

2-Ch RAID0 (NVMe-IP) Demo Instruction

Rev2.1 30-Mar-21

1 Overview

This document describes the instruction to run NVMe-IP 2-ch RAID0 demo on FPGA development board for accessing two NVMe SSDs as RAID0. The demo is designed to run 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, main menu is displayed. The user can set the input to the console for selecting test operation.

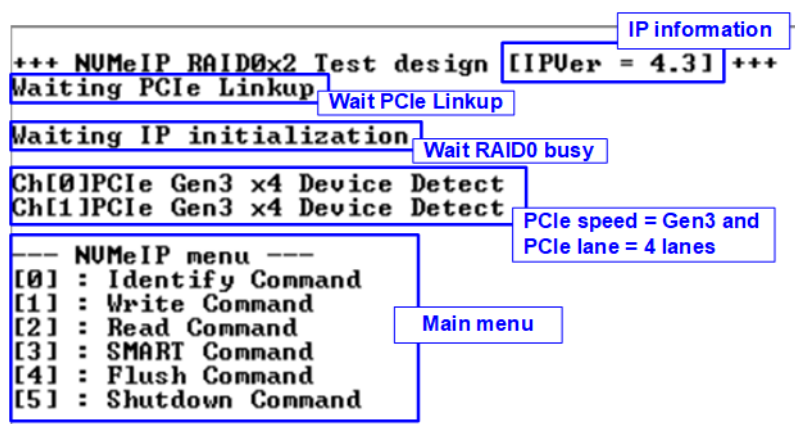


Figure 1-1 NVMe IP 2-ch RAID0 demo main menu

2 Test Menu

2.1 Identify Command

Select '0' to send Identify command to RAID0.

```

+++ Identify Command selected +++
Ch[0]Model Number : Samsung SSD 970 PRO 512GB
Ch[1]Model Number : Samsung SSD 970 PRO 512GB
RAID Capacity= 1024[GB]

--- NUMeIP 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.

When unsupported LBA size SSD 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.

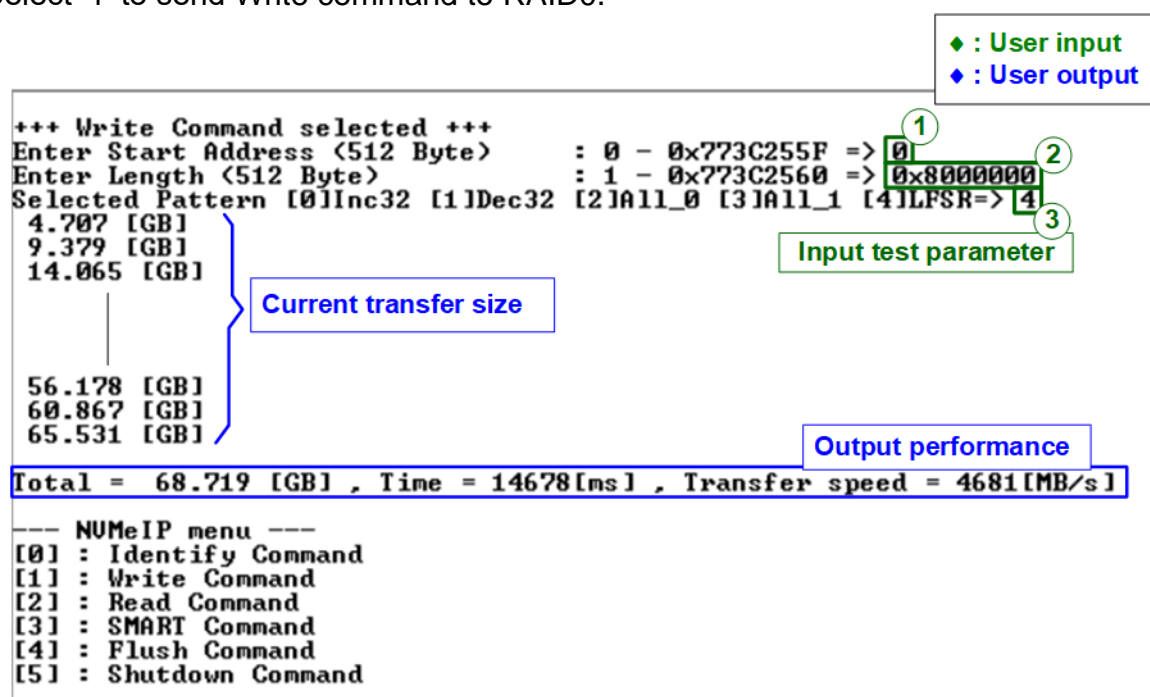
```

