



rmNVMe-IP for Gen4 Demo Instruction

Rev1.1 28-Jun-23

1	Overview.....	2
2	Test Menu	3
2.1	Identify Command.....	3
2.2	Write/Read Command	4
2.2.1	Write command	4
2.2.2	Read Command	8
2.2.3	Mixed Write/Read Command	10
2.3	SMART Command.....	12
2.4	Flush Command	13
2.5	Shutdown Command	13
3	Revision History.....	14

1 Overview

This guide provides instructions on how to use the rmNVMe-IP for Gen4 demo on an FPGA development board to access an NVMe Gen4 SSD. The demo supports performing Identify, Write, Read, SMART, Flush, and Shutdown commands. It allows the user to control test operation through FPGA console.

To get started, follow the instructions in the “dg_nvmeip_fpgasetup” document to set up the FPGA board and the test system. Once setup is complete, the rmNVMeIP for Gen4 menu will be displayed and the user can select the desired test operation through the FPGA console.

```

+++ rmNUMeIP for Gen4 Test design [IPVer = 1.1] +++
Waiting PCIe Linkup
Waiting IP initialization
PCIe Gen4 x4 Device Detect
--- rmNUMeIP for Gen4 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Shutdown Command
  
```

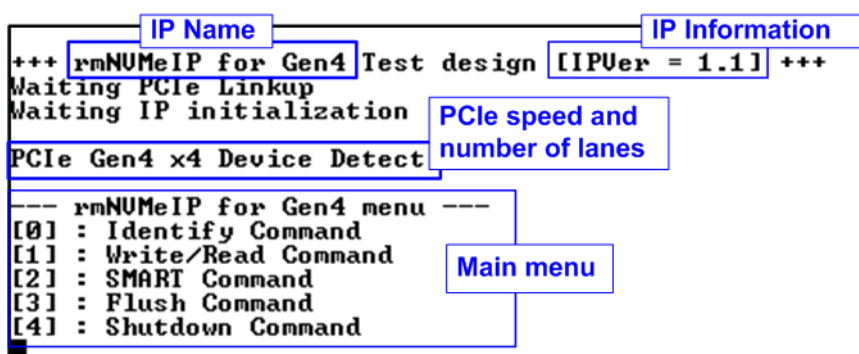


Figure 1-1 rmNVMe-IP for Gen4 menu

As shown in Figure 1-1, the welcome screen shows IP name and IP version number. After that, The PCIe speed, the number of PCIe lanes, and the test menu are 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      : INTEL SSDPF21Q800GB
SSD Capacity      = 800[GB]
--- rmNUMeIP for Gen4 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Shutdown Command
  
```

Model name and SSD Capacity
(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 two values.

- 1) SSD model number: This value is decoded from Identify controller data.
- 2) SSD capacity: This value is signal output from rmNVMe-IP.

