available at Xilinx website.

<https://docs.xilinx.com/r/en-US/pg239-pcie-phy>



### 2.2.3 Dual port RAM

Two dual port RAMs, CtmRAM and IdenRAM, store the returned data from Identify command and SMART command, respectively. IdenRAM has an 8 KB size and is used to store the 8 KB output from the Identify command.

The data bus size for NVMeG4-IP and LAXi2Reg differ, with NVMeG4-IP having a 256-bit size and LAXi2Reg having a 32-bit size. As a result, IdenRAM is configured as an asymmetric RAM with different bus sizes for its Write and Read interfaces.

NVMeG4-IP also has a double-word enable, which allows it to write only 32-bit data in certain cases. The RAM setting on Xilinx IP tool supports write byte enable, so a small logic circuit was designed to convert the double word enable to be write byte enable, as shown in Figure 2-8.

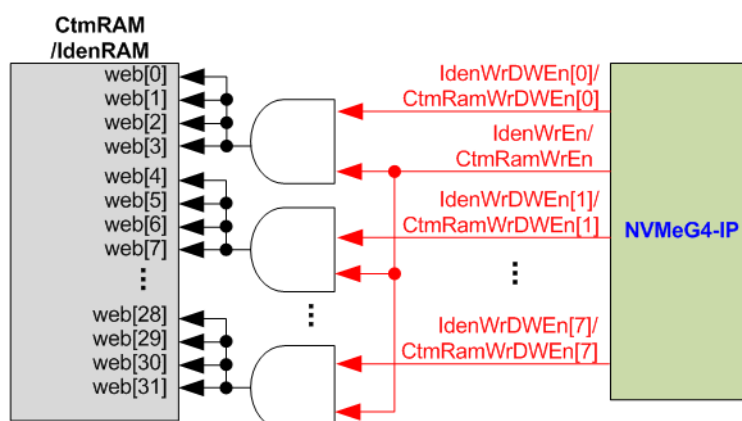


Figure 2-8 Byte write enable conversion logic

The input to the AND logic is bit[0] of WrDWEEn and the WrEn signal. The output of the AND logic is fed to bits[3:0] of IdenRAM byte write enable. Bit[1], [2], ..., [7] of WrDWEEn are then applied to bits[7:4], [11:8], ..., [31:28] of IdenRAM write byte enable, respectively.

On the other hand, CtmRAM is implemented as a true dual-port RAM with two read ports and two write ports, and with byte write enable. A small logic circuit must be used to convert the double word enable of Custom interface to byte write enable, similar to IdenRAM. The true dual-port RAM is used to support additional features when a customized Custom command requires data input. A simple dual-port RAM is sufficient to support the SMART command, even though the data size returned from the SMART command is 512 bytes. However, CtmRAM is implemented with an 8KB RAM for the customized Custom command.

### 2.3 CPU and Peripherals

The CPU system uses a 32-bit AXI4-Lite bus as the interface to access peripherals such as the Timer and UART. The system also integrates an additional peripheral to access NVMeG4-IP test logic by assigning a unique base address and address range. To support CPU read and write operations, the hardware logic must comply with the AXI4-Lite bus standard. LAXi2Reg module, as shown in Figure 2-9, is designed to connect the CPU system via the AXI4-Lite interface, in compliance with the standard.