Not supported LBA size is detected
Warning : LBA Size Not Support!
Please Check Devices
  
```

Figure 2-2 Error message when LBA does not support

2.2 Write Command

Select '1' to send Write command to RAID0.



The screenshot shows the Write Command test interface. It includes a legend at the top right: a green diamond for 'User input' and a blue diamond for 'User output'. The interface displays the following text:

```

+++ Write Command selected +++
Enter Start Address (512 Byte) : 0 - 0x773C255F => 0
Enter Length (512 Byte) : 1 - 0x773C2560 => 0x80000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR=> 4
4.707 [GB]
9.379 [GB]
14.065 [GB]
56.178 [GB]
60.867 [GB]
65.531 [GB]
Total = 68.719 [GB] , Time = 14678[ms] , Transfer speed = 4681[MB/s]
--- NUmEIP menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command
  
```

Annotations on the screenshot include:

- A green box labeled '1' pointing to the Start Address input '0'.
- A green box labeled '2' pointing to the Length input '0x80000000'.
- A green box labeled '3' pointing to the Pattern input '4'.
- A green box labeled 'Input test parameter' pointing to the three input fields.
- A blue box labeled 'Current transfer size' pointing to the list of transfer sizes (4.707 to 65.531 GB).
- A blue box labeled 'Output performance' pointing to the total size, time, and transfer speed line.

Figure 2-3 Test result when running Write command

User inputs 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. During writing data, current transfer size 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.

←64-bit header of each sector→															
48-bit LBA Address = 0 0x0000								32-bit LFSR pattern							
Offset	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E
0000000000	00	00	00	00	00	00	00	00	01	00	00	00	02	00	00
0000000010	04	00	00	00	09	00	00	00	12	00	00	00	24	00	00
0000000020	49	00	00	00	92	00	00	00	24	01	00	00	49	02	00
0000000030	92	04	00	00	24	09	00	00	49	12	00	00	92	24	00
0000000040	24	49	00	00	49	92	00	00	92	24	01	00	24	49	02
0000000050	49	92	04	00	92	24	09	00	24	49	12	00	49	92	24
48-bit LBA Address = 2															
00000001A0	A0	E1	C6	74	41	C3	8D	E9	83	86	1B	D3	06	0D	37
00000001B0	0C	1A	6E	4C	18	34	DC	98	30	68	B8	31	60	D0	70
00000001C0	C0	A0	E1	C6	81	41	C3	8D	03	83	86	1B	07	06	0D
00000001D0	0F	0C	1A	6E	1F	18	34	DC	3F	30	68	B8	7F	60	D0
00000001E0	FF	C0	A0	E1	FF	81	41	C3	FE	03	83	86	FD	07	06
00000001F0	FA	0F	0C	1A	F4	1F	18	34	E9	3F	30	68	D3	7F	60
0000000200	02	00	00	00	00	00	00	00	04	00	00	00	09	00	00
0000000210	12	00	00	00	24	00	00	00	49	00	00	00	92	00	00
Disk#0															
48-bit LBA Address = 1															
Offset	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E
0000000000	01	00	00	00	00	00	00	00	00	00	00	00	00	00	00
0000000010	09	00	00	00	12	00	00	00	24	00	00	00	49	00	00
0000000020	92	00	00	00	24	01	00	00	49	02	00	00	92	04	00
0000000030	24	09	00	00	49	12	00	00	92	24	00	00	24	49	00
0000000040	49	92	00	00	92	24	01	00	24	49	02	00	49	92	04
0000000050	92	24	09	00	24	49	12	00	49	92	24	00	93	24	49
48-bit LBA Address = 3															
0000000000	41	C3	8D	E9	83	86	1B	D3	06	0D	37	A6	0C	1A	6E
0000000010	18	34	DC	98	30	68	B8	31	60	D0	70	63	C0	A0	E1
0000000020	81	41	C3	8D	03	83	86	1B	07	06	0D	37	0F	0C	1A
0000000030	1F	18	34	DC	3F	30	68	B8	7F	60	D0	70	FF	C0	A0
0000000040	FF	81	41	C3	FE	03	83	86	FD	07	06	0D	FA	0F	0C
0000000050	F4	1F	18	34	E9	3F	30	68	D3	7F	60	D0	A7	FF	C0
0000000000	03	00	00	00	00	00	00	00	00	00	00	00	0F	00	00
0000000010	1F	00	00	00	3F	00	00	00	7F	00	00	00	FF	00	00
Disk#1															

Figure 2-4 Example Test data of the 1st 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 1st stripe of RAID0 is mapped to the first 512-byte of SSD#0 while the 2nd 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.

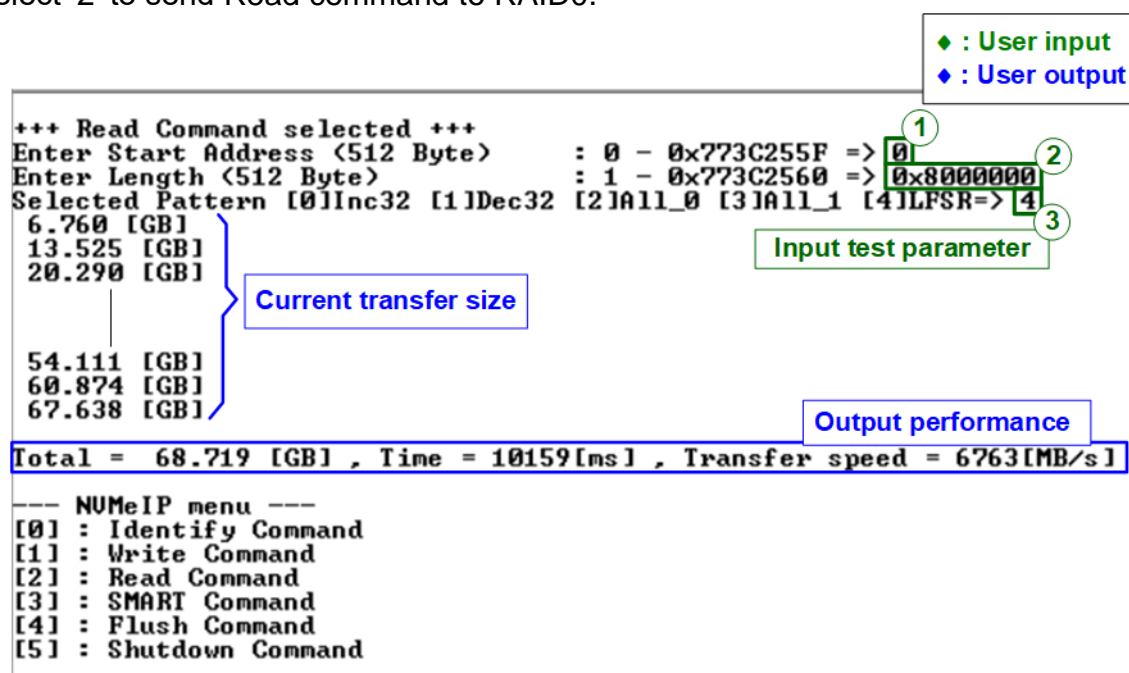
Error input	
+++ Write Command selected +++ Enter Start Address <512 Byte> Invalid input	: 0 - 0x773C255F => 0x00000000 Out of range address
+++ Write Command selected +++ Enter Start Address <512 Byte> Enter Length <512 Byte> Invalid input	: 0 - 0x773C255F => 0 Out of range length : 1 - 0x773C2560 => 0x00000000
+++ Write Command selected +++ Enter Start Address <512 Byte> Enter Length <512 Byte> Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR=> Invalid input	: 0 - 0x773C255F => 0 : 1 - 0x773C2560 => 0x80000000 Invalid pattern

Figure 2-5 Error message from the invalid input

Figure 2-5 shows the example when the input is not in the recommended range for each parameter. The console displays “Invalid input” and then the operation is cancelled.

2.3 Read Command

Select '2' to send Read command to RAID0.



