




OSS PCIe 4.0 Quad M.2 Carrier

OSS-PCIe4-ADPT-x16-M.2-4



User Manual

SKU: OSS-PCIe4-ADPT-x8-M.2-4



OSS
ONE STOP SYSTEMS

Table of Contents

Preface	4
Advisories	4
Safety Instructions	5
Protecting Against Electrostatic Discharge	6
1 Product Information	7
1.1 Removal Tray	7
1.2 PCIe M.2 Carrier Card.....	8
1.2.1 OSS-537 Board.....	8
1.2.2 OSS-568 Board.....	8
1.3 Features	9
1.4 General Specification	9
1.5 Block Diagram	10
1.6 Dimensions.....	10
1.7 M Key Connector.....	11
2 Hardware Requirements	12
3 Software Requirements	12
4 Hardware Installation	13
4.1 Installation-Procedures Overview	13
4.4 Prepare PCIe Carrier Card	14
4.5 Remove the canister	14
4.6 Install M.2 module	15
4.7 Secure the Media	15
4.8 Configure SW1 Dipswitch.....	16
4.8.1 SW1 Dipswitch	16
4.9 Connect Data Cables	16
4.10 Plugin the PCIe carrier board	17
4.11 Connect Power to Mezzanine Card.....	18
4.13 Install Canister	19
4.14 Power ON the computer.....	19
5 Hardware Check	20

5.1	LED Definition	20
6	Verify OSS Device on Windows OS	21
6.1	Device Manager	21
7	Verify M.2 Media on Windows OS.....	22
7.1	Device Manager	22
7.2	Disk Management	23
7.3	Storage Spaces	23
7.4	Check Temperature.....	24
8	Verify OSS Device on Linux OS.....	25
9	Verify M.2 Media on Linux OS	26
9.1	Check Temperature.....	28
9.2	Check OSS Device Speed and Linkwidth	29
9.3	Check M.2 Speed and Linkwidth	30
9.4	PCIe Device Detailed Info	31
9.5	Smartctl.....	32
10	Benchmark Performance (Read & Write).....	33
10.1	Disk Management Tool.....	33
11	FAQ's	35
12	Contacting Technical Support	36
11	Returning Merchandise to One Stop Systems	36
12	Shipping / transporting the card.....	37
11	APPENDIX A Compliance.....	37

Preface

Advisories

Five types of advisories are used throughout this manual to provide helpful information, or to alert you to the potential for hardware damage or personal injury.

**NOTE**

Used to amplify or explain a comment related to procedural steps or text.

**IMPORTANT**

Used to indicate an important piece of information or special “tip” to help you

**CAUTION**

Used to indicate and prevent the following procedure or step from causing damage to the equipment.

**WARNING**

Used to indicate and prevent the following step from causing injury.

**DANGER or STOP**

Used to indicate and prevent the following step from causing serious injury or significant data loss

Disclaimer: We have attempted to identify most situations that may pose a danger, warning, or caution condition in this manual. However, One Stop Systems does not claim to have covered all situations that might require the use of a Caution, Warning, or Danger indicator.

Safety Instructions

Always use caution when servicing any electrical component. Before handling the One Stop Systems Expansion chassis, read the following instructions and safety guidelines to prevent damage to the product and to ensure your own personal safety. Refer to the “Advisories” section for advisory conventions used in this manual, including the distinction between Danger, Warning, Caution, Important, and Note. Always use caution when handling/operating the computer. Only qualified, experienced, authorized electronics personnel should access the interior of the computer and expansion chassis.



WARNING

Never modify or remove the radio frequency interference shielding from your workstation or expansion unit. To do so may cause your installation to produce emissions that could interfere with other electronic equipment around your system.

When Working Inside a Computer

Before taking covers off a computer, perform the following steps:

Turn off the computer and any peripheral devices.

Disconnect the computer and peripheral power cords from their AC outlets or inlets to prevent electric shock or system board damage.

In addition, take note of these safety guidelines when appropriate:

To help avoid possible damage to systems boards, wait five seconds after turning off the computer before removing a component, removing a system board, or disconnecting a peripheral device from the computer.

When you disconnect a cable, pull on its connector or on its strain-relief loop, not on the cable itself. Some cables have a connector with locking tabs. If you are disconnecting this type of cable, press in on the locking tabs before disconnecting the cable. As you pull connectors apart, keep them evenly aligned to avoid bending any connector pins. Also, before connecting a cable, make sure both connectors are correctly oriented and aligned.



CAUTION

Do not attempt to service the system yourself except as explained in this manual. Follow installation instructions closely.

Protecting Against Electrostatic Discharge



Electrostatic Discharge (ESD) Warning

Electrostatic Discharge (ESD) is the enemy of semiconductor devices. You should always take precautions to eliminate any electrostatic charge from your body and clothing before touching any semiconductor device or card by using an electrostatic wrist strap and/or rubber mat.

Static electricity can harm system boards. Perform service at an ESD workstation and follow proper ESD procedures to reduce the risk of damage to components. One Stop Systems strongly encourages you to follow proper ESD procedures, which can include wrist straps and smocks, when servicing equipment.

You can also take the following steps to prevent damage from electrostatic discharge (ESD):

When unpacking a static-sensitive component from its shipping carton, do not remove the component's anti-static packaging material until you are ready to install the component in a computer. Just before unwrapping the anti-static packaging, be sure you are at an ESD workstation or are grounded.

When transporting a sensitive component, first place it in an anti-static container or packaging.


Handle all sensitive components at an ESD workstation. If possible, use anti-static floor pads and workbench pads.

Handle components and boards with care. Do not touch the components or contacts on a board. Hold a board by its edges or by its metal mounting bracket.


1 Product Information

PCIe x16 Gen 4 add-in card supporting dual Gen 4 hot-swap M.2 drives in removable carriers and dual SFF-8643 connectors supporting additional NVMe expansion. The hot-swap removable carriers provide interchangeability and flexibility at Gen 4 M.2 speeds for edge applications while providing scalability through the additional SFF-8643 internal connections.

