

NVMe IP for Gen4/NVMeG4-IP demo instruction

Rev1.5 3-Aug-22

1 Overview

This document describes the instruction to run NVMe-IP for Gen4/NVMeG4-IP demo on FPGA development board. One NVMe SSD is accessed for running the demo. There are six commands that can run, i.e., Identify, Write, Read, SMART, Flush, and Shutdown command. User controls test operation via FPGA console.

After user finishes FPGA board setup following “dg_nvmeip_fpgasetup” document, main menu is displayed and the user sets the input to the console for selecting test operation.

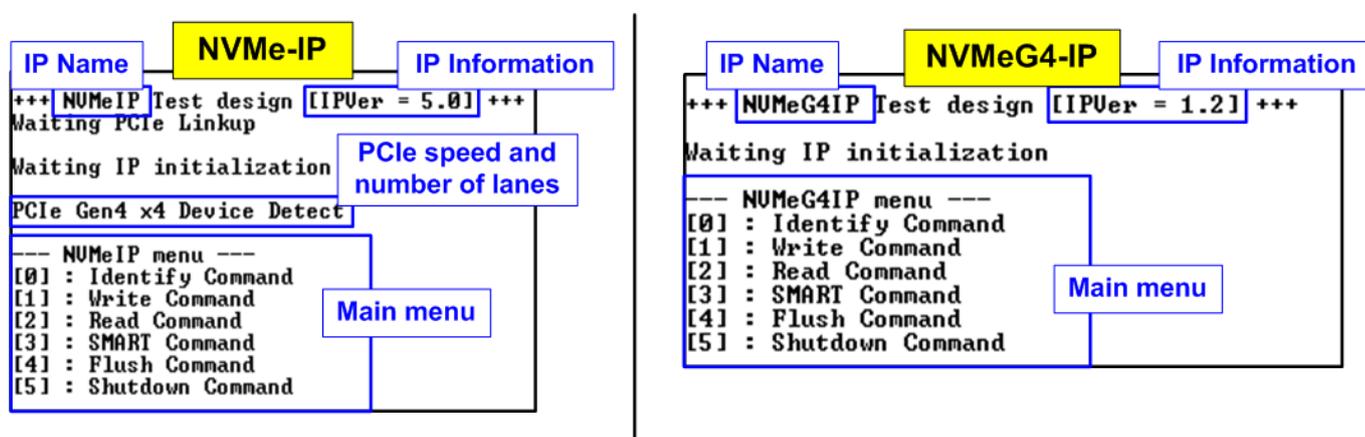


Figure 1-1 Main menu after finishing initialization

On welcome screen, IP name and IP version number are displayed. For standard IP (NVMe-IP), the PCIe speed and number of PCIe lanes are displayed after IP finishing initialization. While NVMeG4-IP does not display because only 4-lane PCIe Gen4 SSD is supported. Finally, the test menu is displayed on the console.

2 Test Menu

2.1 Identify Command

Select '0' to send Identify command to NVMe SSD.

```

+++ Identify Command selected +++
Model Number      : addlink M.2 PCIE G4x4 NVMe
SSD Capacity      = 2000[GB]
Data size per LBA = 512[Byte]

--- NVMeIP menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command

```

Model name, SSD Capacity, and LBA unit
(Output from Identify command)

Figure 2-1 Test result when running Identify command

After finishing the operation, the SSD information output from Identify command is displayed. The console shows three values.

- 1) SSD model number : This value is decoded from Identify controller data.
- 2) SSD capacity : This value is signal output from NVMe(G4)-IP.
- 3) Data size per LBA : This value is signal output from NVMe(G4)-IP. Two values are supported, i.e., 512 byte and 4 Kbyte.

2.2 Write Command

Select '1' to send Write command to NVMe SSD.

◆ : User input
◆ : User output

Normal

```

+++ Write Command selected +++
Enter Start Address (512 Byte)      : 0 - 0xE8E088AF => 0
Enter Length (512 Byte)           : 1 - 0xE8E088B0 => 0x40000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR=> 2
5.674 [GB]
11.347 [GB]
17.088 [GB]
22.760 [GB]
28.499 [GB]
34.171 [GB]
Total = 34.359 [GB] , Time = 6032[ms] , Transfer speed = 5696[MB/s]

--- NUmEIP menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command
    
```

Current transfer size
Input parameters
Output performance

Figure 2-2 Input and test result when running Write command

User inputs three parameters as follows.

