



# rmNVMe-IP for Gen5 Demo Instruction

Rev1.0 8-Aug-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 .....	9
2.2.3	Mixed Write/Read Command .....	11
2.3	SMART Command.....	13
2.4	Flush Command .....	14
2.5	Secure Erase Command .....	14
2.6	Shutdown Command .....	15
3	Revision History.....	16

# 1 Overview

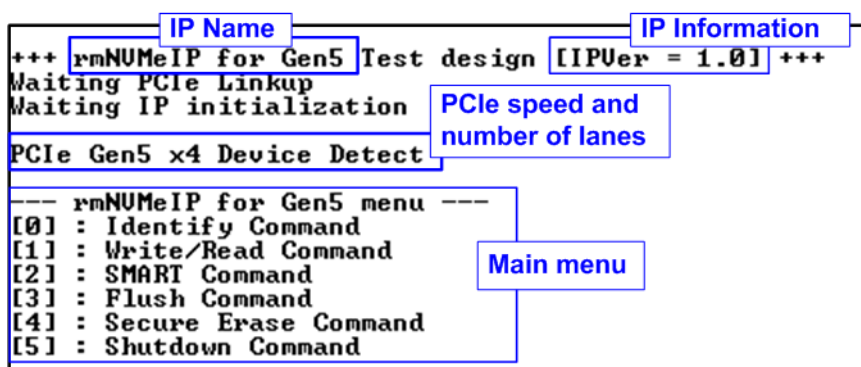
This guide provides instructions on how to use the rmNVMe-IP for Gen5 demo on an FPGA development board, which involves accessing an NVMe Gen5 SSD. The demo supports performing Identify, Write, Read, SMART, Flush, Secure Erase and Shutdown commands. It allows the user to control test operation through FPGA console.

To get started, follow the instructions in the “dg\_nvmeip\_fpgasetup\_g5” document to set up the FPGA board and the test system. Once setup is completed, the rmNVMeIP for Gen5 menu will be displayed and the user can select the desired test operation through the FPGA console.

```

+++ rmNUMeIP for Gen5 Test design [IPVer = 1.0] +++
Waiting PCIe Linkup
Waiting IP initialization
PCIe Gen5 x4 Device Detect

--- rmNUMeIP for Gen5 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Secure Erase Command
[5] : Shutdown Command
  
```



The diagram shows a terminal window with the following text and callouts:

- IP Name**: Points to "rmNUMeIP for Gen5 Test design"
- IP Information**: Points to "[IPVer = 1.0]"
- PCIe speed and number of lanes**: Points to "PCIe Gen5 x4 Device Detect"
- Main menu**: Points to the menu list starting with "--- rmNUMeIP for Gen5 menu ---"

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

The Identify command is used to retrieve information about the NVMe SSD. To send the Identify command, select '0' from the console menu. Once the command operation is completed, the console displays the following three pieces of information.

```

+++ Identify Command selected +++
Model Number      : CSSD-M2M2TPG5NFZ
SSD Capacity      = 2000[GB]
Secure Erase Command : Support

--- rmNUMeIP for Gen5 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Secure Erase Command
[5] : Shutdown Command
  
```

Model name, SSD Capacity and Secure Erase Command Support (Output from Identify command)