Primary Card (OSS-537 board)




1.1 Removal Tray





1.2 PCIe M.2 Carrier Card

1.2.1 OSS-537 Board



1.2.2 OSS-568 Board



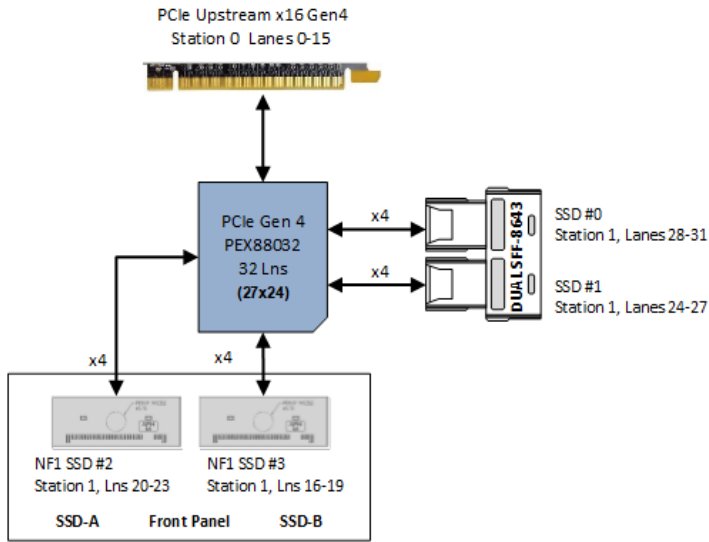
1.3 Features

- Quad PCIe 4.0 NVMe M.2 slots
- Hot-swap removable drive carriers
- Operates at up to 512Gb/s at PCIe 4.0 speeds
- Supports M.2 2242/2260/2280/22110/E1.S drives
- Supports PCIe 3.0 backward compatibility


1.4 General Specification

Form Factor	PCIe 4.0 x16 dual-width add-in card
Dimensions	8.10" x 4.38" (20.6 x 11.1 cm)
Bandwidth	Up to 512Gb/s
Drive Form Factors	Drive form factors supported: o <ul style="list-style-type: none"> • PCIe 4.0 NVMe M.2 2242, 2260, 2280, 22110 • PCIe 3.0 NVMe M.2 2242, 2260, 2280, 22110
Connectors	PCIe x16 card edge connector Dual SFF-8643 internal connectors (per board) Compliant to PCI-SIG PCI Express® External Cable Specification 4.0
Bracket	Standard full-height, dual-width
PCIe Switch	Broadcom PLX PEX 88032 16 GT/s 32-Lane PCIe 4.0 Switch DMA controller SSC Isolation
Power Consumption	60W
Operating Temperature	0°C to 40°C (200LFM airflow required) based on 1.7°C/W
Storage Temperature	-40°C to 85°C
Operating Humidity	20% to 80% relative humidity non-condensing
Storage Humidity	20% to 80% non-condensing
Agency Compliance	Designed to meet the following agency standards: <ul style="list-style-type: none"> • FCC – Part 15 Class A, 47CFR; Canada ICES-003, issue 4, Class A; Japan: VCCI, Class A' CE Emissions 2004-108EC • UL/IEC 60950-1; Canada: CSA C22.2 No. 60950-1; Argentina: IEC60950-1; IEC 60950-1 (CB Certificate and CB Test Report) • CE Mark (EN55022 Class A, EN60950-1, EN55024, EN61000-3-2, EN61000-3-3) • CISPR 22, CISPR 24, Class A; Australia/New Zealand AS/NZS CISPR 22, Class A • RoHS 6 of 6 compliance (Directive 2002/95/EC) • WEEE (EU 2012/19) & RoHS 3 (EU 2015/863)
Supported Operating Systems	Windows 10 & Windows 10 Pro Windows 2012 Server Ubuntu 16x

1.5 Block Diagram





1.6 Dimensions



1.7 M Key Connector

The OSS-537 and OSS-568 boards support M key edge connector



2 Hardware Requirements

This section provides the hardware parts needed for the OSS-537 board to work. It is strictly recommended to follow and use the hardware requirements listed below for the board to operate properly.

1. M.2 NVME SSD (Gen 4)
2. Recommended: Server type computer
3. Optional: Standard workstation (with x16 Gen4 PCIe slot) with good air flow and cooling.

Your computer must have sufficient cooling and airflow to prevent overheating of the M.2 media.

Operating Temperature	0°C to 40°C (200LFM airflow required) based on 1.7°C/W. For standard workstation, it is highly recommended using 20 CFM Fan or higher
-----------------------	--



You can measure the airflow by using an inexpensive gadget called "Anemometer Handheld Wind Speed Mete."

3 Software Requirements

1. Computer running Windows Server or Windows 10 and Windows 10 Pro
 - a. Windows Pro, driver is loaded automatically when card is installed and detected.
 - b. On Windows Server, requires no driver. OS loads the driver automatically.
2. Centos 7
3. Ubuntu 16

4 Hardware Installation

The following steps will guide you through the installation of your M.2 PCIe carrier board.


4.1 Installation-Procedures Overview

Following steps provide the exact sequence that needs to be followed:

1. Set the Quad M.2 Carrier board on a sturdy surface
2. Remove the module canister
3. Install the M.2 onto the circuit board
4. Secure the M.2 onto the circuit board
5. Turn OFF computer before installation
6. Remove cover from the computer
7. Remove the corresponding slot cover from computer chassis
8. Configure SW1 Dipswitch
9. Plug-in PCIe carrier board and secure it.
10. Slide the canister back into the PCIe carrier board
11. Power ON the computer
12. Perform Hardware check (Verify LED indicators)
13. Verify device installation (i.e., Windows Device Manager or Linux lspci tree)



4.4 Prepare PCIe Carrier Card

Place the “Quad M.2 Carrier” board on a sturdy surface




4.5 Remove the canister

Remove the tray / canister from the carrier board by pulling the tab (lever or ejector handle) to disengage.




Flip the canister to access the circuit board



4.6 Install M.2 module

Align the M.2 Key edge connector to the M.2 key edge socket on the circuit board




Slowly insert the media at ~30-degree position into the connector socket until it is fully seated.



4.7 Secure the Media

Secure the M.2 media.




Follow the steps below on how to install the canister in the PCIe card carrier in the PCIe cards in the computer.

IMPORTANT! It is important to install the PCIe Quad M.2 Carrier first in the computer before installing the canister.

4.8 Configure SW1 Dipswitch

Set the SW1 Dipswitch using the settings below before you install the card in the computer. All toggle switches are set to OFF.

1= OFF, 2= OFF, 3= OFF, 4= OFF





4.8.1 SW1 Dipswitch

Toggle Switch#	Description / Purpose	ON	OFF
1	Flash 0 or Flash 1 SBR select	Flash 1 (default)	Flash 0
2	Spare switch		
3	Spare switch		
4	PCIe Switch Debug Serial Port select	UART port	SDB port



4.9 Connect Data Cables

Plug in the HD Mini SAS cables to the board.





Use two HD Mini SAS Cables.

- Plug in the first cable to the SFF 8643 connectors, see photos below.
- Pay attention to the orientation of the cable connector, it has a small plastic latch that snaps in place when you connect it to the 8643 connectors.






