

# 2-Ch RAID0 (NVMe IP for Gen 4/NVMeG4-IP)

## Demo Instruction

Rev1.2 3-Aug-22

### 1 Overview

This document describes the instruction to run 2-ch RAID0 demo by using NVMe-IP for Gen4/NVMeG4-IP. The demo is run on FPGA development board for accessing two NVMe Gen4 SSDs as RAID0. There are six test menus for running six commands - Identify, Write, Read, SMART, Flush, and Shutdown command. User controls test operation via FPGA console.

After user finishes FPGA board setup following “dg\_nvmeip\_raid0x2\_fpgasetup” document.

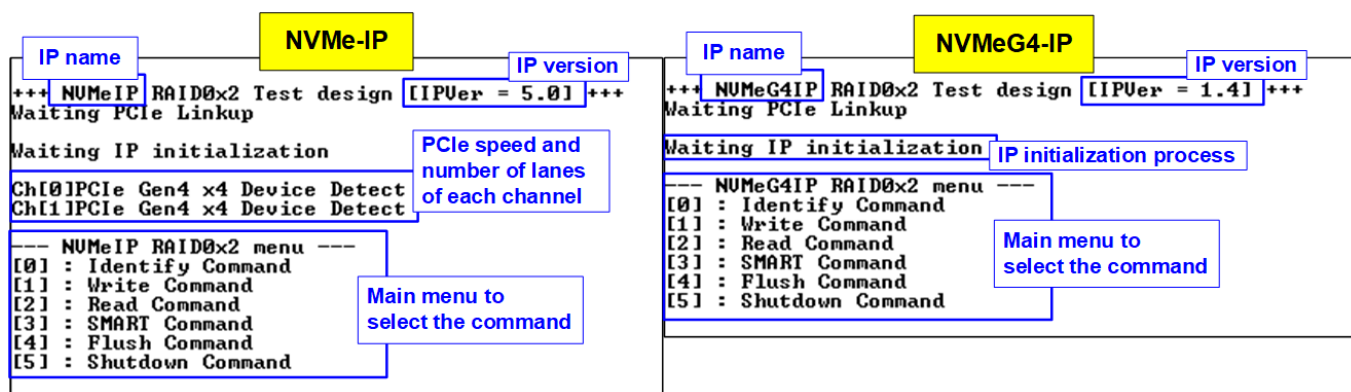


Figure 1-1 RAID0x2 menu after finish initialization

On welcome screen, IP name and IP version number are displayed. For standard IP (NVMe-IP), the PCIe speed and number of PCIe lanes of each channel are displayed in the next message. 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 RAID0.

```

+++ Identify Command selected +++
Ch[0]Model Number : WDS100T1X0E-00AFY0
Ch[1]Model Number : WDS100T1X0E-00AFY0
RAID Capacity= 2000[GB]
--- NUMeIP RAID0x2 menu
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command
  
```

Model name and RAID0 capacity

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 of each SSD.
- 2) RAID capacity : This value is calculated by multiplying device capacity in channel#0 by 2. Therefore, it is recommended to connect two SSDs which have the same size.

*Note: If two SSDs are different model which has different capacity, please connect the smallest capacity SSD to CH#0.*

When unsupported LBA size is detected, the error message is displayed on the console, as shown in Figure 2-2.

*Note: In RAID0 design, LBA size of SSD must be equal to 512-byte. Other size can be supported by modifying RAID0 controller hardware.*

```

Warning : LBA Size Not Support!
Please Check Devices
  
```

Unsupported LBA size is detected

Figure 2-2 Error message when LBA size does not support

## 2.2 Write Command

Select '1' to send Write command to RAID0.

```

+++ Write Command selected +++
Enter Start Address (512 Byte)      : 0 - 0xE8E0DB5F => 0
Enter Length (512 Byte)            : 1 - 0xE8E0DB60 => 0x8000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR => 4
9.614 [GB]
19.282 [GB]
28.926 [GB]
38.614 [GB]
48.338 [GB]
58.010 [GB]
67.677 [GB]
Total = 68.719 [GB] , Time = 7106[ms] , Transfer speed = 9670[MB/s]
--- NUMeIP RAID0x2 menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command

```

◆: User input  
◆: User output

1: Start Address input  
2: Length input  
3: Pattern input

Current transfer size

Output performance

Figure 2-3 Test result when running Write command

User sets three parameters as follows.