```

+++ Identify Command selected +++
User#0 Error Detect
ErrorType = 0x00000040
ErrorType[6]: LBA does not support
PCIe Status = 0x0F2D
User#0 TestPin[47:0] = 0x0000_0000_001F
User#1 TestPin[15:0] = 0x0000
  
```

Error message when LBA unit
detected does not support

Figure 2-2 Error when LBA unit does not support

When the SSD operates in LBA unit which is not equal to 512 bytes, the error message is displayed on the console, as shown in Figure 2-2. After that, the system enters to the inactive status and user needs to reset the system to restart the demo.

2.2 Write/Read Command

Select '1' to send Write and Read command request to rmNVMme-IP via Write interface and Read interface, respectively. The parameters of Write command and Read command can be configured individually. Additionally, there is the feature to limit the maximum data rate for Write/Read command, expressed as a percentage. This allows the user to evaluate the performance balance and relationship between the Write and Read commands of each SSD. Furthermore, the menu also provides the option to disable either the Write or Read, enabling the user to test the peak performance of the system when running only the Write or Read command.

2.2.1 Write command

This topic shows an example when operating only Write command.

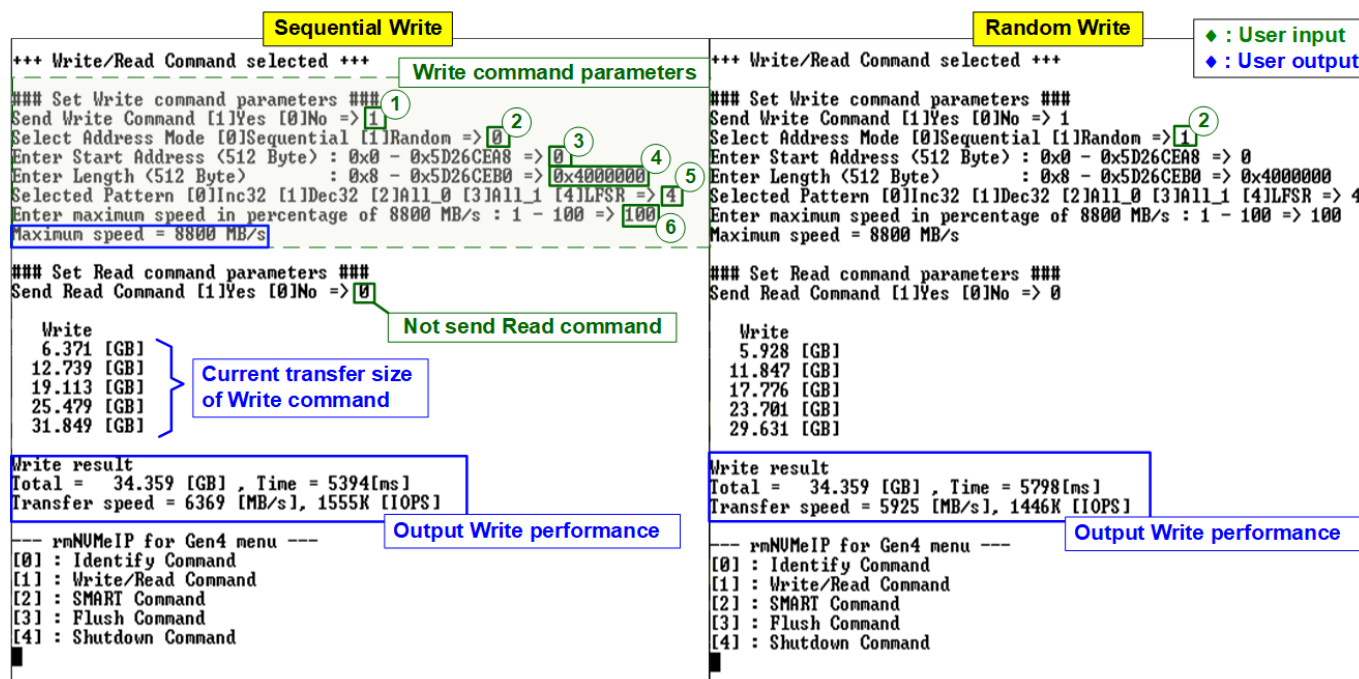


Figure 2-3 Test result of Write command using two Address modes

There are six input parameters for the Write menu, described as follows.

- 1) Send Write Command: Yes or No. If 'No' is selected, the write operation is disabled.
- 2) Address Mode: Sequential or Random mode. The initial address is assigned by Start Address input while the next addresses are calculated based on the selected mode. In sequential mode, the addresses are increased in a continuous sequence. In Random mode, addresses are generated using LFSR equation.
- 3) Start Address: Input start address to write SSD as 512-byte unit. The input is decimal unit when user enters only digit number. User can add "0x" to be prefix for hexadecimal unit. This input must be aligned to 8 for 4-Kbyte alignment.
- 4) Transfer Length: Input total transfer size as 512-byte unit. The input is decimal unit when user enters only digit number. User can add "0x" to be prefix for hexadecimal unit. This input must be aligned to 8 for 4-Kbyte alignment.

- 5) 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.
- 6) Maximum Speed: Setting maximum speed in percentage unit of 8,800 MB/s. Valid from 1 to 100.

Note: 8,800 MB/s is calculated by user clock frequency (275 MHz) x data width (256-bit). This resulting value exceeds the maximum throughput of 4-lane PCIe Gen4 speed.