Plug in the second HD Mini SAS cable. Make sure both cables are firmly seated and latched in.



4.10 Plugin the PCIe carrier board


Plug-in the PCIe carrier board to an open / available x16 slot. You can install the PCIe carrier board in a computer (motherboard) or in an OSS compatible Gen4 expansion unit (i.e., 4UP or EB4400).



Align the x16 and x1 edge connectors on top of the x16 slot connectors on the motherboard and gently push it down until both card edge connectors are firmly seated.





Secure the bracket with two screws.



4.11 Connect Power to Mezzanine Card


Connect Power to the Mezzanine Card. Use available external “4pin ATX Molex Female” power cable from the power supply and connect that to the J2 connector on the mezzanine card (OSS-568 board).




NOTE:

The x1 card edge connector is mainly for power. If a PCIe card slot is not available to use for the x1 card edge connector, you must connect the external Aux power cable to J2 connector. Photos below show the x1 card edge connector plugged-in vs not plugged-in to a card slot.

x1 PCIe Connector: **Not** plugged-in to PCIe card slot



x1 PCIe Connector: Plugged-in to PCIe card slot



If the x1 PCIe connector is already plugged-in to a card slot, the external power cable is no longer required.



4.13 Install Canister

- Slide the canister back into the enclosure of the PCIe carrier board.
- Push the lever forward to latch the canister in place.




4.14 Power ON the computer

Prior to powering ON the computer, verify that the PCIe card edge connector is fully seated as shown from the photos below.



Note:

When PCIe 4.0 Quad M.2 Carrier is installed in an OSS expansion unit or plugged-in to an OSS Gen4 expansion backplane, you must apply power to the OSS expansion unit or board first before powering on the host computer. See diagram below for supported expansion system setup.




5 Hardware Check

Once the host computer has booted up, verify that all LEDs are correctly illuminated on the carrier card. An operational PCIe board will show the following LEDs illuminated.

1. CE LINK, SSD RST and RESET LED
2. M.2 LINK LEDs
3. D2 LED

Power, Reset & CE LEDs



5.1 LED Definition

LED Name	When ON	When OFF
A and B Stat	Swap status still in development.	Normal for now
Core Stat	ARM core running when flashing	ARM core not running
CE	Card edge link status: solid on when Gen4 <ul style="list-style-type: none"> • On for Gen4, • Flashing fast for Gen3, • Flashing slow for Gen2 	Not linked
SSD RST	One or more M.2 is seated in the carrier	Stuck in reset
PWR GOOD	Power is present on the board	No Power
M.2 Link LED	M.2 media is present / detected can flash at different rates	M.2 media is missing / not detected. Data cables are not connected.


6 Verify OSS Device on Windows OS

6.1 Device Manager

Verify the hardware device in Windows Device Manager. As your Windows OS starts up, you will see a small message box popping-up in the lower-right corner of the screen to alert you that Windows has found new hardware.

Follow these steps on how to view and verify the OSS device via Windows Device Manager.

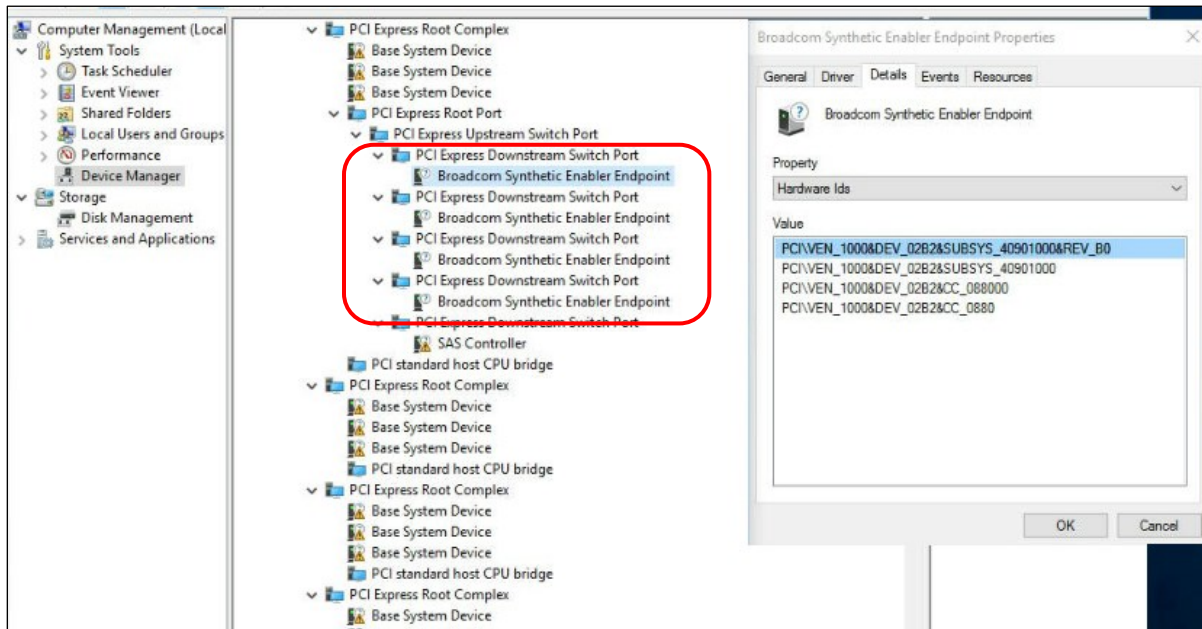
- Go to Device Manager
- Right Click on the Windows Logo on the left bottom of the system tray
- Click "View:
- Click "Devices by connection"



- Select and expand all the "PCI Express Root Complex"
- Select and expand "PCI Express Root Port"
- Select and expand "PCI Express Upstream Switch Port"
 - Four or five instances of "PCI Express Downstream Switch Port" will show up
 - Select and expand all "PCI Express Downstream Port"
 - Under each "PCI Express Downstream Port", you will see "Broadcom Synthetic Enabler Endpoint".