Figure 2-1 Test result when running Identify command

- 1) SSD model number : This value is decoded from Identify controller data of SSD.
- 2) SSD capacity : This value is signal output from rmNVMe-IP.
- 3) Secure Erase Command Support: This value is decoded from the Identify controller data to show whether the SSD supports the Secure Erase command.

```

+++ 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 rmNVMe-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 command of each SSD. Furthermore, the menu also provides the option to disable either the Write or Read operation, 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.

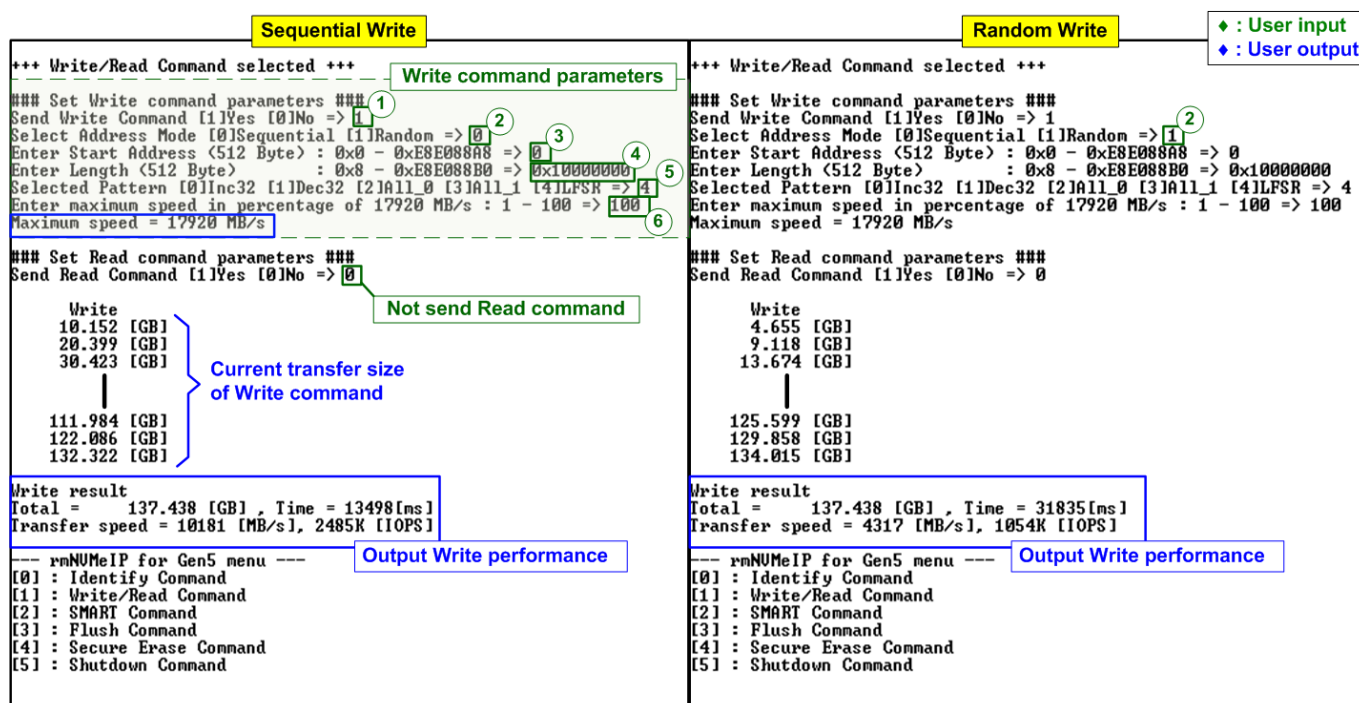


Figure 2-3 Test result of Write command using two Address modes

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 incrementally counted in order to store data in a contiguous memory area. 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 17,920 MB/s. Valid from 1 to 100.

*Note: 17,920 MB/s is calculated by user clock frequency (280 MHz) x data width (512-bit). This resulting value exceeds the maximum throughput of 4-lane PCIe Gen5 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.

In Figure 2-3, the results of the Write command are presented, showcasing the performance achieved through both sequential and random access. The sequential access demonstrates a significantly higher write performance, reaching 10181 MB/s, compared to the random access, which achieves 4317 MB/s. It is commonly observed that the performance of random access in most SSDs is much lower than that of sequential access. This menu allows the user to explore the characteristics of each SSD.