- 1) Start Address: Input start address to write SSD as 512-byte unit. The input is decimal unit when user inputs only digit number. User can add "0x" to be a prefix for hexadecimal unit. When LBA unit of SSD is 4 Kbyte, this input must be aligned to 8.
- 2) Transfer Length: Input total transfer size as 512-byte unit. The input is decimal unit when user inputs only digit number. User can add "0x" to be a prefix for hexadecimal unit. When LBA unit of SSD is 4 Kbyte, this input must be aligned to 8.
- 3) Test pattern: Select test data pattern for writing to SSD. There are five patterns, i.e., 32-bit incremental, 32-bit decremental, all 0, all 1, and 32-bit LFSR counter.

When all inputs are valid, the operation begins. While writing data, current amount of write data is displayed on the console every second to show that system is still alive. Finally, total size, total time usage, and test speed are displayed on the console as a test result.

Note:

Some SSDs shows slower performance after writing large size data to SSD. It needs to recover SSD performance by filling zero pattern or using Format command which is customized command. Please contact our sales for more information if Format command is required.

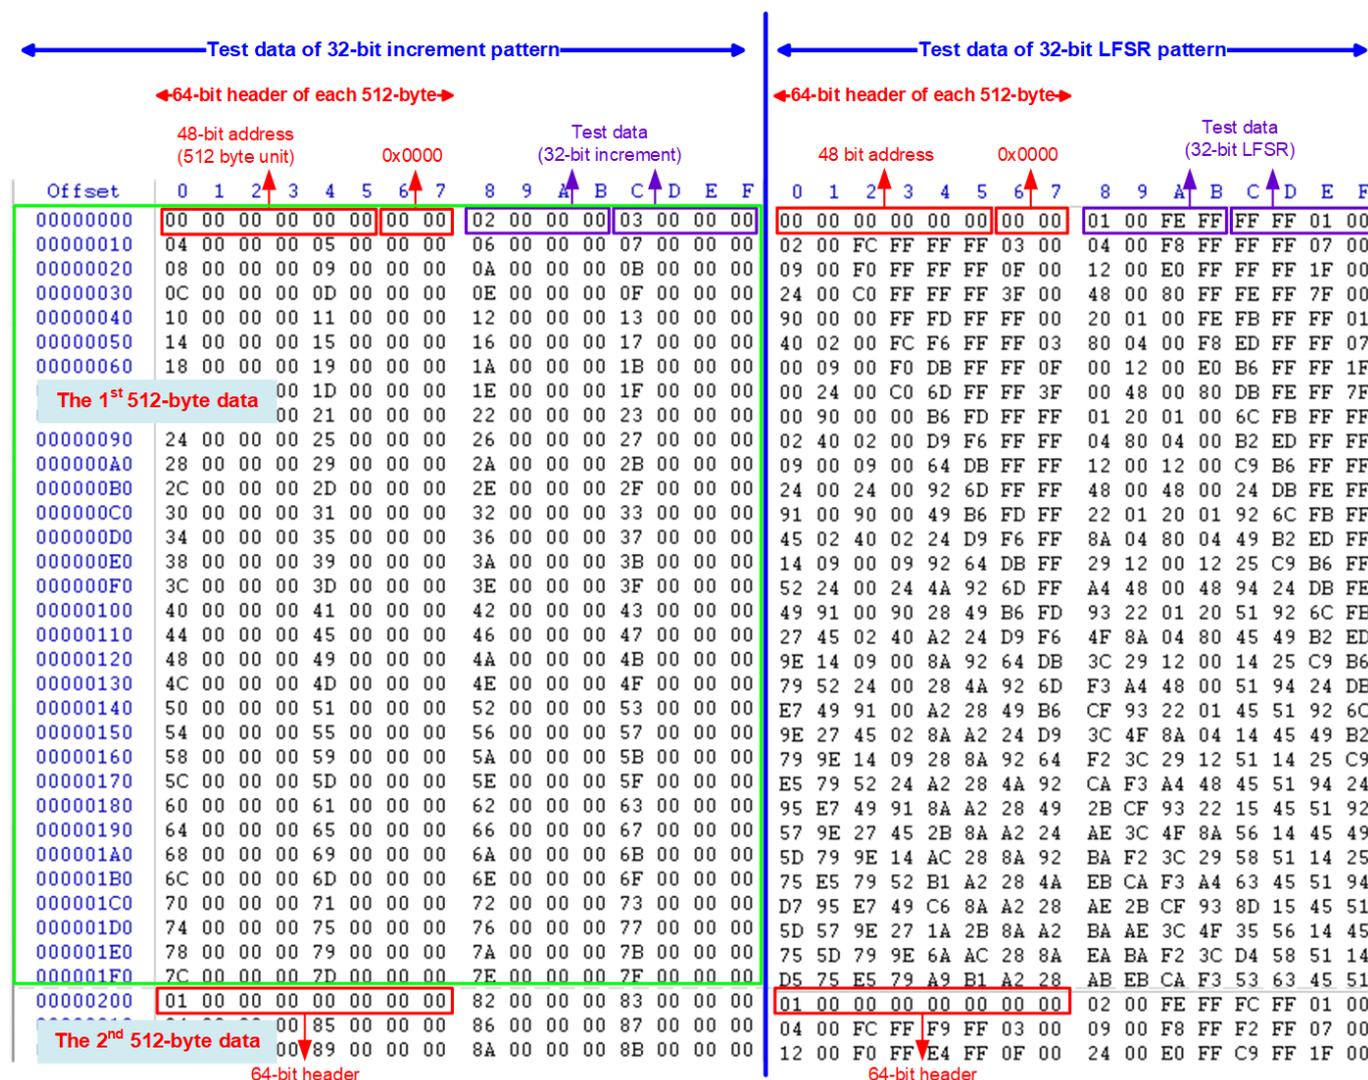


Figure 2-3 Example Test data of the 1st and 2nd 512 byte by using increment/LFSR pattern

Test data in SSD is split into 512-byte unit. For incremental, decremental, and LFSR pattern, each 512-byte data has unique 64-bit header consisting of 48-bit address (in 512-byte unit) and 16-bit zero value. The data after 64-bit header is the test pattern which is selected by user.