Note:

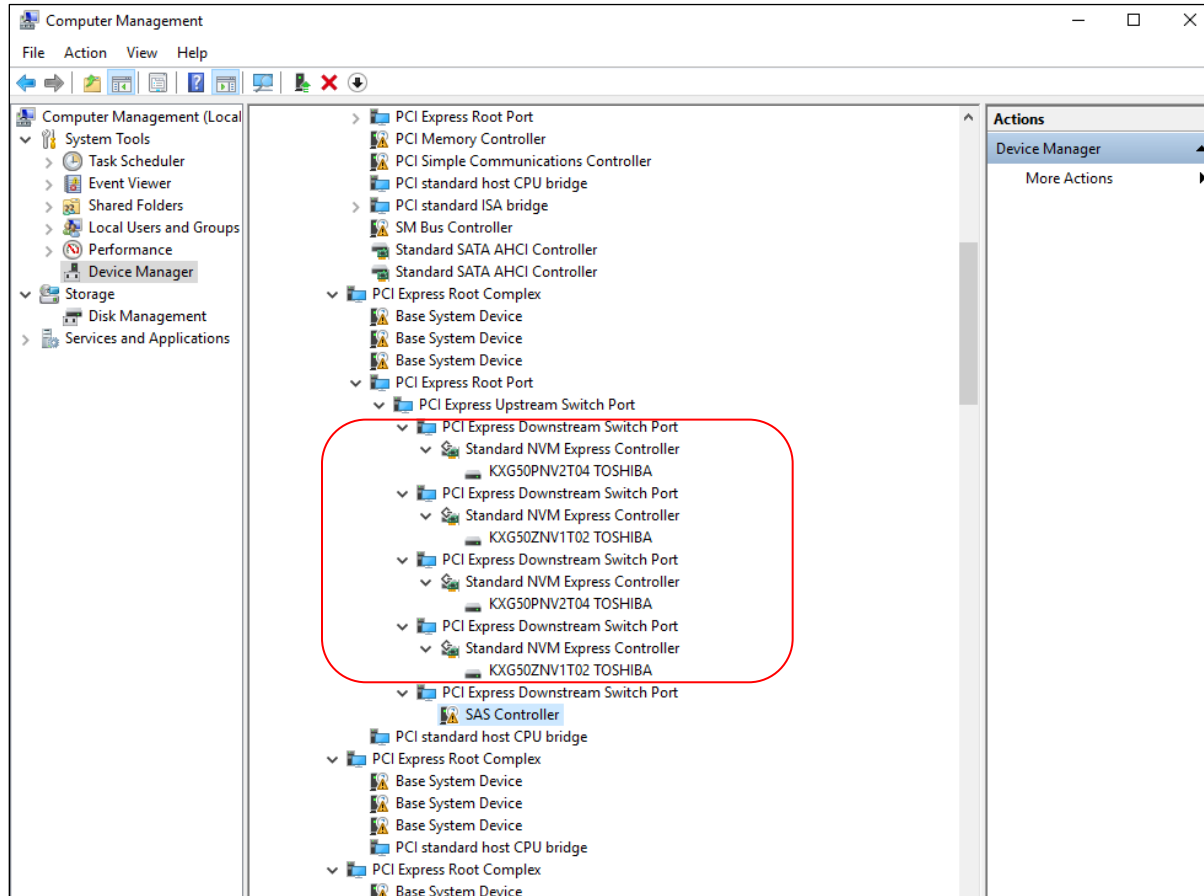
The "Broadcom Synthetic Enabler Endpoint" "PCI Vendor ID 1000, Device ID 02B2" will appear when no M.2 module is inserted or installed.



7 Verify M.2 Media on Windows OS

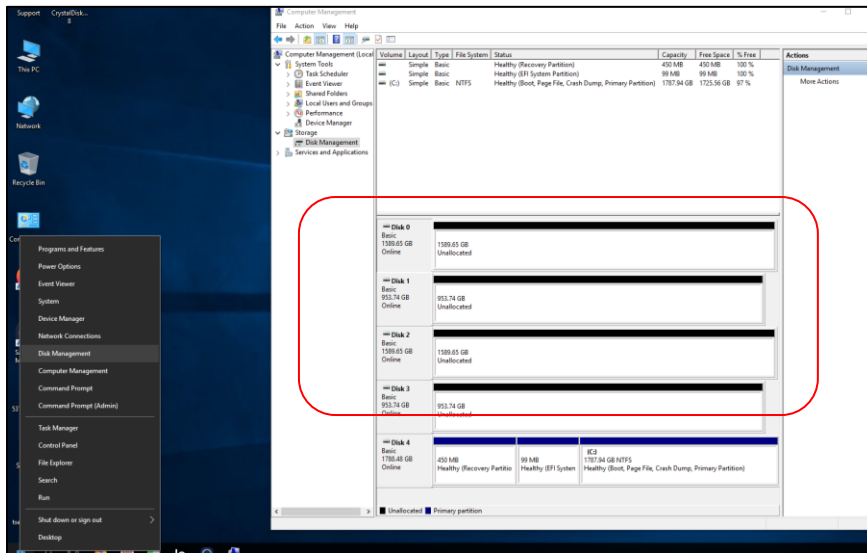
7.1 Device Manager

When installed properly, you will see the four **M.2 media devices and Standard NVM Express Controllers**. The screenshots below represent the hierarchy of a single “OSS PCIe 4.0 Quad M.2 Carrier” board with four M.2 media devices detected on Windows 10 host computer.



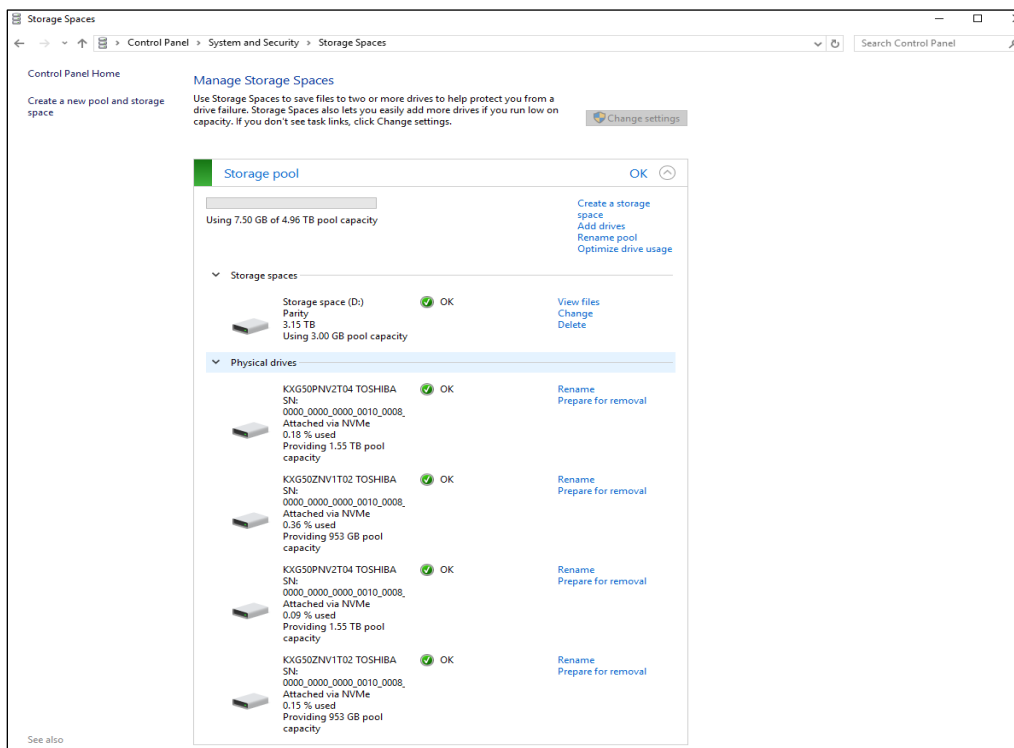
7.2 Disk Management

You can check and verify the M.2 media by using the Disk Management on Windows OS (i.e., Windows 10 Pro), see photo below.