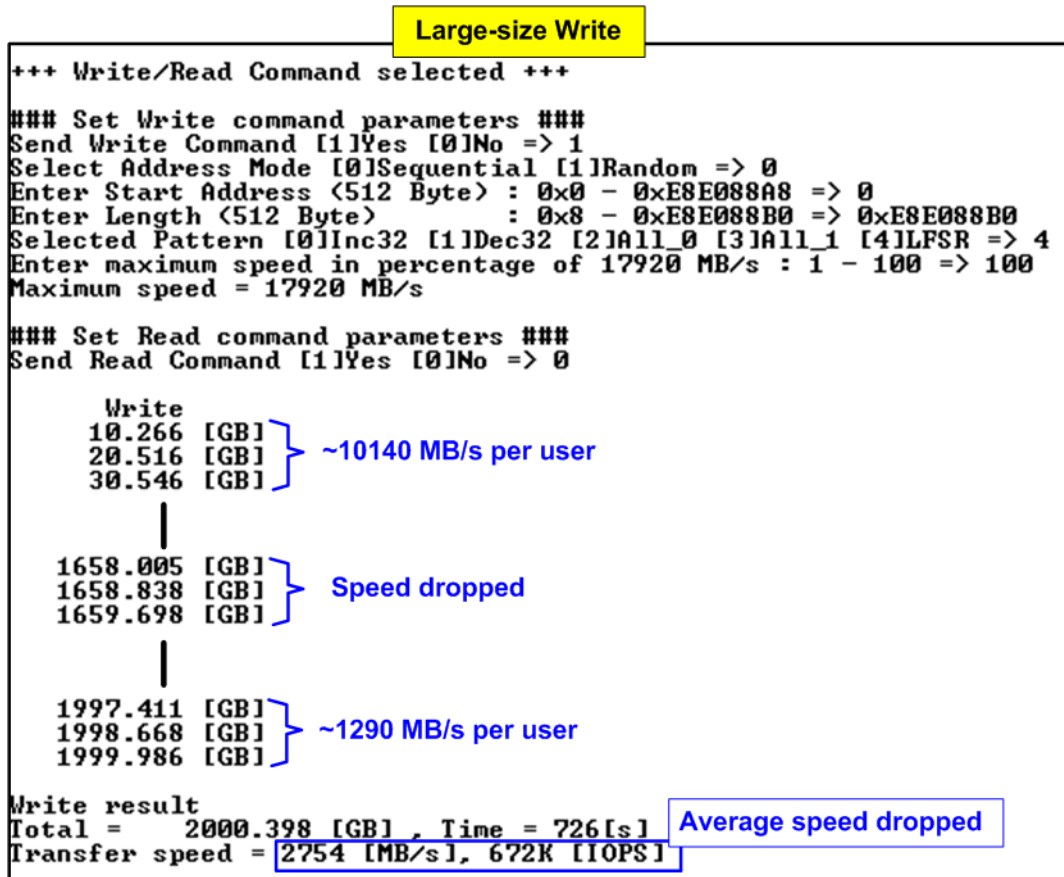


Figure 2-4 Write performance when using large transfer size

SSDs often incorporate a cache to enhance the performance of Write command. However, when transferring a large amount of data that fills up the cache, the performance tends to decrease. In Figure 2-4, the write performance result of the test SSD demonstrate high initial performance, reaching 10140 MB/s per user. However, as the SSD continues to receive write data, the performance decreases to 1290 MB/s. One key feature of enterprise SSDs is their ability to maintain stable performance during the transfer of data that fills the entire disk capacity. Therefore, it is recommended for users to verify the characteristics of the SSD to ensure they align with the system requirements.

*Note: The write performance of SSDs may decrease after long data transfer. In some cases, the performance can be restored by executing the Secure Erase command.*



Figure 2-5 Example Test data of the 1<sup>st</sup> and 2<sup>nd</sup> 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	
<pre>+++ Write/Read Command selected +++ ### Set Write command parameters ### Send Write Command [1]Yes [0]No =&gt; 6 Invalid input</pre>	<div style="border: 1px solid red; padding: 2px; display: inline-block;">Invalid input</div>
<pre>### Set Write command parameters ### Send Write Command [1]Yes [0]No =&gt; 1 Select Address Mode [0]Sequential [1]Random =&gt; 3 Invalid input</pre>	<div style="border: 1px solid red; padding: 2px; display: inline-block;">Invalid input</div>