```

+++ Read Command selected +++
Enter Start Address (512 Byte) : 0 - 0x773C255F => 0
Enter Length (512 Byte) : 1 - 0x773C2560 => 0x80000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_1 [4]LFSR=> 4
6.760 [GB]
13.525 [GB]
20.290 [GB]
54.111 [GB]
60.874 [GB]
67.638 [GB]
Total = 68.719 [GB] , Time = 10159[ms] , Transfer speed = 6763[MB/s]
--- NUMeIP menu ---
[0] : Identify Command
[1] : Write Command
[2] : Read Command
[3] : SMART Command
[4] : Flush Command
[5] : Shutdown Command

```

Figure 2-6 Input and result of Read Command menu

User inputs three parameters as follows.

- 1) Start Address: Start address to read SSD as 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 as 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. During reading data, current transfer size is displayed on the console every second to show that system is still alive. Total size, total time usage, and test speed are calculated and displayed after finishing data transferring.

Figure 2-7 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 or wait until finishing Read command. Without cancelling the operation, the read operation runs until finishing and then displaying the performance on the console as a test result.

When cancelling the operation, the read command still runs as the background process. So, the operation does not complete in the good sequence. It is recommended to power-off/on AB18/AB16 and then presses “RESET” button to restart system.

Verification error without cancellation