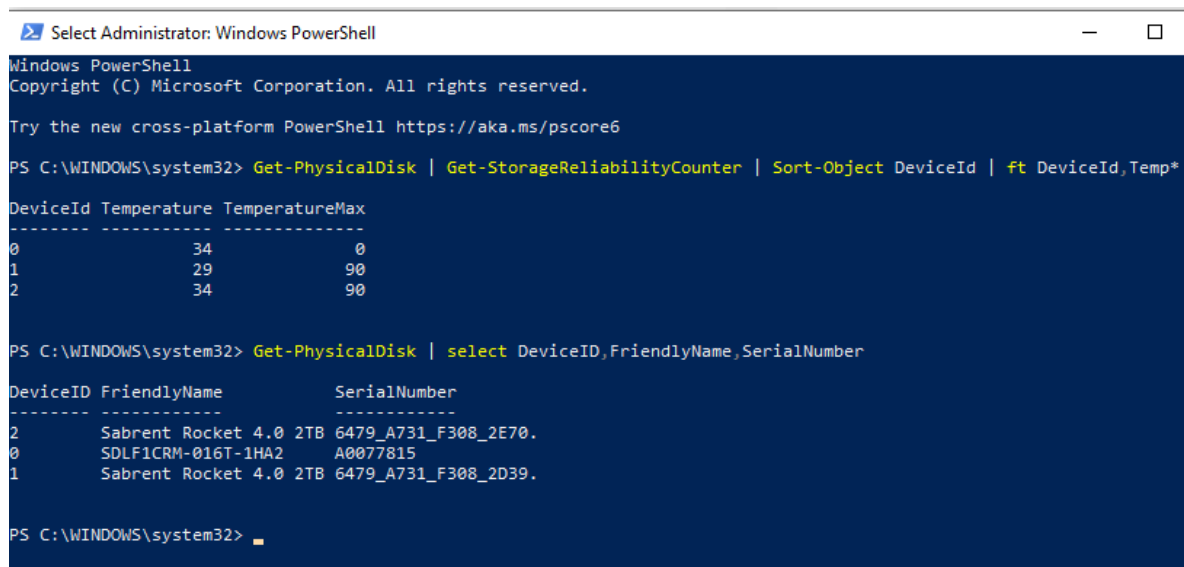
7.3 Storage Spaces

You can verify and manage the disk and create storage pool by using Microsoft "Storage Spaces" tool, see screenshot below.



7.4 Check Temperature

You can check the temperature of the M.2 media via “Windows PowerShell,” see screenshot below.



```
Select Administrator: Windows PowerShell
Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

Try the new cross-platform PowerShell https://aka.ms/pscore6

PS C:\WINDOWS\system32> Get-PhysicalDisk | Get-StorageReliabilityCounter | Sort-Object DeviceId | ft DeviceId,Temp*
DeviceId Temperature TemperatureMax
-----
0          34          90
1          29          90
2          34          90

PS C:\WINDOWS\system32> Get-PhysicalDisk | select DeviceID,FriendlyName,SerialNumber
DeviceID FriendlyName          SerialNumber
-----
2        Sabrent Rocket 4.0 2TB 6479_A731_F308_2E70.
0        SDLF1CRM-016T-1HA2  A0077815
1        Sabrent Rocket 4.0 2TB 6479_A731_F308_2D39.

PS C:\WINDOWS\system32>
```


8 Verify OSS Device on Linux OS

To check and verify if the OSS PCIe 4.0 Quad M.2 Carrier card is properly detected, use the following command lines.

lspci -vvt | grep Broadcom, see output below

```
root@Support:~# lspci -vvt | grep Broadcom
|          \-02.0-[17-1d]----00.0-[18-1d]--+00.0-[19]----00.0  Broadcom / LSI Device 02b2
|
|          +-01.0-[1a]----00.0  Broadcom / LSI Device 02b2
|          +-02.0-[1b]----00.0  Broadcom / LSI Device 02b2
|          +-03.0-[1c]----00.0  Broadcom / LSI Device 02b2
|          \-1f.0-[1d]----00.0  Broadcom / LSI Device 00b2
```

lspci -vvv | grep 02b2, see output below.

```
root@Support:~# lspci -vvv | grep 02b2
19:00.0 System peripheral: Broadcom / LSI Device 02b2 (rev b0)
1a:00.0 System peripheral: Broadcom / LSI Device 02b2 (rev b0)
1b:00.0 System peripheral: Broadcom / LSI Device 02b2 (rev b0)
1c:00.0 System peripheral: Broadcom / LSI Device 02b2 (rev b0)
```

lspci -vvt | grep 02b2, see output below.

```
root@Support:~# lspci -vvt | grep 02b2
|          \-02.0-[17-1d]----00.0-[18-1d]--+00.0-[19]----00.0  Broadcom / LSI Device 02b2
|
|          +-01.0-[1a]----00.0  Broadcom / LSI Device 02b2
|          +-02.0-[1b]----00.0  Broadcom / LSI Device 02b2
|          +-03.0-[1c]----00.0  Broadcom / LSI Device 02b2
```

lspci -vvv | grep c010 , see output below

```
root@Support:~# lspci -vvv | grep c010
17:00.0 PCI bridge: Broadcom / LSI Device c010 (rev b0) (prog-if 00 [Normal decode])
18:00.0 PCI bridge: Broadcom / LSI Device c010 (rev b0) (prog-if 00 [Normal decode])
18:01.0 PCI bridge: Broadcom / LSI Device c010 (rev b0) (prog-if 00 [Normal decode])
18:02.0 PCI bridge: Broadcom / LSI Device c010 (rev b0) (prog-if 00 [Normal decode])
18:03.0 PCI bridge: Broadcom / LSI Device c010 (rev b0) (prog-if 00 [Normal decode])
18:1f.0 PCI bridge: Broadcom / LSI Device c010 (rev b0) (prog-if 00 [Normal decode])
```