<pre>### Set Read command parameters ### Send Read Command [1]Yes [0]No =&gt; 1 Select Address Mode [0]Sequential [1]Random =&gt; 0 Enable verification [0]Disable [1]Enable =&gt; 2 Invalid input</pre>	<div style="border: 1px solid red; padding: 2px; display: inline-block;">Invalid input</div>
<pre>### Set Write command parameters ### Send Write Command [1]Yes [0]No =&gt; 1 Select Address Mode [0]Sequential [1]Random =&gt; 0 Enter Start Address &lt;512 Byte&gt; : 0x0 - 0x5D26CEA8 =&gt; 0xFFFFFFFF Invalid input</pre>	<div style="border: 1px solid red; padding: 2px; display: inline-block;">Out of range input</div>
<pre>### Set Write command parameters ### Send Write Command [1]Yes [0]No =&gt; 1 Select Address Mode [0]Sequential [1]Random =&gt; 0 Enter Start Address &lt;512 Byte&gt; : 0x0 - 0x5D26CEA8 =&gt; 0 Enter Length &lt;512 Byte&gt; : 0x8 - 0x5D26CEB0 =&gt; 6 Invalid input</pre>	<div style="border: 1px solid red; padding: 2px; display: inline-block;">Input is not aligned to 8</div>
<pre>### Set Write command parameters ### Send Write Command [1]Yes [0]No =&gt; 1 Select Address Mode [0]Sequential [1]Random =&gt; 0 Enter Start Address &lt;512 Byte&gt; : 0x0 - 0x5D26CEA8 =&gt; 0 Enter Length &lt;512 Byte&gt; : 0x8 - 0x5D26CEB0 =&gt; 0x4000000 Selected Pattern [0]Inc32 [1]Dec32 [2]A11_0 [3]A11_1 [4]LFSR =&gt; 9 Invalid input</pre>	<div style="border: 1px solid red; padding: 2px; display: inline-block;">Invalid input</div>