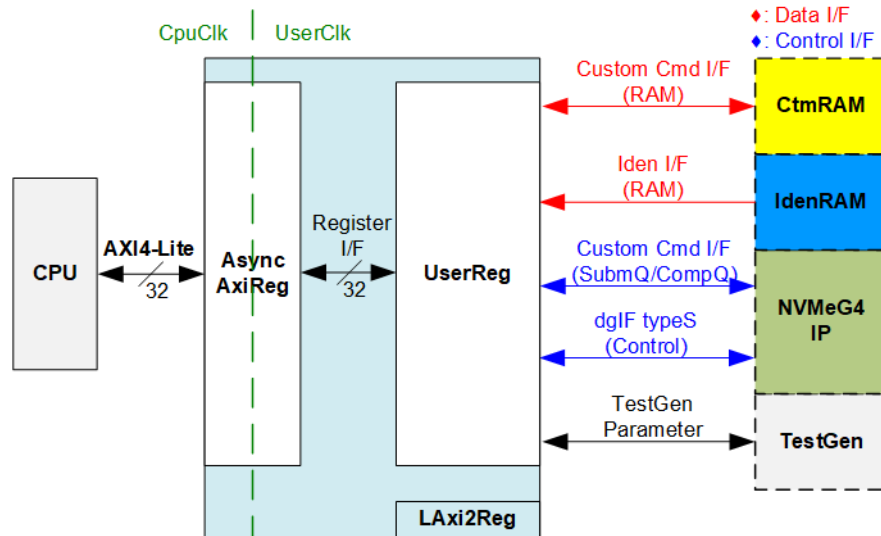


Figure 2-9 CPU and peripherals hardware

LAXi2Reg consists of AsyncAxiReg and UserReg. AsyncAxiReg converts AXI4-Lite signals into a simple Register interface with a 32-bit data bus size, similar to the AXI4-Lite data bus size. It also includes asynchronous logic to handle clock domain crossing between the CpuClk and UserClk domains.

UserReg includes the register file of parameters and the status signals of other modules in the test system, including the CtmRAM, IdenRAM, NVMeG4-IP, and TestGen. More details of AsyncAxiReg and UserReg are explained below.

### 2.3.1 AsyncAxiReg

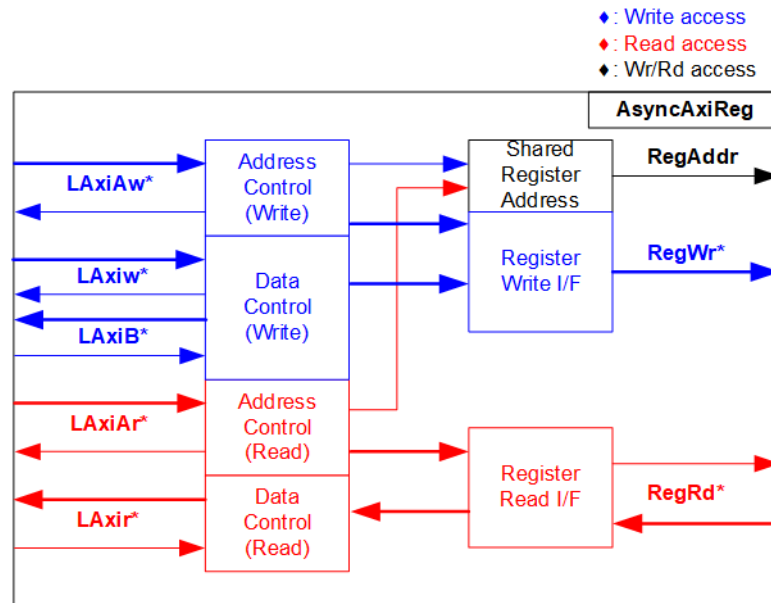


Figure 2-10 AsyncAxiReg Interface

The signal on AXI4-Lite bus interface can be grouped into five categories, i.e., LAXiAw\* (Write address channel), LAXiw\* (Write data channel), LAXiB\* (Write response channel), LAXiAr\* (Read address channel), and LAXir\* (Read data channel). More details to build Custom logic for AXI4-Lite bus can be found in the following document.

[https://github.com/Architech-Silica/Designing-a-Custom-AXI-Slave-Peripheral/blob/master/designing\\_a\\_custom\\_axi\\_slave\\_rev1.pdf](https://github.com/Architech-Silica/Designing-a-Custom-AXI-Slave-Peripheral/blob/master/designing_a_custom_axi_slave_rev1.pdf)

According to AXI4-Lite standard, the write channel and read channel operate independently for both control and data interfaces. Therefore, the logic within AsyncAxiReg to interface with AXI4-Lite bus is divided into four groups, i.e., Write control logic, Write data logic, Read control logic, and Read data logic, as shown in the left side of Figure 2-10. The Write control I/F and Write data I/F of the AXI4-Lite bus are latched and transferred to become the Write register interface with clock domain crossing registers. Similarly, the Read control I/F of AXI4-Lite bus is latched and transferred to the Read register interface, while Read data is returned from Register interface to AXI4-Lite bus via clock domain crossing registers. In the Register interface, RegAddr is a shared signal for write and read access, so it loads the value from LAXiAw for write access or LAXiAr for read access.

The Register interface is compatible with single-port RAM interface for write transaction. The read transaction of the Register interface has been slightly modified from RAM interface by adding the RdReq and RdValid signals to control read latency time. The address of Register interface is shared for both write and read transactions, so user cannot write and read the register at the same time. The timing diagram of the Register interface is shown in Figure 2-11.