9 Verify M.2 Media on Linux OS

The M.2 storage device can be verified by typing the following command lines:

```
# nvme -list
```

```
root@Support:~# nvme list
Node          SN                      Model                      Namespace Usage                Format                FW Rev
-----
/dev/nvme0n1  774B600AK4DS           KXG50ZNV1T02 TOSHIBA                    1          1.02 TB / 1.02 TB      512 B + 0 B      AAGA4101
/dev/nvme1n1  774B600PK4DS           KXG50ZNV1T02 TOSHIBA                    1          1.02 TB / 1.02 TB      512 B + 0 B      AAGA4101
/dev/nvme2n1  38KF708MF3JP           KXG50PNV2T04 TOSHIBA                    1          1.71 TB / 1.71 TB      512 B + 0 B      AFGA4103
/dev/nvme3n1  38KF707IF3JP           KXG50PNV2T04 TOSHIBA                    1          1.71 TB / 1.71 TB      512 B + 0 B      AFGA4103
/dev/nvme4n1  11F0A005T4X8           KCD61LUL1T92                                1          39.71 GB / 1.92 TB      512 B + 0 B      0102
root@Support:~#
```

```
# lsblk -a
```

```
root@Support:~# lsblk -a
NAME        MAJ:MIN RM  SIZE RO TYPE MOUNTPOINT
loop0       7:0      0    4K  1 loop /snap/bare/5
loop1       7:1      0   55M  1 loop /snap/core18/1705
loop2       7:2      0   51M  1 loop /snap/snap-store/547
loop3       7:3      0  62.1M  1 loop /snap/gtk-common-themes/1506
loop4       7:4      0  219M  1 loop /snap/gnome-3-34-1804/72
loop5       7:5      0  56.4M  1 loop /snap/kdiskmark/53
loop6       7:6      0  32.5M  1 loop /snap/snapd/13640
loop7       7:7      0 260.7M  1 loop /snap/kde-frameworks-5-core18/32
loop8       7:8      0 240.8M  1 loop /snap/gnome-3-34-1804/24
loop9       7:9      0  49.8M  1 loop /snap/snap-store/433
loop10      7:10     0  27.1M  1 loop /snap/snapd/7264
loop11      7:11     0  65.2M  1 loop /snap/gtk-common-themes/1519
loop12      7:12     0  55.5M  1 loop /snap/core18/2246
loop13      7:13     0     0 loop
nvme0n1     259:0    0 953.9G  0 disk
├─nvme0n1p1 259:3    0  128M  0 part
nvme1n1     259:1    0 953.9G  0 disk
├─nvme1n1p1 259:5    0  128M  0 part
nvme3n1     259:2    0   1.6T  0 disk
├─nvme3n1p1 259:6    0  128M  0 part
nvme2n1     259:4    0   1.6T  0 disk
├─nvme2n1p1 259:7    0  128M  0 part
nvme4n1     259:8    0   1.8T  0 disk
├─nvme4n1p1 259:9    0  512M  0 part /boot/efi
└─nvme4n1p2 259:10   0   1.8T  0 part /
root@Support:~#
```

```
# lspci -vtt | grep NameOfM.2
```

Output below shows four Toshiba M.2 storage devices. You can replace the "NameofM.2" with your preferred M.2 device brand name.

```
root@Support:~# lspci -vtt | grep Toshiba
|          \-02.0-[17-1d]---00.0-[18-1d]--+00.0-[19]----00.0  Toshiba Corporation Device 0116
|
|          +-01.0-[1a]----00.0  Toshiba Corporation Device 0116
|          +-02.0-[1b]----00.0  Toshiba Corporation Device 0116
|          +-03.0-[1c]----00.0  Toshiba Corporation Device 0116
root@Support:~#
```

```
# lshw -short -class storage
```

```
root@Support:~# lshw -short -class storage
H/W path          Device              Class      Description
=====
/0/100/11.5      storage            C620 Series Chipset Family SSATA Controller [AHCI mode]
/0/100/17        storage            C620 Series Chipset Family SATA Controller [AHCI mode]
/0/101/2/0/0/0   storage            Toshiba Corporation
/0/101/2/0/1/0   storage            Toshiba Corporation
/0/101/2/0/2/0   storage            Toshiba Corporation
/0/101/2/0/3/0   storage            Toshiba Corporation
/0/101/2/0/1f/0  scsi0             storage    Broadcom / LSI
/0/104/2/0       storage            NVMe SSD Controller Cx6
root@Support:~#
```

```
# lsblk -f
```

```
root@Support:~# lsblk -f
NAME                FSTYPE LABEL UUID                                 FSAVAIL FSUSE% MOUNTPOINT
loop0               squashfs
loop1               squashfs
loop2               squashfs
loop3               squashfs
loop4               squashfs
loop5               squashfs
loop6               squashfs
loop7               squashfs
loop8               squashfs
loop9               squashfs
loop10              squashfs
loop11              squashfs
loop12              squashfs
loop13              squashfs
nvme0n1
├─nvme0n1p1
nvme1n1
├─nvme1n1p1
nvme3n1
├─nvme3n1p1
nvme2n1
├─nvme2n1p1
nvme4n1
├─nvme4n1p1 vfat      BA6A-40D0          503.2M    2% /boot/efi
├─nvme4n1p2 ext4       a2021cb5-b175-4cf2-b65d-9202256ff0a0 1.6T     0% /
root@Support:~#
```

```
# fdisk -l | grep nvme
```

```
root@Support:~# fdisk -l | grep nvme
Disk /dev/nvme0n1: 953.89 GiB, 1024209543168 bytes, 2000409264 sectors
/dev/nvme0n1p1 34 262177 262144 128M Microsoft reserved
Disk /dev/nvme1n1: 953.89 GiB, 1024209543168 bytes, 2000409264 sectors
/dev/nvme1n1p1 34 262177 262144 128M Microsoft reserved
Disk /dev/nvme3n1: 1.57 TiB, 1707006873600 bytes, 3333997800 sectors
/dev/nvme3n1p1 34 262177 262144 128M Microsoft reserved
Disk /dev/nvme2n1: 1.57 TiB, 1707006873600 bytes, 3333997800 sectors
/dev/nvme2n1p1 34 262177 262144 128M Microsoft reserved
Disk /dev/nvme4n1: 1.76 TiB, 1920383410176 bytes, 3750748848 sectors
/dev/nvme4n1p1 2048 1050623 1048576 512M EFI System
/dev/nvme4n1p2 1050624 3750748159 3749697536 1.8T Linux filesystem
root@Support:~#
```