```

+++ Read Command selected +++
Enter Start Address <512 Byte>      : 0 - 0x773C255F => 0
Enter Length <512 Byte>            : 1 - 0x773C2560 => 0x80000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_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_000000000_000000000
Read Data[255:128]     = 0x0007FFFFF_FFF800004_0003FFFFF_FFFC00002
Read Data[127:0]       = 0x0001FFFFF_FFFE00001_000000000_000000000
Press any key to cancel operation

6.760 [GB]
13.526 [GB]
20.290 [GB]

54.115 [GB]
60.880 [GB]
67.646 [GB]

Total = 68.719 [GB] , Time = 10158[ms] , Transfer speed = 6764[MB/s]

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

```

Wrong pattern

Message when verification is failed

Output performance

Total = 68.719 [GB] , Time = 10158[ms] , Transfer speed = 6764[MB/s]

Verification error with cancellation

```

+++ Read Command selected +++
Enter Start Address <512 Byte>      : 0 - 0x773C255F => 0
Enter Length <512 Byte>            : 1 - 0x773C2560 => 0x80000000
Selected Pattern [0]Inc32 [1]Dec32 [2]All_0 [3]All_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_000000000_000000000
Read Data[255:128]     = 0x0007FFFFF_FFF800004_0003FFFFF_FFFC00002
Read Data[127:0]       = 0x0001FFFFF_FFFE00001_000000000_000000000
Press any key to cancel operation

6.762 [GB]
Operation is cancelled
Please reset system before starting a new test

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

```

User enters some keys to cancel the operation

Message when operation is cancelled

Figure 2-7 Data verification is failed

2.4 SMART Command

Select '3' to send SMART command to RAID0.

```

+++ SMART Command selected +++

++ ch[0]
<< Health Status >>
Remaining Life : 97%

<< SMART Log Information >>
Percentage Used : 3%
Temperature : 28 Degree Celsius
Total Data Read : 41886 GB
Total Data Read <Raw data> : 0x00000000_00000000_00000000_04E03866
Total Data Written : 52321 GB
Total Data Written <Raw data> : 0x00000000_00000000_00000000_0617306C
Power On Cycles : 1431 Times
Power On Hours : 104 Hours
Unsafe Shutdowns : 742 Times

++ ch[1]
<< Health Status >>
Remaining Life : 98%

<< SMART Log Information >>
Percentage Used : 2%
Temperature : 28 Degree Celsius
Total Data Read : 45640 GB
Total Data Read <Raw data> : 0x00000000_00000000_00000000_05501CDD
Total Data Written : 47900 GB
Total Data Written <Raw data> : 0x00000000_00000000_00000000_0593732A
Power On Cycles : 1145 Times
Power On Hours : 144 Hours
Unsafe Shutdowns : 580 Times

SMART Command Complete

--- NUMeIP menu ---

```

Data output decoded from SMART command (SSD Channel#0)

Data output decoded from SMART command (SSD Channel#1)

Figure 2-8 Test result when running SMART command

After 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. The 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 follow.

- 1) Percentage Used: Display SSD usage in percent unit.
- 2) Temperature in °C unit.
- 3) Total Data Read decoded as GB/TB unit. Additionally, raw data without decoding is displayed in 128-bit hexadecimal unit. The unit size of raw data is 512,000 bytes.
- 4) Total Data Written decoded as GB/TB unit. Additionally, raw data without decoding is displayed in 128-bit hexadecimal unit. 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 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 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 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 NVMe-IP and SSD are inactive status

Figure 2-10 Test result when running in Shutdown command

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 the test system to start new test operation.

3 Revision History

Revision	Date	Description
1.0	9-Oct-17	Initial version release
2.0	26-Jun-20	Remove FPGA setup from the document
2.1	30-Mar-21	Update SMART log information