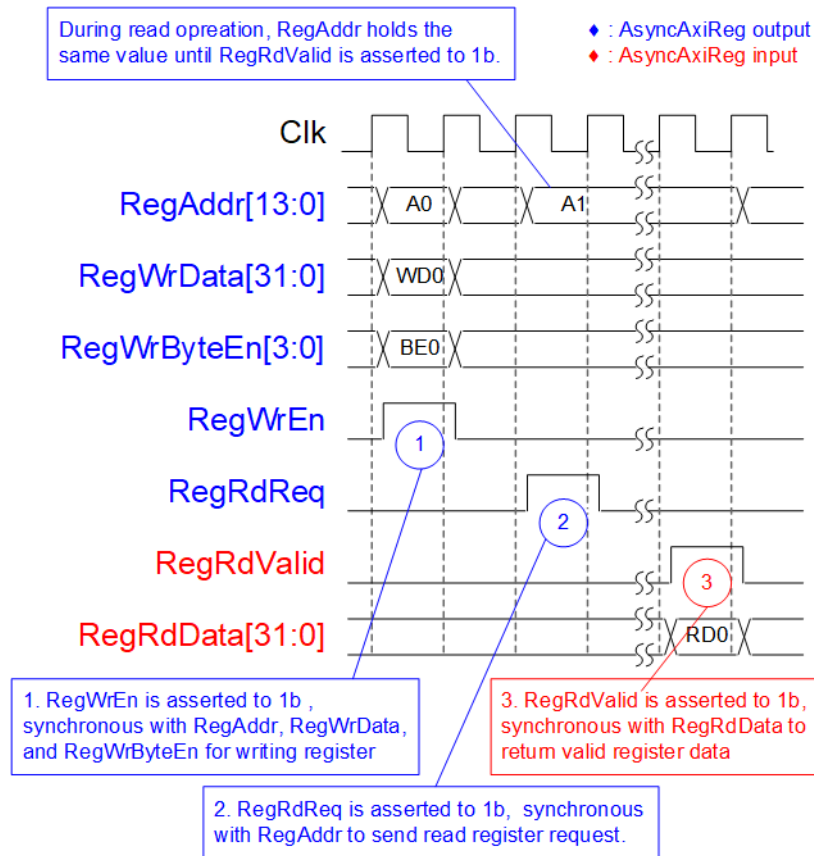


Figure 2-11 Register interface timing diagram

- 1) Timing diagram to write register is similar to that of a single-port RAM. The RegWrEn signal is set to 1b, along with a valid RegAddr (Register address in 32-bit units), RegWrData (write data for the register), and RegWrByteEn (write byte enable). The byte enable consists of four bits that indicate the validity of the byte data. For example, bit[0], [1], [2], and [3] are set to 1b when RegWrData[7:0], [15:8], [23:16], and [31:24] are valid, respectively.
- 2) To read register, AsyncAxiReg sets the RegRdReq signal to 1b with a valid value for RegAddr. The 32-bit data is returned after the read request is received. The slave detects the RegRdReq signal being set to start the read transaction. In the read operation, the address value (RegAddr) remains unchanged until RegRdValid is set to 1b. The address can then be used to select the returned data using multiple layers of multiplexers.
- 3) The slave returns the read data on RegRdData bus by setting the RegRdValid signal to 1b. After that, AsyncAxiReg forwards the read value to the LAxir\* interface.