When all inputs are valid, the operation starts. During the operation, the current amount of write data is displayed on the console every second to indicate that the system is functioning. Finally, total size, total time usage, test speed, and IOPS of the Write command are displayed on the console after the operation is done.

Figure 2-3 shows the results of the Write command when executing sequential and random access. The write performance of the sequential access, at 6369 MB/s, is slightly higher than that of the random access, at 5925 MB/s. However, the performance of random access in most SSDs is typically much less than that of the sequential access. The user can investigate each SSD characteristics by using this menu.

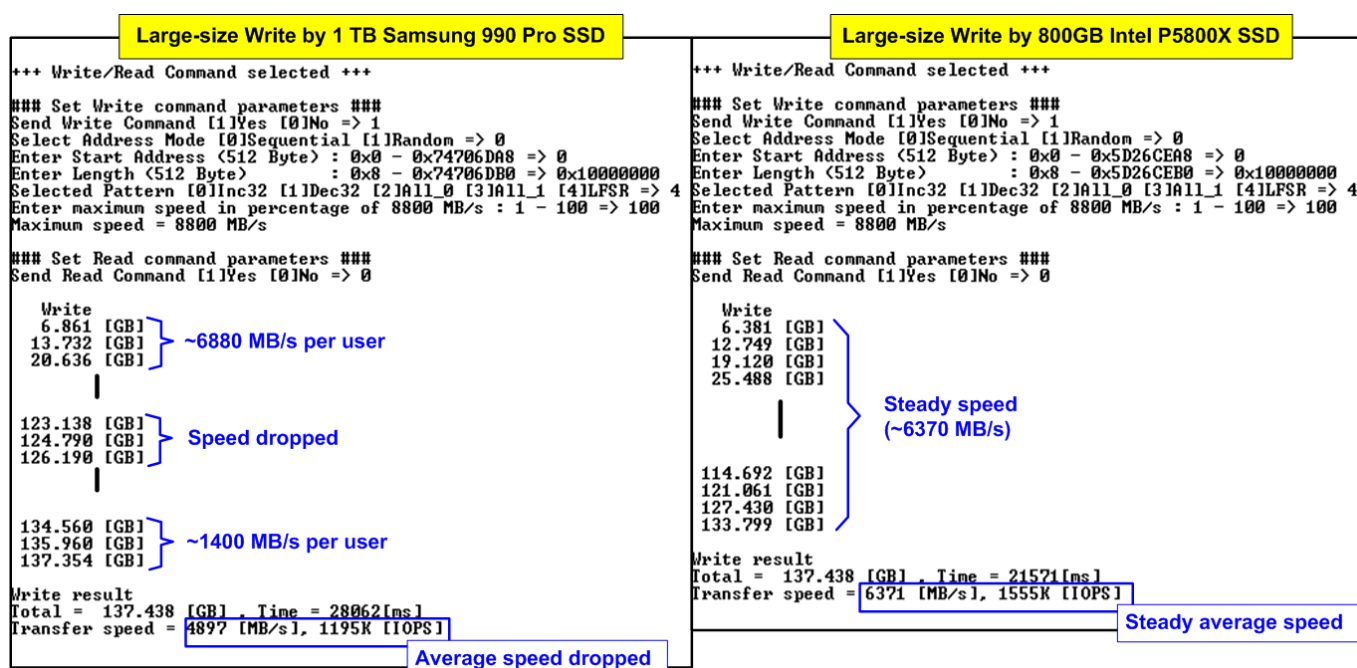


Figure 2-4 Write performance when using large transfer size

Most SSDs incorporate a cache to showcase impressive Write command performance. However, when transferring large amount of data that fill up the cache, performance tends to decrease. In Figure 2-4, the Write performance results of two different SSDs are illustrated. The SSD on the left initially exhibits higher performance, reaching 6880 MB/s per user. However, as data is continuously written to the SSD, the performance drops to 1400 MB/s. On the other hand, the right-side SSD starts with slightly lower performance but maintains a stable transfer rate of 6370 MB/s throughout the test. It is recommended for users to verify the SSD characteristics to ensure they align with the system requirements.

Note: Some SSDs may show slower write performance after large size data is written. To address this issue, running commands such as Write zero, Format, or Sanitized may help recover performance. Please contact our sales team for this customized feature.



Figure 2-5 Example Test data of the 1st and 2nd 4Kbyte data by using incremental/LFSR pattern

Test data in SSD is split into 4096-byte (4K-byte) units. For incremental, decremental, and LFSR patterns, each 4K-byte data has a unique 64-bit header that consists of a 48-bit address in 512-byte units and a 16-bit zero value. The data following the 64-bit header is the test pattern selected by the user.

The left window of Figure 2-5 shows an example of the 32-bit incremental pattern, while the right window displays an example of the 32-bit LFSR pattern. When running all-0 or all-1 patterns, the unique header is not included.

Figure 2-6 shows an example of the inputs that are out of the recommended range for each parameter. The console displays “Invalid input” and then the operation is cancelled.

Error input