The left window of Figure 2-3 shows the example when using 32-bit incremental pattern while the right window shows the example when using 32-bit LFSR pattern. The unique header is not included when running all-0 or all-1 pattern.

When user runs Write or Read command with 4-Kbyte LBA SSD, there is the message displayed on the console to show the input limitation which must be aligned to 8 as shown in Figure 2-4. When the input is not aligned to 8, "Invalid input" is displayed and the operation is cancelled.

Figure 2-5 shows the example when the input is out of the recommended range for each parameter. The console displays "Invalid input" and then the operation is cancelled.

LBA alignment error

```

+++ Write Command selected +++
Please input [Start Address] and [Length] in unit of 8
Enter Start Address (512 Byte) : 0 - 0x2E9390AF => 8
Enter Length (512 Byte) : 1 - 0x2E9390A8 => 7
Invalid input
--- NUMeIP me
[0] : Identif
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command
    
```

Recommended message when LBA unit = 4 Kbyte

When LBA unit = 4 KB and length is not aligned to 8, error message is displayed.

Figure 2-4 Error message when the input is unaligned for 4-Kbyte LBA SSD

Error input

```

+++ Write Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E088AF => 0xFFFFFFFF
Invalid input
+++ Write Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E088AF => 0
Enter Length (512 Byte) : 1 - 0xE8E088B0 => 0xFFFFFFFF
Invalid input
+++ Write Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E088AF => 0
Enter Length (512 Byte) : 1 - 0xE8E088B0 => 0x4000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR=> 5
Invalid input
    
```

Out of range address

Out of range Length

Invalid pattern

Figure 2-5 Error message from the invalid input

2.3 Read Command

Select '2' to send Read command to NVMe SSD.