### 2.3.2 UserReg

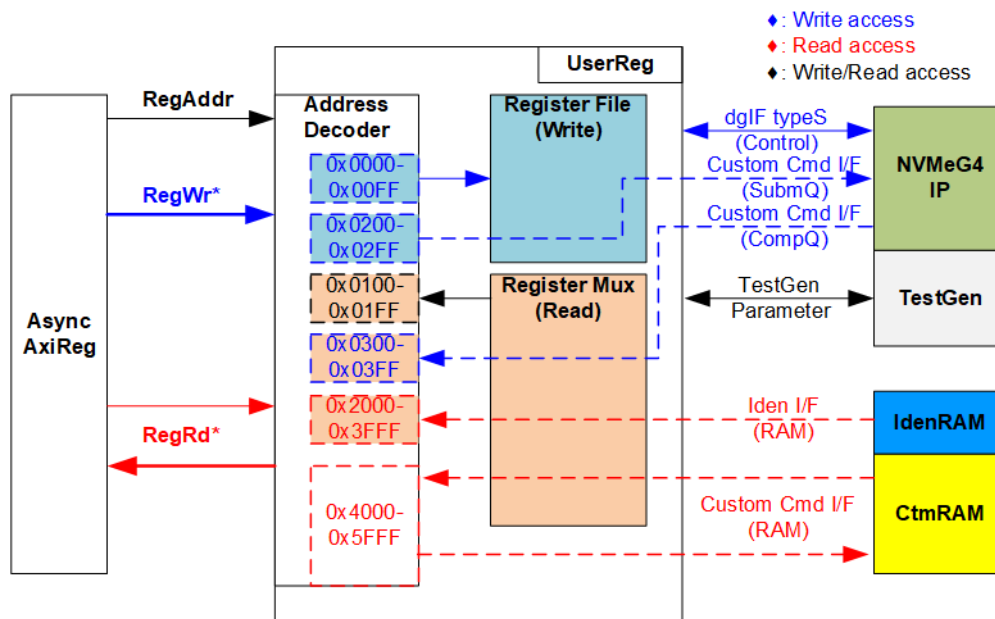


Figure 2-12 UserReg Interface

The UserReg module consists of an Address decoder, a Register File, and a Register Mux. The Address decoder decodes the address requested by AsyncAxiReg and selects the active register for either write or read transactions. The assigned address range in UserReg is divided into six areas, as shown in Figure 2-12.

- 1) 0x0000 – 0x00FF: Mapped to set the command with the parameters of NVMeG4-IP and TestGen. This area is write-access only.
- 2) 0x0200 – 0x02FF: Mapped to set the parameters for the Custom command interface of NVMeG4-IP. This area is write-access only.
- 3) 0x0100 – 0x01FF: Mapped to read the status signals of NVMeG4-IP and TestGen. This area is read-access only.
- 4) 0x0300 – 0x03FF: Mapped to read the status of Custom command interface (NVMeG4-IP). This area is read-access only.
- 5) 0x2000 – 0x3FFF: Mapped to read data from IdenRAM. This area is read-access only.
- 6) 0x4000 – 0x5FFF: Mapped to write or read data using Custom command RAM interface. This area allows both write-access and read access. However, the demo shows only read-access by running the SMART command.

The Address decoder decodes the upper bits of RegAddr to select the active hardware (NVMeG4-IP, TestGen, IdenRAM, or CtmRAM). The Register File within UserReg has a 32-bit bus size, so the write byte enable (RegWrByteEn) is not used in the test system, and the CPU uses a 32-bit pointer to set the hardware register.

To read the register, multi-level multiplexers (mux) select the data to return to the CPU by using the address. The lower bits of RegAddr are fed to the submodule to select the active data from each submodule. While the upper bits are used in UserReg to select the returned data from each submodule. Consequently, the latency time of read data equals to two clock cycles, and RegRdValid is created by RegRdReq, with two D Flip-flops asserted. More details of the address mapping within the UserReg module are shown in Table 2-1.