<pre>### Set Write command parameters ### Send Write Command [1]Yes [0]No =&gt; 1 Select Address Mode [0]Sequential [1]Random =&gt; 0 Enter Start Address &lt;512 Byte&gt; : 0x0 - 0x5D26CEA8 =&gt; 0 Enter Length &lt;512 Byte&gt; : 0x8 - 0x5D26CEB0 =&gt; 0x4000000 Selected Pattern [0]Inc32 [1]Dec32 [2]A11_0 [3]A11_1 [4]LFSR =&gt; 4 Enter maximum speed in percentage of 11200 MB/s : 1 - 100 =&gt; 101 Invalid input</pre>	<div style="border: 1px solid red; padding: 2px; display: inline-block;">Out of range transfer rate</div>

Figure 2-6 Error message from the invalid input



## 2.2.2 Read Command

This topic shows an example when operating only Read command.

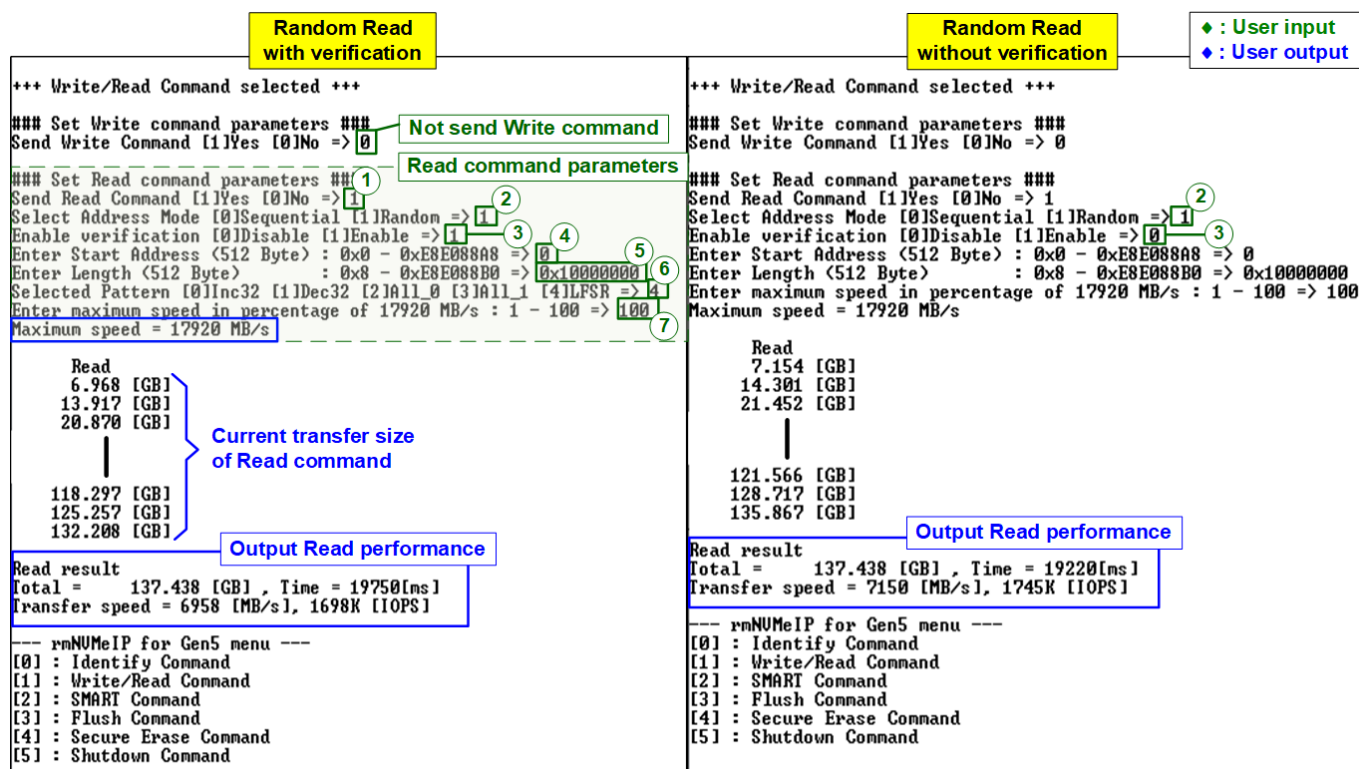


Figure 2-7 Test result of Read command using Random Address mode

The Read command operation receives 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 by 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 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 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 17,920 MB/s. Valid from 1 to 100.

*Note: 17,920 MB/s is calculated by user clock frequency (280 MHz) x data width (512-bit).*

Once all valid inputs are provided, the operation starts. During the operation, the console displays the current amount of received data every second, indicating that the system is functioning properly. In case of invalid input or failed data verification, an error message will be displayed instead. Upon completion of the operation, the console will present the total size, total time usage, test speed, and IOPS of the Read command.