```

+++ Write/Read Command selected +++

### Set Write command parameters ###
Send Write Command [1]Yes [0]No => 6
Invalid input

### Set Write command parameters ###
Send Write Command [1]Yes [0]No => 1
Select Address Mode [0]Sequential [1]Random => 3
Invalid input

### Set Read command parameters ###
Send Read Command [1]Yes [0]No => 1
Select Address Mode [0]Sequential [1]Random => 0
Enable verification [0]Disable [1]Enable => 2
Invalid input

### Set Write command parameters ###
Send Write Command [1]Yes [0]No => 1
Select Address Mode [0]Sequential [1]Random => 0
Enter Start Address (512 Byte) : 0x0 - 0x5D26CEA8 => 0xFFFFFFFF
Invalid input

### Set Write command parameters ###
Send Write Command [1]Yes [0]No => 1
Select Address Mode [0]Sequential [1]Random => 0
Enter Start Address (512 Byte) : 0x0 - 0x5D26CEA8 => 0
Enter Length (512 Byte) : 0x8 - 0x5D26CEB0 => 6
Invalid input

### Set Write command parameters ###
Send Write Command [1]Yes [0]No => 1
Select Address Mode [0]Sequential [1]Random => 0
Enter Start Address (512 Byte) : 0x0 - 0x5D26CEA8 => 0
Enter Length (512 Byte) : 0x8 - 0x5D26CEB0 => 0x40000000
Selected Pattern [0]Inc32 [1]Dec32 [2]A11_0 [3]A11_1 [4]LFSR => 9
Invalid input

### Set Write command parameters ###
Send Write Command [1]Yes [0]No => 1
Select Address Mode [0]Sequential [1]Random => 0
Enter Start Address (512 Byte) : 0x0 - 0x5D26CEA8 => 0
Enter Length (512 Byte) : 0x8 - 0x5D26CEB0 => 0x40000000
Selected Pattern [0]Inc32 [1]Dec32 [2]A11_0 [3]A11_1 [4]LFSR => 4
Enter maximum speed in percentage of 11200 MB/s : 1 - 100 => 101
Invalid input
    