**Table 2-1 Register Map**

Address Rd/Wr	Register Name (Label in the “nvmeiptest.c”)	Description
<b>0x0000 – 0x00FF: Control signals of NVMeG4-IP and TestGen (Write access only)</b>		
BA+0x0000	User Address (Low) Reg (USRADRL_INTREG)	[31:0]: Input to be bits[31:0] of start address as 512-byte unit (UserAddr[31:0] of dglF typeS)
BA+0x0004	User Address (High) Reg (USRADRH_INTREG)	[15:0]: Input to be bits[47:32] of start address as 512-byte unit (UserAddr[47:32] of dglF typeS)
BA+0x0008	User Length (Low) Reg (USRLENL_INTREG)	[31:0]: Input to be bits[31:0] of transfer length as 512-byte unit (UserLen[31:0] of dglF typeS)
BA+0x000C	User Length (High) Reg (USRLENH_INTREG)	[15:0]: Input to be bits[47:32] of transfer length as 512-byte unit (UserLen[47:32] of dglF typeS)
BA+0x0010	User Command Reg (USRCMD_INTREG)	[2:0]: Input to be user command (UserCmd of dglF typeS for NVMeG4-IP) 000b: Identify, 001b: Shutdown, 010b: Write SSD, 011b: Read SSD, 100b: SMART/Secure Erase, 110b: Flush, 101b/111b: Reserved When this register is written, the command request is sent to NVMeG4-IP to start the operation.
BA+0x0014	Test Pattern Reg (PATTSEL_INTREG)	[2:0]: Select test pattern 000b-Increment, 001b-Decrement, 010b-All 0, 011b-All 1, 100b-LFSR
BA+0x0020	NVMe Timeout Reg (NVMTIMEOUT_INTREG)	[31:0]: Mapped to TimeOutSet[31:0] of NVMeG4-IP
<b>0x0100 – 0x01FF: Status signals of NVMeG4-IP and TestGen (Read access only)</b>		
BA+0x0100	User Status Reg (USRSTS_INTREG)	[0]: UserBusy of dglF typeS (0b: Idle, 1b: Busy) [1]: UserError of dglF typeS (0b: Normal, 1b: Error) [2]: Data verification fail (0b: Normal, 1b: Error)
BA+0x0104	Total disk size (Low) Reg (LBASIZEL_INTREG)	[31:0]: Mapped to LBASize[31:0] of NVMeG4-IP
BA+0x0108	Total disk size (High) Reg (LBASIZEH_INTREG)	[15:0]: Mapped to LBASize[47:32] of NVMeG4-IP [31]: Mapped to LBAMode of NVMeG4-IP
BA+0x010C	User Error Type Reg (USRERRTYPE_INTREG)	[31:0]: Mapped to UserErrorType[31:0] of NVMeG4-IP to show error status
BA+0x0110	PCIe Status Reg (PCIESTS_INTREG)	[7:0]: Unused for NVMeG4-IP [15:8]: Mapped to MACStatus[7:0] of NVMeG4-IP
BA+0x0114	Completion Status Reg (COMPSTS_INTREG)	[15:0]: Mapped to AdmCompStatus[15:0] of NVMeG4-IP [31:16]: Mapped to IOCompStatus[15:0] of NVMeG4-IP
BA+0x0118	NVMe CAP Reg (NVMCAP_INTREG)	[31:0]: Mapped to NVMeCAPReg[31:0] of NVMeG4-IP
BA+0x011C	NVMe Test pin Reg (NVMTESTPIN_INTREG)	[31:0]: Mapped to TestPin[31:0] of NVMeG4-IP
BA+0x0120 – BA+0x012F	MAC Test pin0-3 Reg (MACTESTPIN0-3_INTREG)	The 128-bit of MACTestPin, output signal from NVMeG4-IP 0x0120: Bit[31:0], 0x0124[31:0]: Bit[63:32], ..., 0x012C[31:0]: Bit[127:96]

Address	Register Name	Description
Rd/Wr	(Label in the "nvmeipstest.c")	
<b>0x0100 – 0x01FF: Status signals of NVMeG4-IP and TestGen (Read access only)</b>		
BA+0x0130 – BA+0x014F	Expected value Word0-7 Reg (EXPPATW0-W7_INTREG)	256-bit of the expected data of the 1 <sup>st</sup> failure when executing Read command. 0x0130: Bit[31:0], 0x0134[31:0]: Bit[63:32], ..., 0x014C[31:0]: Bit[255:224]
BA+0x0150 – BA+0x016F	Read value Word0-7 Reg (RDPATW0-W7_INTREG)	256-bit of the read data of the 1 <sup>st</sup> failure when executing Read command. 0x0150: Bit[31:0], 0x0154[31:0]: Bit[63:32], ..., 0x016C[31:0]: Bit[255:224]
BA+0x0170	Data Failure Address(Low) Reg (RDFAILNOL_INTREG)	[31:0]: Bit[31:0] of the byte address of the 1 <sup>st</sup> failure when executing Read command
BA+0x0174	Data Failure Address(High) Reg (RDFAILNOH_INTREG)	[24:0]: Bit[56:32] of the byte address of the 1 <sup>st</sup> failure when executing Read command
BA+0x0178	Current test byte (Low) Reg (CURTESTSIZE_L_INTREG)	[31:0]: Bit[31:0] of the current test data size in TestGen module
BA+0x017C	Current test byte (High) Reg (CURTESTSIZE_H_INTREG)	[24:0]: Bit[56:32] of the current test data size of TestGen module
<b>Other interfaces (Custom command of NVMeG4-IP, IdenRAM, and Custom RAM)</b>		
BA+0x0200 – BA+0x023F	Custom Submission Queue Reg (CTMSUBMQ_STRUCT)	[31:0]: Submission queue entry of SMART, Secure Erase, and Flush command. Input to be CtmSubmDW0-DW15 of NVMeG4-IP. 0x200: DW0, 0x204: DW1, ..., 0x23C: DW15
BA+0x0300 – BA+0x030F	Custom Completion Queue Reg (CTMCOMPQ_STRUCT)	[31:0]: CtmCompDW0-DW3 output from NVMeG4-IP. 0x300: DW0, 0x304: DW1, ..., 0x30C: DW3
BA+0x0800	IP Version Reg (IPVERSION_INTREG)	[31:0]: Mapped to IPVersion[31:0] of NVMeG4-IP
BA+0x2000 – BA+0x2FFF	Identify Controller Data (IDENCTRL_CHARREG)	4KB Identify Controller Data structure
BA+0x3000 – BA+0x3FFF	Identify Namespace Data (IDENNAME_CHARREG)	4KB Identify Namespace Data structure
BA+0x4000 – BA+0x5FFF	Custom command RAM (CTMRAM_CHARREG)	Connect to 8KB CtmRAM interface for storing 512-byte data output from SMART Command.