Figure 2-7 showcases the performance results of the Read command operation using Random addressing. In this SSD, the optimal read performance of 7150 MB/s is achieved after executing a Secure Erase command to erase the data. Therefore, the verification feature must be disabled to receive the null data after erasing the SSD. The read performance is slightly reduced (6958 MB/s) after filling the SSD with the LFSR pattern, as shown on the left side of Figure 2-7.

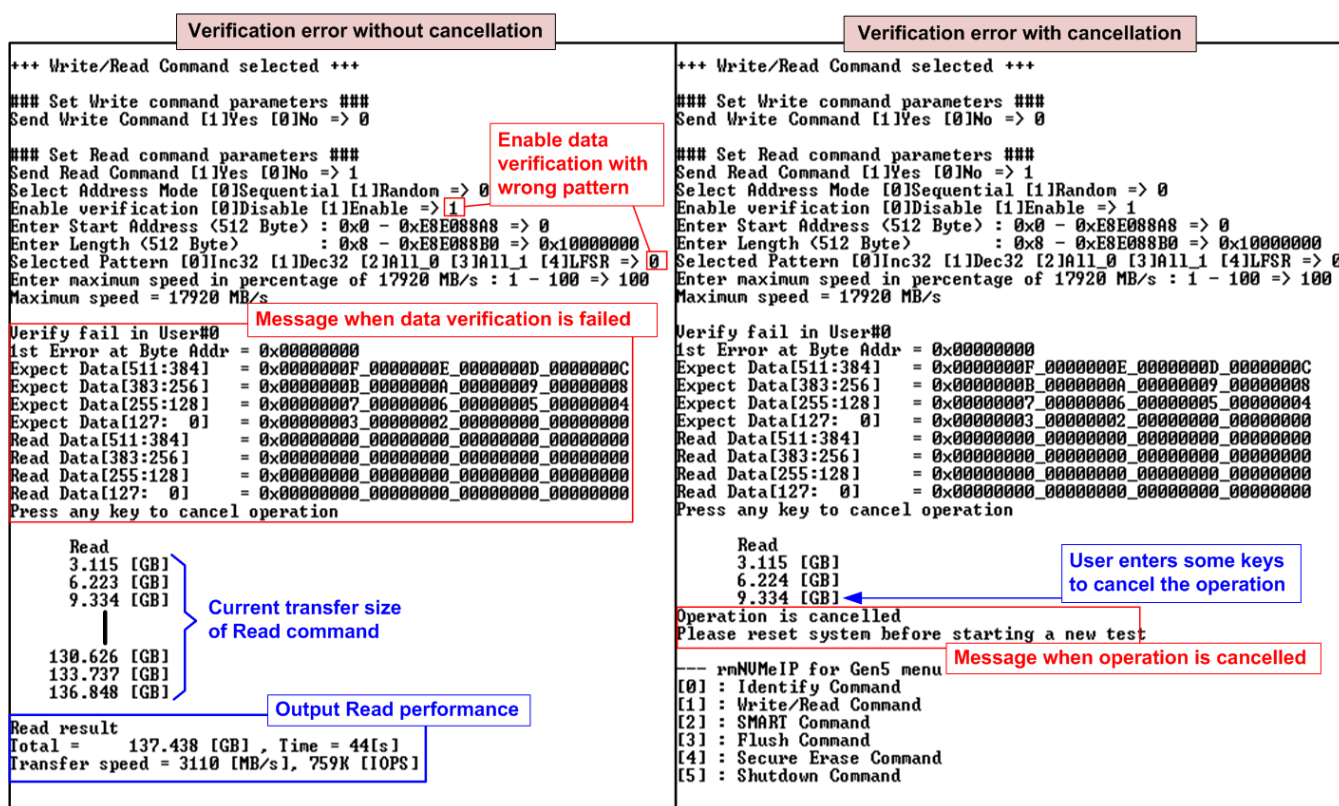


Figure 2-8 Data verification has failed

Figure 2-8 provides an example of error message displayed when data verification fails. The message “Verify fail” is displayed with detailed information about the 1<sup>st</sup> 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 input for cancellation, the operation will continue until the Read command is completed. 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 the expected 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.

*Note: The read performance of this SSD using Sequential addressing mode is half value compared to Random addressing mode. This phenomenon is dependent on the characteristics of the SSD controller.*

### 2.2.3 Mixed Write/Read Command

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

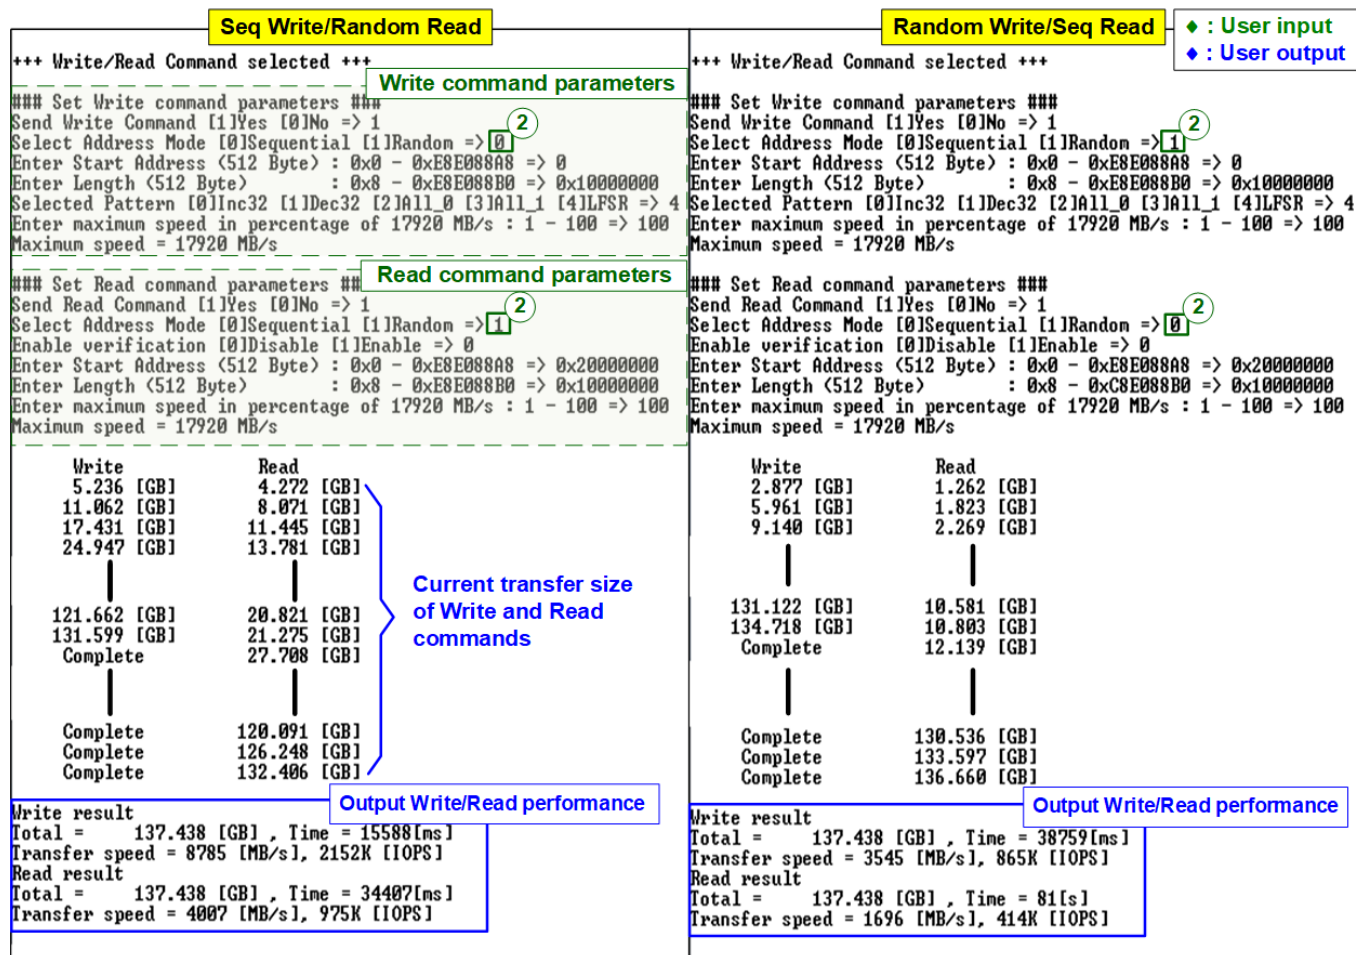


Figure 2-9 Test result of mixed Write-Read command without speed limitation

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