```

Figure 2-6 Error message from the invalid input

2.2.2 Read Command

This topic shows an example when operating only Read command.

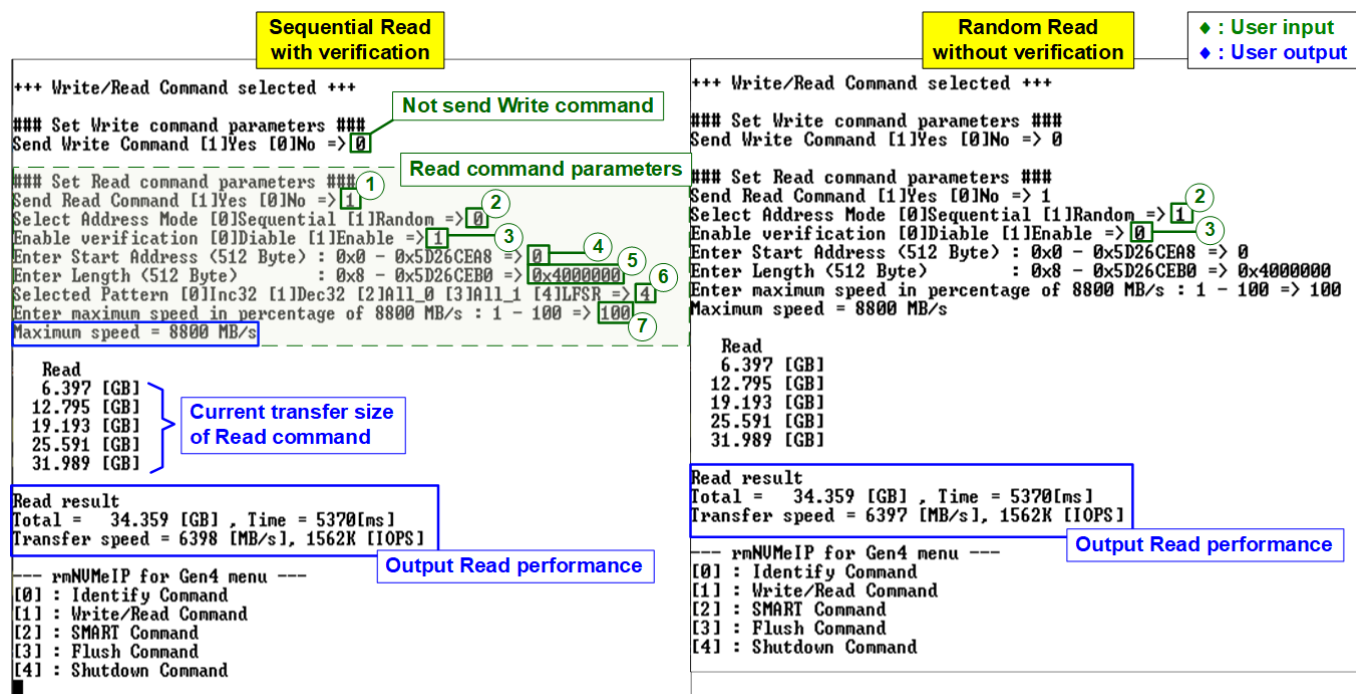


Figure 2-7 Test result of Read command by using two Address modes

The Read command operation requires six or seven parameters, described as follows.

- 1) Send Read Command: Yes or No. If 'No' is selected, the read operation is disabled.
- 2) Address Mode: Sequential or Random mode. The initial address is assigned by Start Address input while the next addresses are calculated based on the selected mode. In sequential mode, the addresses are increased in a continuous sequence. In Random mode, addresses are generated using LFSR equation. Both modes generate the same test data for each memory address, but the order of addresses is different. If the user does not write data to the SSD by full-disk size, the unwritten space will be different for each mode. To verify the written data inside the SSD in such cases, the address mode used in Both the Write menu and Read menu must be the same.
- 3) Enable Verification: Enable or disable data verification for the received data.
- 4) Start Address: Input the start address to read SSD as 512-byte unit. The input is decimal unit when user enters only digit number. User can add "0x" to be prefix for hexadecimal unit. This input must be aligned to 8 for 4-Kbyte alignment.
- 5) Transfer Length: Input total transfer size as 512-byte unit. The input is decimal unit when user enters only digit number. User can add "0x" to be prefix for hexadecimal unit. This input must be aligned to 8 for 4-Kbyte alignment.