blkid

```

root@Support:~# blkid
/dev/nvme4n1p2: UUID="a2021cb5-b175-4cf2-b65d-9202256ff0a0" TYPE="ext4" PARTUUID="de86726a-0812-4922-ab89-7375f7cb3adb"
/dev/loop0: TYPE="squashfs"
/dev/loop1: TYPE="squashfs"
/dev/loop2: TYPE="squashfs"
/dev/loop3: TYPE="squashfs"
/dev/loop4: TYPE="squashfs"
/dev/loop5: TYPE="squashfs"
/dev/loop6: TYPE="squashfs"
/dev/loop7: TYPE="squashfs"
/dev/nvme0n1p1: PARTLABEL="Microsoft reserved partition" PARTUUID="c6103456-33c0-11ec-83c0-b03af2b6059f"
/dev/nvme1n1p1: PARTLABEL="Microsoft reserved partition" PARTUUID="c610345d-33c0-11ec-83c0-b03af2b6059f"
/dev/nvme3n1p1: PARTLABEL="Microsoft reserved partition" PARTUUID="c6103453-33c0-11ec-83c0-b03af2b6059f"
/dev/nvme2n1p1: PARTLABEL="Microsoft reserved partition" PARTUUID="c610345a-33c0-11ec-83c0-b03af2b6059f"
/dev/nvme4n1p1: UUID="BA6A-40D0" TYPE="vfat" PARTLABEL="EFI System Partition" PARTUUID="87586ba2-8beb-445c-8804-604ba944e880"
/dev/loop8: TYPE="squashfs"
/dev/loop9: TYPE="squashfs"
/dev/loop10: TYPE="squashfs"
/dev/loop11: TYPE="squashfs"
/dev/loop12: TYPE="squashfs"
/dev/loop13: TYPE="squashfs"
root@Support:~# █

```

9.1 Check Temperature

To check on the temperature of the M.2 storage devices on Linux (i.e., Ubuntu 18x), use the command line "nvme smart-log /dev/nvme". You would need to identify the storage device name first, use the command line "nvme list".

```

root@Support:~# nvme list
Node      SN                      Model                      Namespace Usage          Format          FW Rev
-----
/dev/nvme0n1  774B600AK4DS          KXG50ZNV1T02 TOSHIBA          1          1.02 TB / 1.02 TB  512 B + 0 B  AAGA4101
/dev/nvme1n1  774B600PK4DS          KXG50ZNV1T02 TOSHIBA          1          1.02 TB / 1.02 TB  512 B + 0 B  AAGA4101
/dev/nvme2n1  38KF708MF3JP          KXG50PNV2T04 TOSHIBA          1          1.71 TB / 1.71 TB  512 B + 0 B  AFGA4103
/dev/nvme3n1  38KF707IF3JP          KXG50PNV2T04 TOSHIBA          1          1.71 TB / 1.71 TB  512 B + 0 B  AFGA4103
/dev/nvme4n1  11F0A005T4X8          KCD61LUL1T92          1          41.58 GB / 1.92 TB  512 B + 0 B  0102
root@Support:~# nvme smart-log /dev/nvme0n1
Smart Log for NVME device:nvme0n1 namespace-id:ffffffff
critical_warning      : 0
temperature           : 46 C
available_spare        : 100%
available_spare_threshold : 10%
percentage_used        : 31%
data_units_read        : 1,300,291,316
data_units_written     : 359,469,447
host_read_commands    : 1,826,706,712
host_write_commands    : 806,385,283
controller_busy_time  : 14,071
power_cycles           : 261
power_on_hours         : 15,125
unsafe_shutdowns      : 219
media_errors           : 0
num_err_log_entries   : 0
Warning Temperature Time : 0
Critical Composite Temperature Time : 0
Temperature Sensor 1   : 46 C
Thermal Management T1 Trans Count : 0
Thermal Management T2 Trans Count : 0
Thermal Management T1 Total Time : 0
Thermal Management T2 Total Time : 0
root@Support:~# █

```

9.2 Check OSS Device Speed and Linkwidth

To check the “Speed and Linkwidth” of the OSS devices use the command lines below.

First, you need to identify the bus address, type "lspci -vtt | grep Broadcom". See output below

```
root@Support:~# lspci -vtt | grep Broadcom
|                                     \-1f.0-[1d]----00.0  Broadcom / LSI Device 00b2
root@Support:~#
```

Once you have identified the bus number, use the command lines below to retrieve the “Link Status and Link Capability”

```
# lspci -vvv -s 1d:00.0 | grep 'LnkSta|LnkCap'
```

```
root@Support:~# lspci -vvv -s 1d:00.0 | grep 'LnkSta|LnkCap'
\pcilib: sysfs_read_vpd: read failed: Input/output error
LnkCap: Port #0, Speed 16GT/s, Width x16, ASPM not supported
LnkSta: Speed 16GT/s (ok), Width x16 (ok)
LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete-, EqualizationPhase1-
root@Support:~#
```