### 3 CPU Firmware

#### 3.1 Test firmware (nvmeiptest.c)

The CPU follows these steps upon system startup to complete the initialization process.

- 1) Initialize UART and Timer settings.
- 2) Wait for the PCIe connection to become active (PCIESTS\_INTREG[0]=1b).
- 3) Wait for NVMeG4-IP to complete its own initialization process (USRSTS\_INTREG[0]=0b).  
If errors are encountered, the process will stop and display an error message.
- 4) Display the status of the PCIe link, including the number of lanes and the speed, by reading PCIESTS\_INTREG[7:2] status.
- 5) Display the main menu with options to run seven commands for NVMeG4-IP, i.e., Identify, Write, Read, SMART, Flush, Secure Erase, and Shutdown.

More details on the sequence for each command in the CPU firmware are described in the following sections.

##### 3.1.1 Identify Command

The sequence for the firmware when the Identify command is selected by user is as follows.

- 1) Set bits[2:0] of USRCMD\_INTREG to 000b to send the Identify command request to NVMeG4-IP. The busy flag (USRSTS\_INTREG[0]) will then change from 0b to 1b.
- 2) The CPU waits until the operation is completed or an error is detected by monitoring USRSTS\_INTREG[1:0].
  - Bit[0] is de-asserted to 0b when the command is completed. The data of Identify command returned by NVMeG4-IP will be stored in IdenRAM.
  - Bit[1] is asserted to 1b, indicating an error. In this case, the error message will be displayed on the console with details decoded from USRERRTYPE\_INTREG[31:0]. The process will then stop.
- 3) Once the busy flag (USRSTS\_INTREG[0]) is de-asserted to 0b, the CPU proceeds to display information that has been decoded from LBASIZEL/H\_INTREG, which includes the SSD capacity and LBA unit size. Besides, further information, such as the SSD model, can be retrieved from the IdenRAM (IDENCTRL\_CHARREG).

### 3.1.2 Write/Read Command

The sequence for the firmware when the Write/Read command is selected is as follows.

- 1) Receive start address, transfer length, and test pattern from Serial console. If any inputs are invalid, the operation will be cancelled.
 

*Note: If LBA unit size = 4 KB, the start address and transfer length must align to 8.*
- 2) After obtaining all the inputs, set them to USRADRL/H\_INTREG, USRLENL/H\_INTREG, and PATTSEL\_INTREG.
- 3) To execute either the Write or Read command, set bits[2:0] of USRCMD\_INTREG to 010b or 011b, respectively. This sends the command request to the NVMeG4-IP. Once the command is issued, the busy flag of NVMeG4-IP (USRSTS\_INTREG[0]) will change from 0b to 1b.
- 4) The CPU waits until the operation is completed or an error (excluding verification error) is detected by monitoring USRSTS\_INTREG[2:0].
  - Bit[0] is de-asserted to 0b when the command is completed.
  - Bit[1] is asserted to 1b, indicating an error. In this case, the error message will be displayed on the console with details decoded from USRERRTYPE\_INTREG[31:0]. The process will then stop.
  - Bit[2] is asserted when data verification fails. In this case, the verification error message will then be displayed on the console, but the CPU will continue to run until the operation is completed or the user inputs any key to cancel the operation.

While the command is running, the current transfer size, read from CURTESTSIZE\_INTREG, will be displayed every second.

- 5) Once the busy flag (USRSTS\_INTREG[0]) is de-asserted to 0b, CPU will calculate and display the test result on the console including the total time usage, total transfer size, and transfer speed.

### 3.1.3 SMART Command

The sequence of the firmware when the SMART command is selected is as follows.