- 6) Test pattern (Available when Verification is enabled): Select the test data pattern to verify the received data from the SSD. There are five patterns: 32-bit incremental, 32-bit decremental, all 0, all 1, and 32-bit LFSR counter.
- 7) Maximum Speed: Setting maximum speed in percentage unit of 8,800 MB/s. Valid from 1 to 100.

Note: 8,800 MB/s is calculated by user clock frequency (275 MHz) x data width (256-bit).

When all inputs are valid, the operation starts. During the operation, the current amount of received data is displayed on the console every second to indicate that the system is functioning. In the event of invalid input or failed data verification, an error message will be displayed instead. Once the operation is complete, the console will present the total size, total time usage, test speed, and IOPS of the Read command.

Figure 2-8 shows an example of error message when data verification fails. The message “Verify fail” is displayed with the information of the 1st failure data including the error byte address, the expected value, and the read value.

User can simply press any key(s) to cancel read operation. However, if no cancellation command is initiated, the operation will continue until the Read command is finished. After completion, the console presents the output performance.

If the operation is cancelled, the Read command will still run as a background process and may not finish in a good sequence. In such cases, it is recommended to power off and then power on both the FPGA board and the adapter board (if connected) to ensure proper functionality.

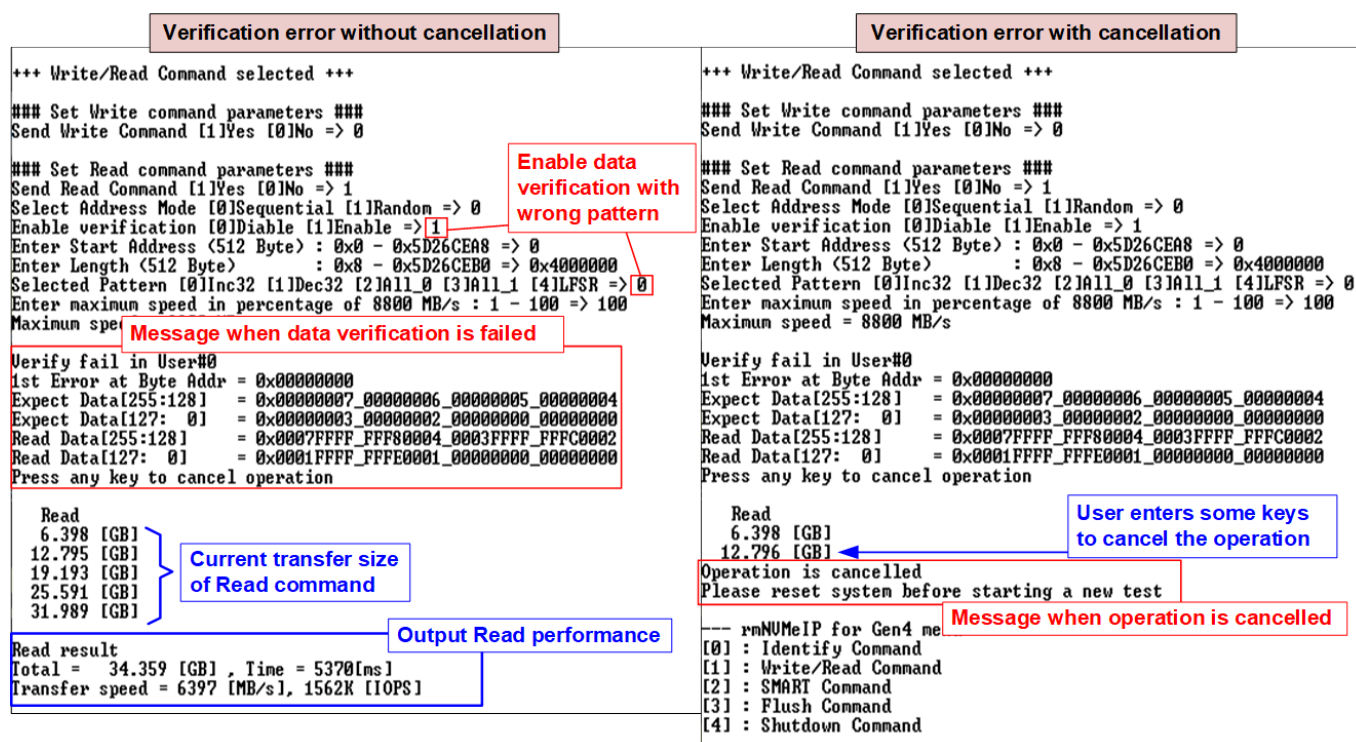


Figure 2-8 Data verification is failed

2.2.3 Mixed Write/Read Command

This section shows an example of operating both Write and Read commands.