Figure 2-9 illustrates an SSD with imbalanced loading for Write and Read commands. When setting the maximum speed of both Write and Read commands to 100, the Write command exhibits better performance, compared to the Read command, regardless of whether sequential or random access is utilized. The optimal total performance of 12792 MB/s (8785 + 4007 MB/s) is achieved when using sequential Write mixed with random Read. Conversely, the slowest total performance of 5241 MB/s (3545 + 1696 MB/s) is observed when using random Write mixed with sequential Read. These specific characteristics are unique to this SSD. The results may differ when using other SSDs.



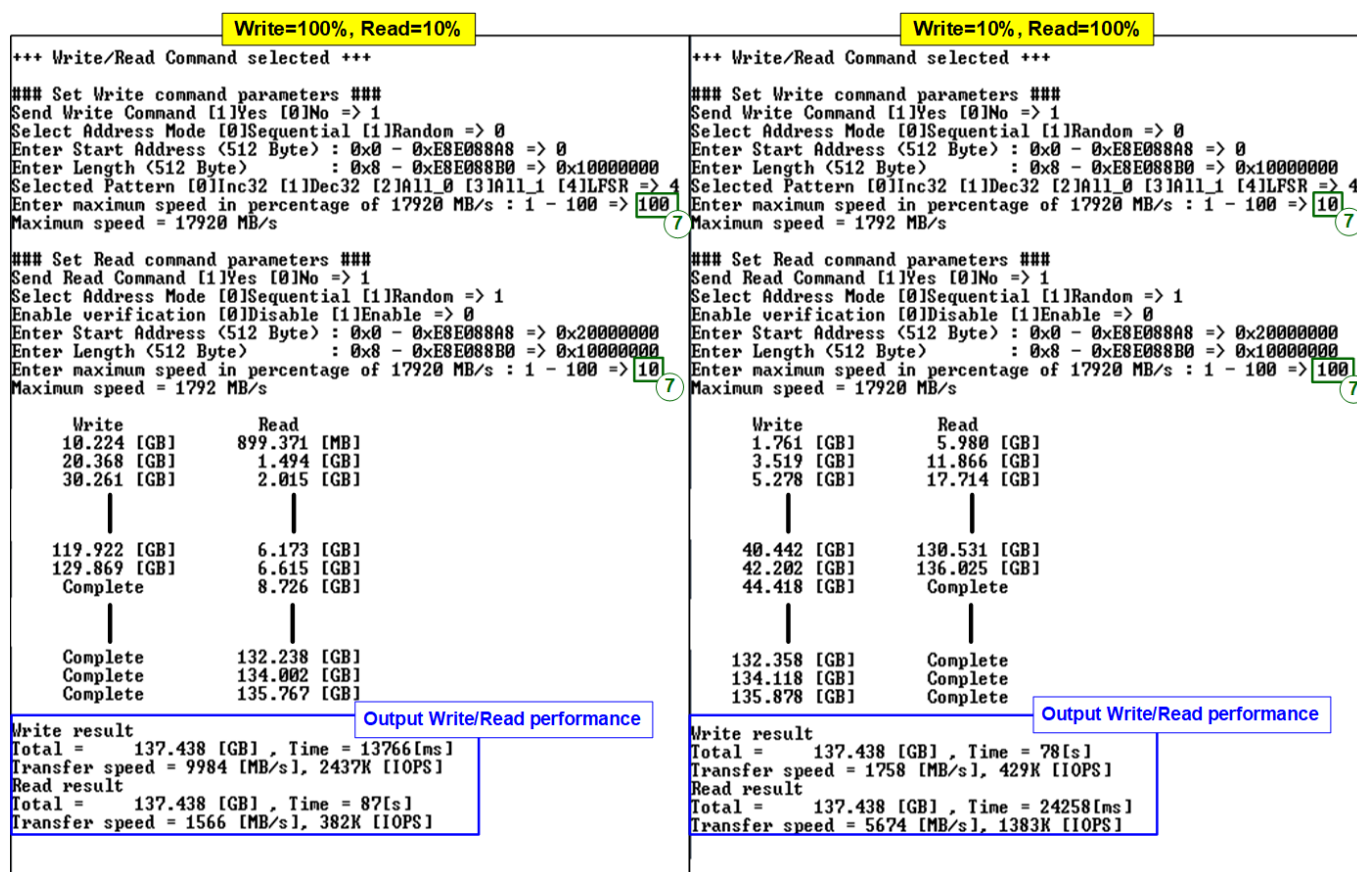


Figure 2-10 Test result when mixed Write-Read command with limited maximum speed

As shown in Figure 2-10, if the maximum speed is set to 10%, the speed will be limited to 10% of 17920 MB/s or 1795 MB/s. However, the actual performance achieved with the limited speed might be slightly lower than the set value. The test result shows that the write operation achieved a performance of 1758 MB/s, while the read operation achieved 1566 MB/s.

Furthermore, the test results indicate that enabling a 10% read operation leads to a slight reduction in write performance, from 10181 MB/s to 9984 MB/s. Similarly, when 10% write operation is enabled, the read performance decreases by approximately 18%, from 6958 MB/s to 5674 MB/s. These performance characteristics are specific to each SSD model and may vary accordingly.

## 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 Figure 2-11. This information includes both the Health status and SMART log information.