- 1) Start Address: Start address to write RAID0 in 512-byte unit. The input is decimal unit when the input is only digit number. User can add “0x” to be a prefix for hexadecimal unit.
- 2) Transfer Length: Total transfer size in 512-byte unit. The input is decimal unit when the input is only digit number. User can add “0x” to be a prefix for hexadecimal unit.
- 3) Test pattern: Select test data pattern for writing RAID0. 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 calculated and displayed on the console to be a test result.

Note:

1. Typically, the performance of 2-ch RAID0 system is about two times of the performance of one SSD. It is recommended to use the same SSD model for all SSD channels for matching SSD characteristic. If the different SSD model are applied in the system, the 2-ch RAID0 performance is limited by two times of the slowest SSD performance.
2. 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.

← 64-bit header of 512-byte data →

Offset	48-bit LBA Address = 0 0x0000																32-bit LFSR pattern																48-bit LBA Address = 1															
	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F																
0000000000	00	00	00	00	00	00	00	FF	FF	00	00	FF	FF	FF	00	01	00	00	00	00	00	00	00	00	00	FE	FF	00	00	FE	FF	FF	00															
0000000010	01	00	00	00	FF	01	00	00	FF	FF	01	00	FE	FF	FF	01	02	00	00	00	FC	01	00	00	FC	FF	01	00	FD	FF	FF	01																
0000000020	02	00	00	00	FF	03	00	00	FF	FF	03	00	FD	FF	FF	03	04	00	00	00	F9	03	00	00	F9	FF	03	00	FB	FF	FF	03																
0000000030	04	00	00	00	FF	07	00	00	FF	FF	07	00	FB	FF	FF	07	09	00	00	00	F2	07	00	00	F2	FF	07	00	F6	FF	FF	07																
0000000040	09	00	00	00	FF	0F	00	00	FF	FF	0F	00	F6	FF	FF	0F	12	00	00	00	E4	0F	00	00	E4	FF	0F	00	ED	FF	FF	0F																
0000000050	12	00	00	00	FF	1F	00	00	FF	FF	1F	00	ED	FF	FF	1F	24	00	00	00	C9	1F	00	00	C9	FF	1F	00	DB	FF	FF	1F																
0000000060	24	00	00	00	FF	3F	00	00	FF	FF	3F	00	DB	FF	FF	3F	49	00	00	00	92	3F	00	00	92	FF	3F	00	B6	FF	FF	3F																
0000000070	49	00	00	00	FF	7F	00	00	FE	FF	7F	00	B6	FF	FF	7F	92	00	00	00	24	7F	00	00	25	FF	7F	00	6D	FF	FF	7F																
0000000080	92	00	00	00	FF	FF	00	00	FD	FF	FF	00	6D	FF	FF	FF	24	01	00	00	49	FE	00	00	4B	FE	FF	00	DB	FE	FF	FF																
	<b>SSD#0</b>																<b>SSD#1</b>																															
00000001D0	F3	24	49	12	C9	B6	FF	FF	25	C9	B6	FF	28	92	A4	ED	E7	49	92	24	DD	DB	24	C9	31	A4	6D	C9	3C	FF	7F	DB																
00000001E0	E7	49	92	24	92	6D	FF	FF	4A	92	6D	FF	51	24	49	DB	CF	93	24	49	BA	B7	49	92	62	48	DB	92	79	FE	FF	B6																
00000001F0	CF	93	24	49	24	DB	FE	FF	94	24	DB	FE	A3	48	92	B6	9E	27	49	92	75	6F	93	24	C5	90	B6	25	F2	FC	FF	6D																
0000000200	02	00	00	00	00	00	00	00	FD	FF	00	00	FD	FF	FF	00	03	00	00	00	00	00	00	00	FC	FF	00	00	FC	FF	FF	00																
0000000210	04	00	00	00	FA	01	00	00	FA	FF	01	00	FB	FF	FF	01	07	00	00	00	F9	01	00	00	F9	FF	01	00	F8	FF	FF	01																
0000000220	09	00	00	00	F4	03	00	00	F4	FF	03	00	F6	FF	FF	03	0F	00	00	00	F2	03	00	00	F2	FF	03	00	F0	FF	FF	03																
0000000230	12	00	00	00	E9	07	00	00	E9	FF	07	00	ED	FF	FF	07	1F	00	00	00	E4	07	00	00	E4	FF	07	00	E0	FF	FF	07																
0000000240	24	00	00	00	D2	0F	00	00	D2	FF	0F	00	DB	FF	FF	0F	3F	00	00	00	C9	0F	00	00	C9	FF	0F	00	C0	FF	FF	0F																
0000000250	49	00	00	00	A4	1F	00	00	A4	FF	1F	00	B6	FF	FF	1F	7F	00	00	00	92	1F	00	00	92	FF	1F	00	80	FF	FF	1F																
0000000260	92	00	00	00	49	3F	00	00	49	FF	3F	00	6D	FF	FF	3F	FF	00	00	00	24	3F	00	00	24	FF	3F	00	00	FF	FF	3F																
0000000270	4	01	00	00	92	7E	00	00	93	FE	7F	00	DB	FE	FF	7F	FF	01	00	00	49	7E	00	00	48	FE	7F	00	00	FE	FF	7F																
0000000280	49	02	00	00	24	FD	00	00	26	FD	FF	00	B6	FD	FF	FF	FF	03	00	00	92	FC	00	00	90	FC	FF	00	00	FC	FF	FF																