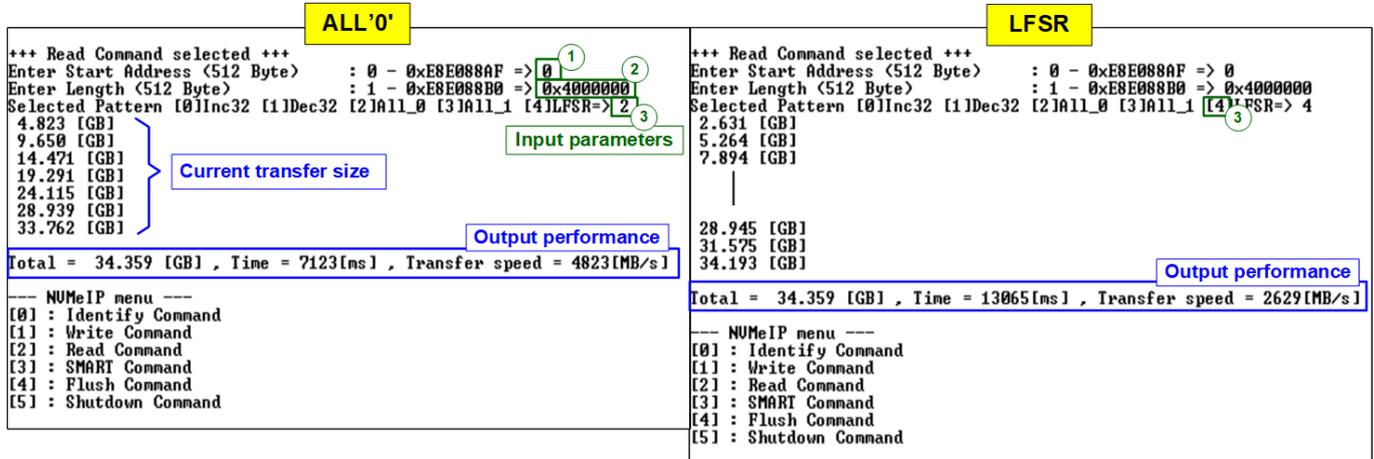


Figure 2-6 Input and test result when running Read Command

User inputs three parameters as follows.

- 1) Start Address: Input start address to read SSD as 512-byte unit. The input is decimal unit when user inputs only digit number. User can add “0x” to be a prefix for hexadecimal unit. When LBA unit of SSD is 4 Kbyte, this input must be aligned to 8.
- 2) Transfer Length: Input total transfer size as 512-byte unit. The input is decimal unit when user inputs only digit number. User can add “0x” to be a prefix for hexadecimal unit. When LBA unit of SSD is 4 Kbyte, this input must be aligned to 8.
- 3) Test pattern: Select test data pattern to verify data from SSD. Test pattern must be matched with the pattern using in Write Command menu. There are five patterns, i.e., 32-bit incremental, 32-bit decremental, all 0, all 1, and 32-bit LFSR counter.

Similar to Write command menu, test system reads data from SSD when all inputs are valid. While reading data, current amount of read data is displayed on the console every second to show that system is still alive. Total size, total time usage, and test speed are displayed after finishing the operation.

“Invalid input” is displayed when some inputs are invalid or unaligned to 8 (when connecting to 4-KB LBA SSD).

Note:

- 1) Some SSDs shows the different performance when changing test pattern. As shown in Figure 2-6, the best read performance is achieved by using all-zero pattern. While the performance is reduced when using LFSR pattern.
- 2) When using customized IPcore which extends the buffer size to 1 Mbyte instead of 256 Kbyte, the read performance is increased to 6396 Mbyte/s as shown in Figure 2-7.

ALL '0' by 1 Mbyte buffer

```

+++ Read Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E088AF => 0
Enter Length (512 Byte) : 1 - 0xE8E088B0 => 0x4000000
Selected Pattern [0]Inc32 [1]Dec32 [2]A11_0 [3]A11_1 [4]LFSR=> 2
6.393 [GB]
12.790 [GB]
19.188 [GB]
25.586 [GB]
31.984 [GB]

Total = 34.359 [GB] , Time = 5371[ms] , Transfer speed = 6396[MB/s]
    
```

Figure 2-7 Test result when running Read command by customized IP (1 Mbyte buffer)

Figure 2-8 shows error message when data verification is failed. "Verify fail" is displayed with the information of the 1st failure data, i.e., the error byte address, the expected value, and the read value.

User can press any key(s) to cancel read operation. Otherwise, the operation is still run until finishing Read command. After that, the output performance is displayed on the console.

When cancelling the operation, the Read command still runs as the background and may not finish in a good sequence. It is recommended to power-off/on FPGA board and adapter board (if connected).