```

+++ SMART Command selected +++
<< Health Status >>
Remaining Life : 87%
<< SMART Log Information >>
Percentage Used : 13%
Temperature : 33 Degree Celsius
Total Data Read : 146979 GB
Total Data Read (Raw data) : 0x00000000_00000000_00000000_111C0C42
Total Data Written : 362068 GB
Total Data Written (Raw data) : 0x00000000_00000000_00000000_2A25CD30
Power On Cycles : 423 Times
Power On Hours : 103 Hours
Unsafe Shutdowns : 29 Times

SMART Command Complete

--- rmNUMeIP for Gen5 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Secure Erase Command
[5] : Shutdown Command

```

Data output decoded from SMART command

Figure 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 Gen5 ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Secure Erase Command
[5] : Shutdown Command
  
```

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 Secure Erase Command

To initiate a Secure Erase command on the NVMe SSD, select option '4' from the menu. Before the operation begins, a confirmation message is displayed on the console, requesting the user to confirm the command. The user needs to enter 'y' or 'Y' to proceed with the operation or any other key to cancel.

```

+++ Secure Erase Command selected +++
Are you sure to erase all Data ?
Press 'y' to confirm : y
Secure Erase Command Complete
--- rmNUMeIP for Gen5 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Secure Erase Command
[5] : Shutdown Command
  
```

Figure 2-13 Test result when Secure Erase command

Once the Secure Erase command is completed, the console displays the message “Secure Erase Command Complete”.

## 2.6 Shutdown Command

To send a Shutdown command to the NVMe SSD, select option '5'.

```

--- rmNUMeIP for Gen5 menu ---
[0] : Identify Command
[1] : Write/Read Command
[2] : SMART Command
[3] : Flush Command
[4] : Secure Erase 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...
  
```

◆ : User input  
◆ : User output

Confirmation message

Press 'y' to confirm

Last message before the IP and SSD are inactive

Figure 2-14 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
1.0	7-Aug-23	Initial version release