64-bit header of the next 512-byte data

Figure 2-4 Example Test data of the 1<sup>st</sup> 512-byte of each SSD by using LFSR pattern

The stripe size in 2-ch RAID0 demo is 512-byte. For incremental, decremental, or LFSR pattern, each 512-byte data has unique 64-bit header which consists 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 1<sup>st</sup> stripe of RAID0 is mapped to the first 512-byte of SSD#0 while the 2<sup>nd</sup> stripe of RAID0 is mapped to the first 512-byte of SSD#1, as shown in Figure 2-4. The unique header is not included when running all-0 or all-1 pattern.

```

Error input
+++ Write Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E0DB5F => 0xFFFFFFFF
Invalid input

+++ Write Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E0DB5F => 0
Enter Length (512 Byte) : 1 - 0xE8E0DB60 => 0xFFFFFFFF
Invalid input

+++ Write Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E0DB5F => 0
Enter Length (512 Byte) : 1 - 0xE8E0DB60 => 0x8000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR => 7
Invalid input
    
```

Figure 2-5 Error message from the invalid input

Figure 2-5 shows the example when the input is out-of-range from the recommended value. The console displays “Invalid input” and then the operation is cancelled.

## 2.3 Read Command

Select '2' to send Read command to RAID0.

```

    ◆ : User input
    ◆ : User output

+++ Read Command selected +++
Enter Start Address (512 Byte) : 0 - 0xE8E0DB5F => 0
Enter Length (512 Byte) : 1 - 0xE8E0DB60 => 0x80000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR => 4
8.580 [GB]
17.164 [GB]
25.749 [GB]
34.323 [GB]
42.907 [GB]
51.490 [GB]
60.075 [GB]
68.664 [GB]
Total = 68.719 [GB] , Time = 8006[ms] , Transfer speed = 8582[MB/s]

--- NUMeIP RAID0x2 menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command
  
```

Annotations in the image:  
 1: Points to the input '0' for Start Address.  
 2: Points to the input '0x80000000' for Length.  
 3: Points to the input '4' for Selected Pattern.  
 A bracket groups the progress bars (8.580 to 68.664 GB) as 'Current transfer size'.  
 A box at the bottom right of the progress bars is labeled 'Output performance'.  
 A box at the bottom of the terminal output is labeled 'Output performance'.

Figure 2-6 Input and result of Read Command menu

User inputs three parameters as follows.

- 1) Start Address: Start address to read RAID0 in 512-byte unit. The input is decimal unit when the input is only digit number. User can add "0x" to be a prefix for hexadecimal unit.
- 2) Transfer Length: Total transfer size in 512-byte unit. The input is decimal unit when the input is only digit number. User can add "0x" to be a prefix for hexadecimal unit.
- 3) Test pattern: Select test data pattern to verify data from RAID0. 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 starts reading data from RAID0 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. Finally, total size, total time usage, and test speed are calculated and displayed on the console to be a test result.

Note:

1. Some SSDs shows the different performance when changing test pattern. For example, the read performance when using all-zero pattern is better than using LFSR pattern.
2. Similar to Write command, when 2 SSDs are different, performance of RAID0 is limited by 2 times of the slowest SSD performance.
3. For the system that is sensitive to the read performance, the read performance of most SSDs is better when using customized IP that extends the buffer size from 256 Kbytes to 1 Mbytes. Please contact our sales for more information.