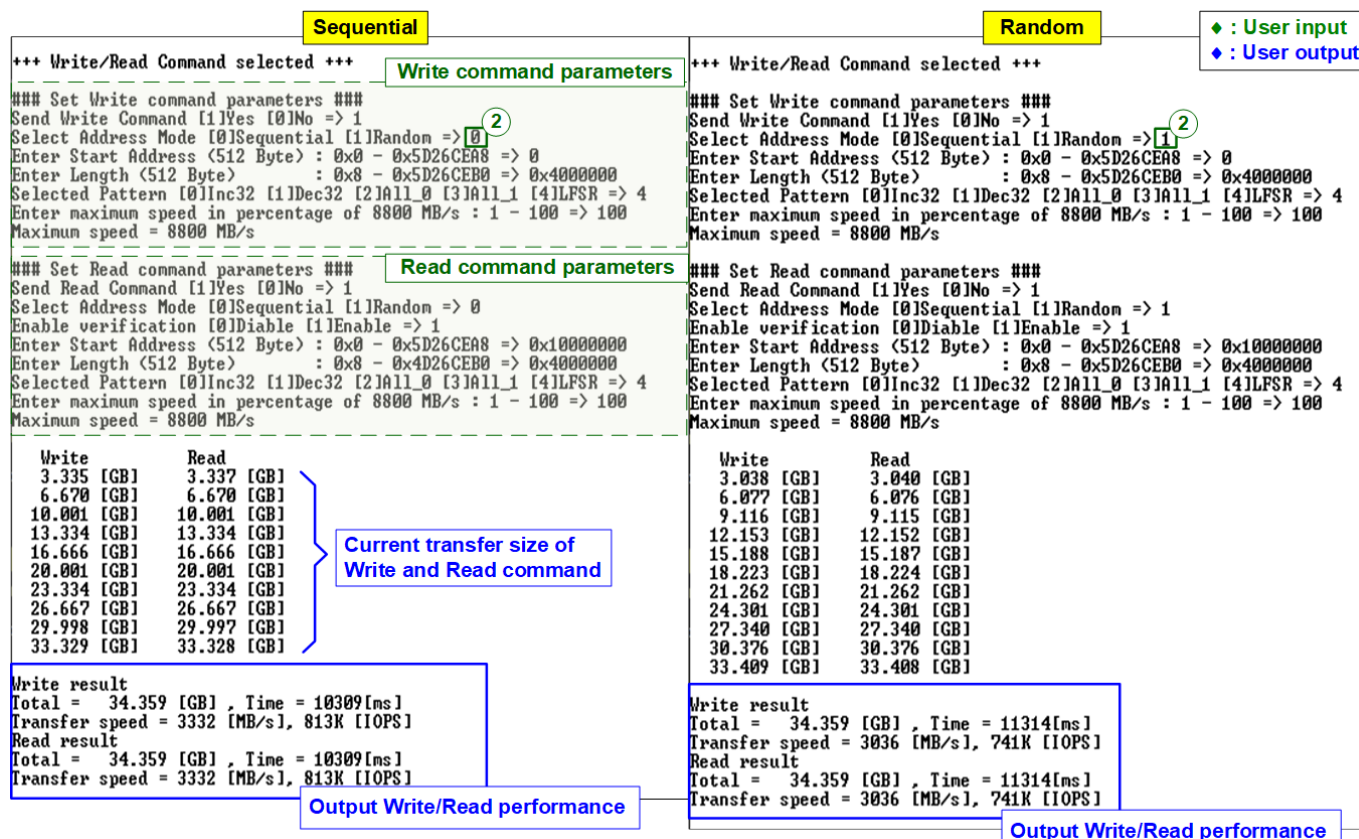


Figure 2-9 Test result of mixed Write-Read command by using two Address modes

The inputs required for the Write command are described in topic 2.2.1, while the inputs for the Read command are described in topic 2.2.2. Note that the Write and Read commands utilize the same SSD resources, which means that increasing the maximum Write speed will result in reduced Read performance, and vice versa.

Figure 2-9 illustrates an SSD that achieves a well-balanced loading for both Write and Read commands. By setting the maximum speed of both Write and Read commands to 100, the resulting Write performance and the Read performance are equal after the operations are completed. In this scenario, the Sequential access performance reaches 3332 MB/s, which is higher than the Random access performance of 3036 MB/s.

Next, the limitations of the maximum speed for both Write command and Read command are demonstrated.

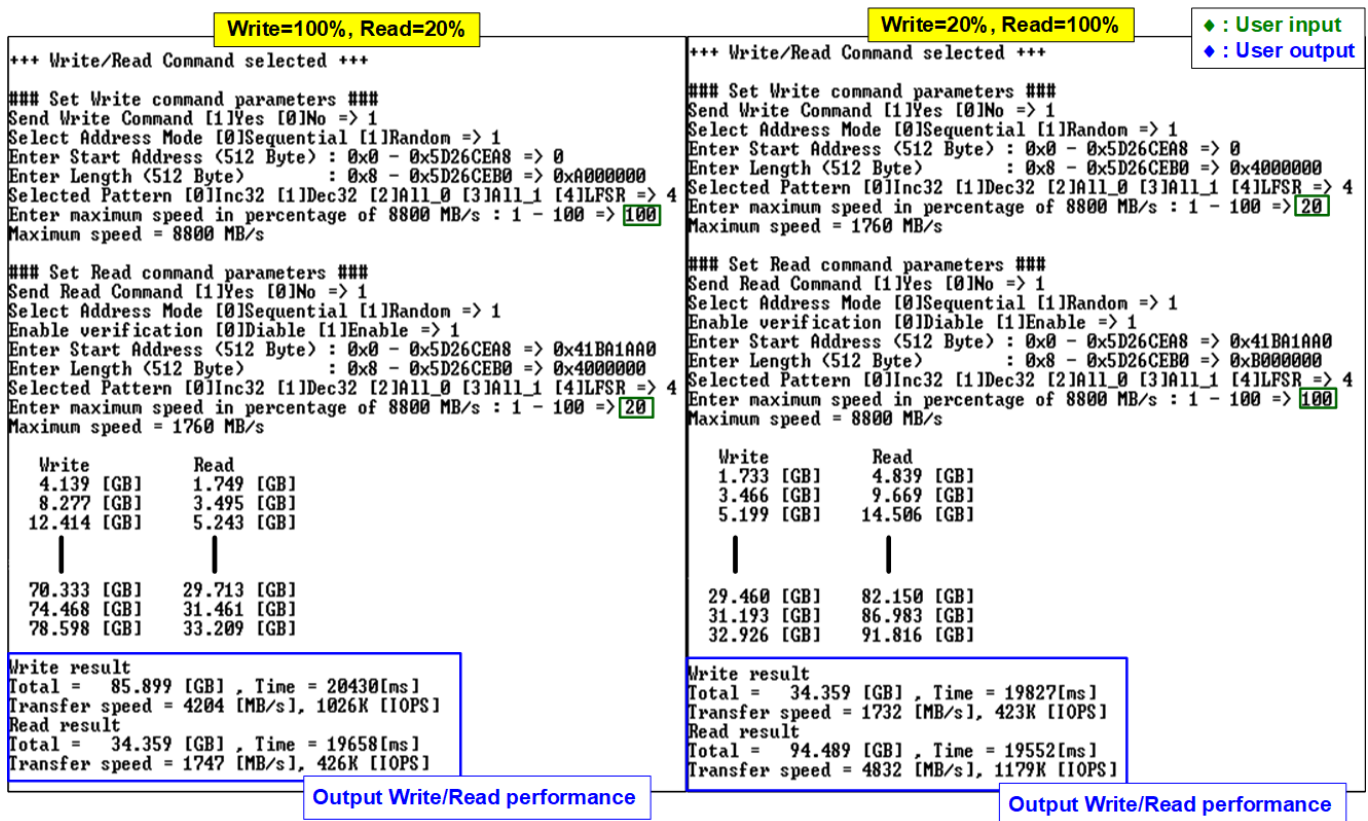


Figure 2-10 Test result when mixed Write-Read command with limited maximum speed

As shown in Figure 2-10, if the maximum speed is set to 20%, it will limit the speed to 20% of 8800 MB/s or 1760 MB/s. However, the actual performance achieved with the limited speed may be slightly lower than the set value. When using random access, the total performance, which includes both write and read operations, is approximately in the range of 6000 to 6500 MB/s.

2.3 SMART Command

To send a SMART command to the NVMe SSD, select option '2'. Once the operation is complete, the console will display the SMART/Health Information output, as shown in Figures 2-11. This information includes both the Health status and SMART log information.