Verification error without cancellation	Verification error with cancellation
<pre> +++ Read Command selected +++ Enter Start Address (512 Byte) : 0 - 0xE8E088AF => 0 Enter Length (512 Byte) : 1 - 0xE8E088B0 => 0x4000000 Selected Pattern [0]Inc32 [1]Dec32 [2]A11_0 [3]A11_1 [4]LFSR=> 1 Verify fail 1st Error at Byte Addr = 0x00000000 Expect Data[255:128] = 0xFFFFFFFF8_FFFFFFFF9_FFFFFFFFA_FFFFFFFFB Expect Data[127: 0] = 0xFFFFFFFFC_FFFFFFFFD_00000000_00000000 Read Data[255:128] = 0x00000000_00000000_00000000_00000000 Read Data[127: 0] = 0x00000000_00000000_00000000_00000000 Press any key to cancel operation 4.774 [GB] 9.552 [GB] 14.332 [GB] 19.110 [GB] 23.888 [GB] 28.666 [GB] 33.443 [GB] Total = 34.359 [GB] , Time = 7191[ms] , Transfer speed = 4777[MB/s] --- NVMeIP menu --- [0] : Identify Command [1] : Write Command [2] : Read Command [3] : SMART Command [4] : Flush Command [5] : Shutdown Command </pre>	<pre> +++ Read Command selected +++ Enter Start Address (512 Byte) : 0 - 0xE8E088AF => 0 Enter Length (512 Byte) : 1 - 0xE8E088B0 => 0x4000000 Selected Pattern [0]Inc32 [1]Dec32 [2]A11_0 [3]A11_1 [4]LFSR=> 1 Verify fail 1st Error at Byte Addr = 0x00000000 Expect Data[255:128] = 0xFFFFFFFF8_FFFFFFFF9_FFFFFFFFA_FFFFFFFFB Expect Data[127: 0] = 0xFFFFFFFFC_FFFFFFFFD_00000000_00000000 Read Data[255:128] = 0x00000000_00000000_00000000_00000000 Read Data[127: 0] = 0x00000000_00000000_00000000_00000000 Press any key to cancel operation 4.776 [GB] Operation is cancelled Please reset system before starting a new test --- NVMeIP menu --- [0] : Identify Command [1] : Write Command [2] : Read Command [3] : SMART Command [4] : Flush Command [5] : Shutdown Command </pre>
<p>Wrong pattern</p> <p>Message when verification is failed</p> <p>Output performance</p>	<p>User enters some keys to cancel the operation</p> <p>Message when operation is cancelled</p>

Figure 2-8 Data verification is failed

2.4 SMART Command

Select '3' to send SMART command to NVMe SSD.

```

+++ SMART Command selected +++
<< Health Status >>
Remaining Life : 91%
<< SMART Log Information >>
Percentage Used : 9%
Temperature : 29 Degree Celsius
Total Data Read : 131062 GB
Total Data Read (Raw data) : 0x00000000_00000000_00000000_0F41BB11
Total Data Written : 146475 GB
Total Data Written (Raw data) : 0x00000000_00000000_00000000_110D0A6E
Power On Cycles : 713 Times
Power On Hours : 78 Hours
Unsafe Shutdowns : 410 Times

SMART Command Complete

--- NUMeIP menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command

```

Data output decoded from SMART command

Figure 2-9 Test result when running SMART command

After finishing the operation, SMART/Health Information (output from SMART command) is displayed as shown in Figure 2-9. The console shows Health status and SMART log information. Health status shows the remaining life of the SSD in percent unit which is calculated from Percentage Used in the SMART log information.

The SMART log information shows seven parameters as follows.

- 1) Percentage used: Display SSD usage in percent unit.
- 2) Temperature in °C unit.
- 3) Total Data Read decoded as GB/TB unit. Also, raw data without decoding is displayed by 32 digits of hex number (128 bits). The unit size of raw data is 512,000 bytes.
- 4) Total Data Written decoded as GB/TB unit. Also, raw data without decoding is displayed by 32 digits of hex number (128 bits). The unit size of raw data is 512,000 bytes.
- 5) Power On Cycles: Display the number of power cycles.
- 6) Power On Hours: Display the period of time in hours to show how long the SSD has been powered on.
- 7) Unsafe Shutdowns: Display the number of unsafe shutdowns of SSD.

2.5 Flush Command

Select '4' to send Flush command to NVMe SSD.

```

+++ Flush Command selected +++
Flush Command Complete
--- NUMeIP menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command
  
```

Message after finishing the operation

Figure 2-10 Test result when running Flush command

“Flush Command Complete” is displayed after finishing Flush operation.

2.6 Shutdown Command

Select '5' to send Shutdown command to NVMe SSD.

```

--- NUMeIP menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command
+++ Shutdown Command selected +++
Are you sure you want to shutdown the device now ?
Press 'y' to confirm : y
Shutdown command is complete
The device has turned off...
  
```

Confirmation message

Press 'y' to confirm

Last message before the IP and SSD are inactive

Figure 2-11 Shutdown Command with confirmation

The confirmation message is displayed on the console. User enters 'y' or 'Y' to continue the operation or enters other keys to cancel the operation.

After finishing Shutdown operation, “Shutdown command is complete” is displayed on the console as the last message. Main menu is not displayed anymore. User needs to power off/on test system to start new test operation.

3 Revision History

Revision	Date	Description
1.5	3-Aug-22	Update read performance by 1 Mbyte buffer
1.4	14-Sep-21	Support NVMe-IP for Gen4
1.3	20-Jul-21	Update SMART info and test performance
1.2	21-Dec-20	Remove FPGA setup from the document
1.1	20-Apr-20	Remove power adapter cable from AB18
1.0	29-Jan-20	Initial version release