- 1) The 16-Dword of the Submission Queue entry (CTMSUBMQ\_STRUCT) is set to the SMART command value.
- 2) Set bits[2:0] of USRCMD\_INTREG[2:0] to 100b to send the SMART command request to NVMeG4-IP. The busy flag (USRSTS\_INTREG[0]) will then change from 0b to 1b.
- 3) The CPU waits until the operation is completed or an error is detected by monitoring USRSTS\_INTREG[1:0].
  - Bit[0] is de-asserted to 0b after the operation is finished. The data of SMART command returned by NVMeG4-IP will be stored in CtmRAM.
  - Bit[1] is asserted to 1b, indicating an error. In this case, the error message will be displayed on the console with details decoded from USRERRTYPE\_INTREG[31:0]. The process will then stop.
- 4) After the busy flag (USRSTS\_INTREG[0]) is de-asserted to 0b, the CPU will display information decoded from CtmRAM (CTMRAM\_CHARREG), including Remaining Life, Percentage Used, Temperature, Total Data Read, Total Data Written, Power On Cycles, Power On Hours, and Number of Unsafe Shutdown.

For more information on the SMART log, refer to the NVM Express Specification.  
<https://nvmexpress.org/specifications/>

### 3.1.4 Flush Command

The sequence of the firmware when the Flush command is selected is as follows.

- 1) The 16-Dword of the Submission Queue entry (CTMSUBMQ\_STRUCT) is set to the Flush command value.
- 2) Set bits[2:0] of USRCMD\_INTREG to 110b to send Flush command request to NVMeG4-IP. The busy flag of NVMeG4-IP (USRSTS\_INTREG[0]) will then change from 0b to 1b.
- 3) The CPU waits until the operation is completed or an error is detected by monitoring USRSTS\_INTREG[1:0].
  - Bit[0] is de-asserted to 0b after the operation is finished. The CPU will then return to the main menu.
  - Bit[1] is asserted to 1b, indicating an error. In this case, the error message will be displayed on the console with details decoded from USRERRTYPE\_INTREG[31:0]. The process will then stop.

### 3.1.5 Secure Erase Command

The sequence of the firmware when the Secure Erase command is selected is as follows.

- 1) The 16-Dword of the Submission Queue entry (CTMSUBMQ\_STRUCT) is set to the Secure Erase command value.
- 2) Set NVMTIMEOUT\_INTREG to 0 to disable timer to prevent the timeout error.
- 3) Set bits[2:0] of USRCMD\_INTREG[2:0] to 100b to send Secure Erase command request to NVMeG4-IP. The busy flag of NVMeG4-IP (USRSTS\_INTREG[0]) will then change from 0b to 1b.
- 4) The CPU waits until the operation is completed or an error is detected by monitoring USRSTS\_INTREG[1:0].
  - Bit[0] is de-asserted to 0b after the operation is finished. The CPU will then proceed to the next step.
  - Bit[1] is asserted to 1b, indicating an error. In this case, the error message will be displayed on the console with details decoded from USRERRTYPE\_INTREG[31:0]. The process will then stop.
- 5) After completing the command, the timer is re-enabled to generate timeout error in NVMeG4-IP by setting NVMTIMEOUT\_INTREG to the default value.

### 3.1.6 Shutdown Command

The sequence of the firmware when the Shutdown command is selected is as follows.

- 1) Set bits[2:0] of USRCMD\_INTREG to 001b to send the Shutdown command request to NVMeG4-IP. The busy flag of NVMeG4-IP (USRSTS\_INTREG[0]) will then change from 0b to 1b.
- 2) The CPU waits until the operation is completed or an error is detected by monitoring USRSTS\_INTREG[1:0].
  - Bit[0] is de-asserted to 0b after the operation is finished. The CPU will then proceed to the next step.
  - Bit[1] is asserted to 1b, indicating an error. The error message will be displayed on the console with details decoded from USRERRTYPE\_INTREG[31:0]. The process will then stop.
- 3) After Shutdown command completes, both the SSD and NVMeG4-IP will become inactive and the CPU will be unable to receive any new commands from the user. To continue testing, the user must power off and power on the system.

### 3.2 Function list in Test firmware

int exec_ctm(unsigned int user_cmd)	
Parameters	user_cmd: 4-SMART command, 6-Flush command
Return value	0: No error, -1: Some errors are found in the NVMeG4-IP
Description	Execute SMART command as outlined in section 3.1.3 (SMART Command), or execute Flush command as outlined in section 3.1.4. (Flush Command).

unsigned long long get_cursize(void)	
Parameters	None
Return value	Read value of CURTESTSIZEL/H_INTREG
Description	The value of CURTESTSIZEL/H_INTREG is read and converted to byte units before being returned as the result of the function.