```

+++ SMART Command selected +++
<< Health Status >>
Remaining Life : 100%

<< SMART Log Information >>
Percentage Used : 0%
Temperature : 32 Degree Celsius
Total Data Read : 18206 GB
Total Data Read <Raw data> : 0x00000000_00000000_00000000_021E8FAD
Total Data Written : 22410 GB
Total Data Written <Raw data> : 0x00000000_00000000_00000000_029BD9F1
Power On Cycles : 125 Times
Power On Hours : 12 Hours
Unsafe Shutdowns : 17 Times

SMART Command Complete

--- rmNUMeIP for Gen4 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Shutdown Command

```

Data output decoded from SMART command

Figures 2-11 Test result when running SMART command

The Health status displays the remaining life of the SSD as a percentage, which is calculated from the Percentage Used value in the SMART log information. The SMART log information displays the following seven parameters.

- 1) Percentage used: The percentage of the SSD's lifespan that has been consumed.
- 2) Temperature: The temperature of the SSD in degree Celsius.
- 3) Total Data Read: The total amount of data that has been read from the SSD, displayed in GB/TB units. Additionally, the raw data without decoding is displayed as a 32-digit hex number (128 bits). The unit size of raw data is 512,000 bytes.
- 4) Total Data Written: The total amount of data that has been written to the SSD, displayed in GB/TB units. Additionally, the raw data without decoding is displayed as a 32-digit hex number (128 bits). The unit size of raw data is 512,000 bytes.
- 5) Power On Cycles: The number of times the SSD has been powered on.
- 6) Power On Hours: The total amount of time in hours that the SSD has been powered on.
- 7) Unsafe Shutdowns: The number of times the SSD has experienced an unsafe shutdown.

2.4 Flush Command

To initiate a Flush command on the NVMe SSD, select option '3' from the menu. The Flush command ensures that all modified data in the cache memory is written to Flash memory in the SSD.

```

+++ Flush Command selected +++
Flush Command Complete
--- rmNUMeIP for Gen4 ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Shutdown Command
  
```

Message after finishing the operation

Figure 2-12 Test result when running Flush command

Once the Flush operation is completed, the console will display the message “Flush Command Complete”.

2.5 Shutdown Command

To send a Shutdown command to the NVMe SSD, select option '4'.

```

--- rmNUMeIP for Gen4 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : 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...
  
```

◆ : User input
◆ : User output

Confirmation message

Press 'y' to confirm

Last message before the IP and SSD are inactive

Figure 2-13 Test result when running Shutdown command

A confirmation message will be displayed on the console, and the user will need to enter 'y' or 'Y' to proceed with the operation. Press any other key to cancel the operation.

Once the Shutdown operation is complete, “Shutdown command is complete” will be displayed as the final message. The console becomes inactive. To begin a new test operation, the user will need to power off and on the test system.

3 Revision History

Revision	Date	Description
1.1	16-Jun-23	Update write performance
1.0	12-May-23	Initial version release