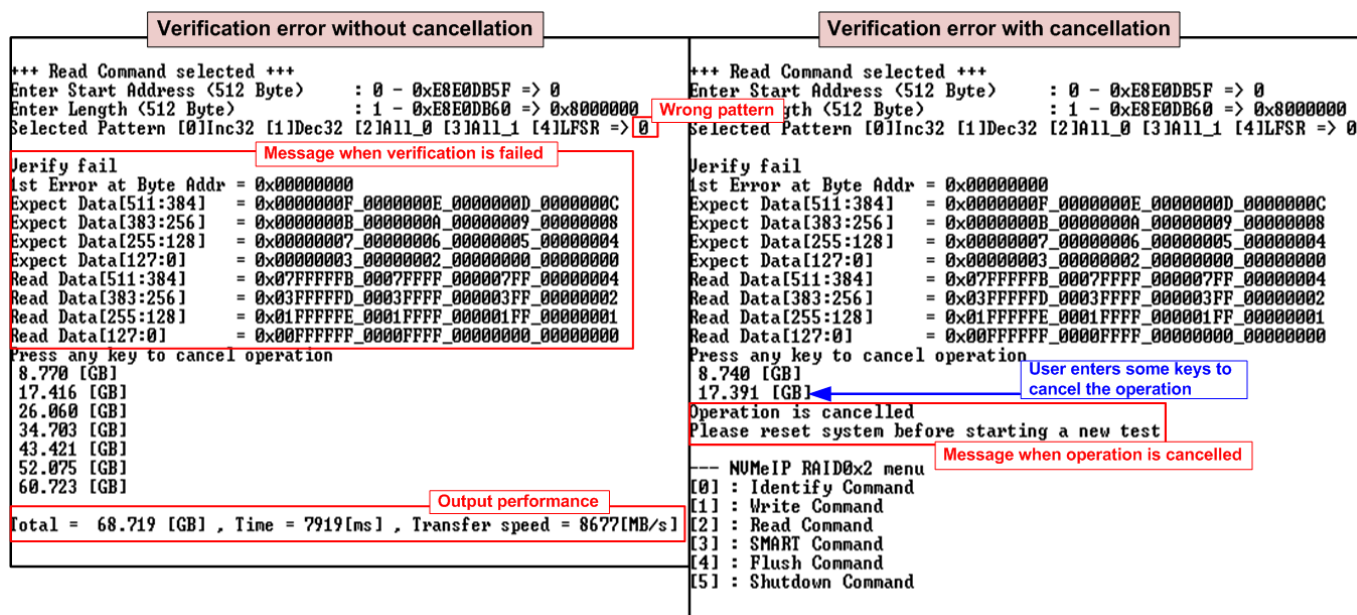


Figure 2-7 Data verification is failed

Figure 2-7 shows error message when data verification is failed. “Verify fail” is displayed with the information of the 1<sup>st</sup> 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 operation. After that, the output performance is displayed on the console.

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

## 2.4 SMART Command

Select '3' to send SMART command to RAID0.

```

+++ SMART Command selected +++
++ ch[0]
<< Health Status >>
Remaining Life : 98%
<< SMART Log Information Ch[0] >>
Percentage Used : 2%
Temperature : 38 Degree Celsius
Total Data Read : 22813 GB
Total Data Read <Raw data> : 0x00000000_00000000_00000000_02A7D6F6
Total Data Written : 43483 GB
Total Data Written <Raw data> : 0x00000000_00000000_00000000_050FD072
Power On Cycles : 427 Times
Power On Hours : 93 Hours
Unsafe Shutdowns : 367 Times
SSD#0

++ ch[1]
<< Health Status >>
Remaining Life : 94%
<< SMART Log Information Ch[1] >>
Percentage Used : 6%
Temperature : 35 Degree Celsius
Total Data Read : 124309 GB
Total Data Read <Raw data> : 0x00000000_00000000_00000000_0E7879EE
Total Data Written : 159989 GB
Total Data Written <Raw data> : 0x00000000_00000000_00000000_129FC0C1
Power On Cycles : 17820 Times
Power On Hours : 701 Hours
Unsafe Shutdowns : 17608 Times
SSD#1

SMART Command Complete

--- NUMeIP RAID0x2 menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command

```

Figure 2-8 Test result when running SMART command

When finishing the operation, SMART/Health Information (output from SMART command) of two SSDs are displayed as shown in Figure 2-8. 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, described 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 RAID0.

```

+++ Flush Command selected +++
Flush Command Complete
--- NUMeIP RAID0x2 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 2-9 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 RAID0.

```

--- NUMeIP RAID0x2 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 RAID0 and all SSDs are inactive

Figure 2-10 Test result when running Shutdown command

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

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



### 3 Revision History

Revision	Date	Description
1.2	3-Aug-22	Support NVMe-IP for Gen4
1.1	20-Jul-21	Remove FPGA setup topic and update test result
1.0	8-May-20	Initial version release