int get_param(userin_struct* userin)	
Parameters	userin: Three inputs from user, i.e., start address, total length in 512-byte unit, and test pattern
Return value	0: Valid input, -1: Invalid input
Description	Receive the input parameters from the user and verify the value. When the input is invalid, the function returns -1. Otherwise, all inputs are updated to userin parameter.

void iden_dev(void)	
Parameters	None
Return value	None
Description	Execute Identify command as outlined in section 3.1.1 (Identify Command).

int setctm_erase(void)	
Parameters	None
Return value	0: No error, -1: Some errors are found in the NVMeG4-IP
Description	Set Secure Erase command to CTMSUBMQ_STRUCT and call exec_ctm function to execute Secure Erase command.

int setctm_flush(void)	
Parameters	None
Return value	0: No error, -1: Some errors are found in the NVMeG4-IP
Description	Set Flush command to CTMSUBMQ_STRUCT and call exec_ctm function to execute Flush command.

int setctm_smart(void)	
Parameters	None
Return value	0: No error, -1: Some errors are found in the NVMeG4-IP
Description	Set SMART command to CTMSUBMQ_STRUCT and call exec_ctm function to execute SMART command. Finally, decode and display SMART information on the console.

void show_error(void)	
Parameters	None
Return value	None
Description	Read USRERRTYPE_INTREG, decode the error flag, and display the corresponding error message. Also, call show_pciestat function to check the hardware's debug signals.

void show_pciestat(void)	
Parameters	None
Return value	None
Description	Read PCIESTS_INTREG until the read value from two read times is stable. After that, display the read value on the console. Also, debug signals are read from NVMTESTPIN_INTREG and MACTESTPIN0-3_INTREG to display on the console.

void show_result(void)	
Parameters	None
Return value	None
Description	Print total transfer size by calling get_cursize and show_size functions. After that, calculate total time usage from global parameters (timer_val and timer_upper_val) and then display in usec, msec, or sec unit. Finally, transfer performance is calculated and displayed in MB/s unit.

void show_size(unsigned long long size_input)	
Parameters	size_input: Transfer size to display on the console
Return value	None
Description	Calculate and display the input value in MB or GB unit.

void show_smart_hex16byte(volatile unsigned char *char_ptr)	
Parameters	*char_ptr: Pointer of 16-byte SMART data
Return value	None
Description	Display 16-byte SMART data as hexadecimal unit.

void show_smart_int8byte(volatile unsigned char *char_ptr)	
Parameters	*char_ptr: Pointer of 8-byte SMART data
Return value	None
Description	When the input value is less than 4 billion (32-bit), the 8-byte SMART data is displayed in decimal units. If the input value exceeds this limit, an overflow message is displayed.

void show_smart_size8byte(volatile unsigned char * char_ptr)	
Parameters	*char_ptr: pointer of 8-byte SMART data
Return value	None
Description	Display 8-byte SMART data as GB or TB unit. When the input value is more than limit (500 PB), the overflow message is displayed instead.

void show_vererr(void)	
Parameters	None
Return value	None
Description	Read RDFAILNOL/H_INTREG (error byte address), EXPPATW0-W7_INTREG (expected value), and RDPATW0-W7_INTREG (read value) to display verification error details on the console.

void shutdown_dev(void)	
Parameters	None
Return value	None
Description	Execute Shutdown command as outlined in section 3.1.6 (Shutdown Command).

int wrdd_dev(unsigned int user_cmd)	
Parameters	user_cmd: 2-Write command, 3-Read command
Return value	0: No error, -1: Receive invalid input or some errors are found.
Description	Execute Write command or Read command as outlined in section 3.1.2 (Write/Read Command). In this function, 'show_result' is called to compute and display transfer performance in Write/Read command.

## 4 Example Test Result

The performance results of executing Write and Read commands on the demo system using a TB Samsung 990 Pro and the NVMeG4-IP with 1 MB buffer are shown in Figure 4-1. The test utilizes an LFSR pattern for data and specifies a transfer size of 64 GB.




Figure 4-1 Test Performance of NVMeG4-IP demo using Samsung 990 Pro SSD

When using the ZCU102 board, the measured write performance reached approximately 7,500 MB/s, while the read performance achieved approximately 7,000 MB/s.



## 5 Revision History

Revision	Date	Description
1.04	22-Dec-23	Add Secure Erase feature
1.03	11-Jul-23	Update register map, and performance result
1.02	19-Jul-21	Add details, update register map, and performance result
1.01	21-Dec-20	Update performance result
1.00	30-Aug-19	Initial Release

Copyright: 2019 Design Gateway Co,Ltd.