

SANBlaze

FEATURE VALIDATION *Using NVM Compare*

SUMMARY

This white paper describes how to do a feature validation with SANBlaze software using NVM Compare as an example.

WRITTEN BY

Haiyan Lin
Senior Software Engineer



INTRODUCTION

When testing NVMe SSDs, some customers like to perform complex testing, tweaking multiple variables to create a variety of real-life scenarios that could potentially be replicated out in the field. For these customers, it is important for them to be able to control individual options in order to complete complex testing to validate their SSDs. Conversely, other customers want to run basic NVMe testing to ensure they comply with various specifications. For these customers, the testing process is greatly simplified using SANBlaze's pre-made scripts.

SANBlaze provides a complete system for NVMe testing, including three hardware options: The SANBlaze Express-DT (Desktop), the SANBlaze Express-RM (Rackmount), and the SANBlaze Express-RM4 (Gen4).



Figure 1: SANBlaze SBExpress-DT



Figure 2: SANBlaze SBExpress-RM / RM4

Whether you want to perform complex testing or testing that requires minimal tweaking of parameters and options, it is beneficial to learn and know what is "under the hood" when using SANBlaze. Knowing what goes on behind the scenes of simplified scripting can help you realize the potential of your testing and possibly save you many Qual hours in the process.

This white paper takes you through a SANBlaze validation process using NVM Compare as the example throughout. Understanding this process will enable you to use any SANBlaze NVMe command in the same manner - determining support in your drive for the command, running the command, checking the output, deciding if there needs to be a re-run, and viewing the output in the log files. Understanding what happens "under the hood" is useful to see the power of SANBlaze and how we can help you test your NVMe drives in the fastest and most accurate way possible.

Verify NVM Compare Feature with SANBlaze

The NVMe Compare command is one of the NVM commands. It reads the logical blocks specified by the command from the medium and compares the data read to a comparison data buffer transferred as part of the command.

If the data read from the controller and the comparison data buffer are equivalent with no mismatches, then the command completes successfully. If there is any mismatch, the command completes with an error of Compare Failure.

If metadata is provided, then a comparison is also performed for the metadata, excluding protection information. The command may specify protection information to be checked as well.

Verify NVM Compare Support

Not all SSD controllers support NVMe Compare. Users can find out through the SSD Optional NVM Command Support (ONCS) in the Identify Controller output which indicates the optional NVM commands and features supported by the controller. The bytes 521:520 are for ONCS and bit 0 of ONCS, if set to '1', this indicates the controller supports the Compare command. If cleared to '0', then the controller does not support the Compare command.

Bytes	O/M ¹	Description
521:520	M	<p>Optional NVM Command Support (ONCS): This field indicates the optional NVM commands and features supported by the controller. Refer to section 6.</p> <p>Bits 15:8 are reserved.</p> <p>Bit 7 if set to '1', then the controller supports the Verify command. If cleared to '0', then the controller does not support the Verify command.</p> <p>Bit 6 if set to '1', then the controller supports the Timestamp feature. If cleared to '0', then the controller does not support the Timestamp feature. Refer to section 5.21.1.14.</p> <p>Bit 5 if set to '1', then the controller supports reservations. If cleared to '0', then the controller does not support reservations. If the controller supports reservations then the following commands associated with reservations shall be supported: Reservation Report, Reservation Register, Reservation Acquire, and Reservation Release. Refer to section 8.8 for additional requirements.</p> <p>Bit 4 if set to '1', then the controller supports the Save field set to a non-zero value in the Set Features command and the Select field set to a non-zero value in the Get Features command. If cleared to '0', then the controller does not support the Save field set to a non-zero value in the Set Features command and the Select field set to a non-zero value in the Get Features command.</p> <p>Bit 3 if set to '1', then the controller supports the Write Zeroes command. If cleared to '0', then the controller does not support the Write Zeroes command.</p> <p>Bit 2 if set to '1', then the controller supports the Dataset Management command. If cleared to '0', then the controller does not support the Dataset Management command.</p> <p>Bit 1 if set to '1', then the controller supports the Write Uncorrectable command. If cleared to '0', then the controller does not support the Write Uncorrectable command.</p> <p>Bit 0 if set to '1', then the controller supports the Compare command. If cleared to '0', then the controller does not support the Compare command.</p>

Figure 3: Identify Controller Data Structure for ONCS

Write Range of LBAs With Specified Pattern

You can write the range of LBAs with specified pattern as follows:

The screenshot shows the SANBlaze VirtualUI interface. On the left is a navigation tree with 'Namespace' selected. The main window is titled 'Namespace 1 Tests' and contains a 'Test Configuration' section. In the 'New Disk Test' group, the following settings are visible and highlighted in red:

- Test:** Write
- # Blocks per I/O:** 1
- # I/Os per thread:** 100
- Seek Type:** Sequential
- Test Pattern:** 0x55AA55AA
- Write Enabled:**
- Start Test:** (button highlighted in red)

Figure 6: Write Range of LBAs with Specific Data Pattern

Click the NameSpace you want to run the command on in the left menu, then select tab “Namespace 1 Tests” on the right. In the “New Disk Test” group under “Test” select “Write”, enter “# Blocks per I/O” as 1, enter “# I/Os per thread” as 100, select “Seek Type” as “Sequential”, choose the “Test Pattern” you want (I chose 0x55AA55AA) above, check the “Write Enabled” box, then click “Start Test”. It will write LBA 0 to 99 with the specified pattern. All buttons/inputs are highlighted in red above.

You will see the writing passed as follows in the top of the same window.

The screenshot shows the 'Test Status' window with the following output:

```

Test started at Sep 17 18:09:22
test index is 662
proc level is 0, index is 660, Write_660 with 1 user (0 active)
device is port 0 target 101 lun 1, used by initiator 0 on this path
Device NVMe SAMSUNG MZVLV128HCGR-00000 BXV7000Q, SN S2J4HX0H903220
Write_660, 1 thread, 1 block per I/O, 100 I/Os, sequential I/O
I/Os done: 100
max outstanding I/Os: 1
status is Passed, 1 thread done, 0 errors
Test ended at Sep 17 18:09:22 (elapsed time 00.00.00)
    
```

Test	# Threads	# Blocks per I/O	Status	Errors	Action
Write_660	1	1	Passed	0	Start_Write_660, Delete_Write_660

Figure 7: Command Output of Writing Range of LBAs

The SCT=2h means Status Code Type = 2 and it is the Media and Data Integrity Error category. Any media specific errors that occur in the NVM or data integrity type errors shall be of this type. SC=85h means Status Code = 0x85 and it is Compare Failure. The command failed due to a miscompare during a Compare command as below:

Value	Description
84h	End-to-end Reference Tag Check Error: The command was aborted due to an end-to-end reference tag check failure.
85h	Compare Failure: The command failed due to a miscompare during a Compare command.
86h	Access Denied: Access to the namespace and/or LBA range is denied due to lack of access rights. Refer to the appropriate security specification (e.g., TCG Storage Interface Interactions specification).
87h	Deallocated or Unwritten Logical Block: The command failed due to an attempt to read from or verify an LBA range containing a deallocated or unwritten logical block.
88h to BFh	Reserved

Figure 15: Status Code – Media and Data Integrity Error Values, NVM Command Set

Run NVM Compare on SSD Not Supporting This Feature

If you tried to run the NVM Compare command on SSDs that don't support this feature, it will fail with "Invalid Opcode". For example, I have one drive not supporting NVM Compare and its ONCS = 0x4 as follows from Identify Controller output:

```

00000180: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00000190: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
000001a0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
000001b0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
000001c0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
000001d0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
000001e0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
000001f0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00000200: 66 44 00 00 01 00 00 00 04 00 00 00 06 00 ff ff
00000210: 00 00 01 00 00 00 00 00 00 00 00 00 00 00 00 00

```

```

. . .
MAXCMD =0x0000
NN = 0x00000001
ONCS = 0x0004
FUSES = 0x0000
FNA = 0x06
. . .

```

Figure 16: Identify Controller Output from SSD Not Supporting NVM Compare

Over 450 Tests!!

Let us certify your NVMe drives *before* they go to market so you can rest easy that your customers are taken care of.

What Do We Test?

- NVMe Generic I/O Commands
- NVMe I/O Tests
- NVMe Resets (all supported reset methods)
- NVMe Namespace Management
- NVMe Basic Management Commands
- NVMe-MI Full Command Set
- NVMe Dual Port Drive Tests
- NVMe SBExpress Hotplug and Link Testing
- NVMe Quarch Testing Pull/Plug Glitch
- NVMe Miscellaneous Commands (e.g. SR-IOV)
- IOL NVMe Certification
- IOL NVMe-MI Certification
- SSD Endurance JEDEC Spec. (long runtime)

Certified by SANBlaze™

Get Started
Today!

Tel.: (978) 679-1400
sales@sanblaze.com
support@sanblaze.com

For more information please contact:

SANBlaze, Inc.
One Monarch Drive, Suite 204
Littleton, MA 01460
www.sanblaze.com

sales@sanblaze.com | support@sanblaze.com