If you need to retrieve the full detail, use “lspci -vvv 1d:00.0”

```
root@Support:~# lspci -vvv -s 1d:00.0
1d:00.0 Serial Attached SCSI controller: Broadcom / LSI Device 00b2 (rev b0)
Subsystem: Broadcom / LSI Device 4090
Control: I/O- Mem+ BusMaster+ SpecCycle- MemWINV- VGASnoop- ParErr- Stepping- SERR- FastB2B- DisINTx-
Status: Cap+ 66MHz- UDF- FastB2B- ParErr- DEVSEL=fast >TAbort- <TAbort- <MAbort- >SERR- <PERR- INTx-
Latency: 0, Cache Line Size: 32 bytes
NUMA node: 0
Region 0: Memory at 201fffe00000 (64-bit, prefetchable) [size=16K]
Capabilities: [40] Power Management version 3
Flags: PMEClk- DSI- D1- D2- AuxCurrent=0mA PME(D0+,D1-,D2-,D3hot+,D3cold+)
Status: D0 NoSoftRst+ PME-Enable- Dscl=0 Dscale=0 PME-
Capabilities: [68] Express (v2) Endpoint, MSI 00
DevCap: MaxPayload 2048 bytes, PhantFunc 0, Latency L0s <64ns, L1 <1us
ExtTag- AttnBtn- AttnInd- PwrInd- RBE+ FLReset+ SlotPowerLimit 0.000W
DevCtl: CorrErr- NonFatalErr- FatalErr- UnsupReq-
RlxdOrd+ ExtTag+ PhantFunc- AuxPwr- NoSnoop+ FLReset-
MaxPayload 256 bytes, MaxReadReq 4096 bytes
DevSta: CorrErr- NonFatalErr- FatalErr- UnsupReq- AuxPwr- TransPend-
LnkCap: Port #0, Speed 16GT/s, Width x16, ASPM not supported
ClockPM- Surprise- LLActRep- BwNot- ASPMOptComp+
LnkCtl: ASPM Disabled; RCB 64 bytes Disabled- CommClk-
ExtSynch- ClockPM- AutWidDis- BWInt- AutBWInt-
LnkSta: Speed 16GT/s (ok), Width x16 (ok)
TrErr- Train- SlotClk- DLActive- BWMgmt- ABWMgmt-
DevCap2: Completion Timeout: Not Supported, TimeoutDis-, NROPrPrP-, LTR-
10BitTagComp-, 10BitTagReq-, OBFF Not Supported, ExtFnt-, EETLPPrefix-
EmergencyPowerReduction Not Supported, EmergencyPowerReductionInit-
FRS-, TPHComp-, ExtTPHComp-
AtomicOpsCap: 32bit- 64bit- 128bitCAS-
DevCtl2: Completion Timeout: 50us to 50ms, TimeoutDis-, LTR-, OBFF Disabled
AtomicOpsCtl: ReqEn-
LnkCtl2: Target Link Speed: 16GT/s, EnterCompliance- SpeedDis-
Transmit Margin: Normal Operating Range, EnterModifiedCompliance- ComplianceSOS-
Compliance De-emphasis: -6dB
LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete-, EqualizationPhase1-
EqualizationPhase2-, EqualizationPhase3-, LinkEqualizationRequest-
Capabilities: [48] MSI: Enable- Count=1/1 Maskable+ 64bit+
Address: 0000000000000000 Data: 0000
Masking: 00000000 Pending: 00000000
Capabilities: [a4] MSI-X: Enable+ Count=1 Masked-
Vector table: BAR=0 offset=00002000
PBA: BAR=0 offset=00003000
Capabilities: [b0] Vital Product Data
pcilib: sysfs_read_vpd: read failed: Input/output error
Not readable
Capabilities: [e0] Vendor Specific Information: Len=20 <?>
Capabilities: [100 v1] Device Serial Number 00-17-6a-00-00-00-b2-00
Kernel driver in use: mpt3sas
Kernel modules: mpt3sas
root@Support:~#
```

9.3 Check M.2 Speed and Linkwidth

Type 'lspci -vtt | grep VendorName. Replace "Toshiba" with the vendor's name of you M.2 media.

```
root@Support:~# lspci -vtt | grep Toshiba
| \-02.0-[17-1d]---00.0-[18-1d]--+00.0-[19]---00.0 Toshiba Corporation Device 0116
|                                     +-01.0-[1a]---00.0 Toshiba Corporation Device 0116
|                                     +-02.0-[1b]---00.0 Toshiba Corporation Device 0116
|                                     +-03.0-[1c]---00.0 Toshiba Corporation Device 0116
```

Then type "lspci -vvv -s 19:00.0 | grep 'LnkSta|LnkCap'

```
root@Support:~# lspci -vvv -s 19:00.0 | grep 'LnkSta|LnkCap'
LnkCap: Port #0, Speed 8GT/s, Width x4, ASPM not supported
LnkSta: Speed 8GT/s (ok), Width x4 (ok)
LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete+, EqualizationPhase1+
root@Support:~#
```

To retrieve the detailed info, use "lspci -vvvv -s 19:00.0". Replace the 19:00.00 with the bus number of your M.2 device (XX:XX.X)

```
root@Support:~# lspci -vvvv -s 19:00.0
19:00.0 Non-Volatile memory controller: Toshiba Corporation Device 0116 (prog-if 02 [NVM Express])
Subsystem: Toshiba Corporation Device 0001
Physical Slot: 16
Control: I/O- Mem+ BusMaster+ SpecCycle- MemWINV- VGASnoop- ParErr+ Stepping- SERR+ FastB2B- DisINTx+
Status: Cap+ 66MHz- UDF- FastB2B- ParErr- DEVSEL=fast >TAbort- <TAbort- <MAbort- >SERR- <PERR- INTx-
Latency: 0, Cache Line Size: 32 bytes
Interrupt: pin A routed to IRQ 18
NUMA node: 0
Region 0: Memory at #6200000 (64-bit, non-prefetchable) [size=16K]
Capabilities: [40] Express (v2) Endpoint, MSI 00
DevCap: MaxPayload 256 bytes, PhantFunc 0, Latency L0s unlimited, L1 unlimited
ExtTag- AttnBtm- AttnInd- PwrInd- RBE+ FLReset- SlotPowerLimit 25.000W
DevCtl: CorrErr- NonFatalErr- FatalErr- UnsupReq-
RlxdOrd- ExtTag- PhantFunc- AuxPwr- NoSnoop- FLReset-
MaxPayload 256 bytes, MaxReadReq 4096 bytes
DevSta: CorrErr- NonFatalErr- FatalErr- UnsupReq- AuxPwr- TransPnd-
LnkCap: Port #0, Speed 8GT/s, Width x4, ASPM not supported
ClockPM- Surprise- LLActRep- BwNot- ASPMOptComp+
LnkCtl: ASPM Disabled; RCB 64 bytes Disabled- ComCtl-
ExtSynch- ClockPM- AutWidDis- BWInt- AutBWInt-
LnkSta: Speed 8GT/s (ok), Width x4 (ok)
TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
DevCap2: Completion Timeout: Range AB, TimeoutDis+, NROPrPrP-, LTR+
10BitTagComp-, 10BitTagReq-, OBFF Not Supported, ExtFmt+, EETLPPrefix-
EmergencyPowerReduction Not Supported, EmergencyPowerReductionInit-
FRS-, TPHComp-, ExtTPHComp-
AtomicOpsCap: 32bit- 64bit- 128bitCAS-
DevCtl2: Completion Timeout: 50us to 50ms, TimeoutDis-, LTR-, OBFF Disabled
AtomicOpsCtl: ReqEn-
LnkCtl2: Target Link Speed: 8GT/s, EnterCompliance- SpeedDis-
Transmit Margin: Normal Operating Range, EnterModifiedCompliance- ComplianceSOS-
Compliance De-emphasis: -6dB
LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete+, EqualizationPhase1+
EqualizationPhase2+, EqualizationPhase3+, LinkEqualizationRequest-
Capabilities: [80] Power Management version 3
Flags: PMEClk- DSI- D1- D2- AuxCurrent=0mA PME(D0-, D1-, D2-, D3hot-, D3cold-)
Status: D0 NoSoftRst+ PME-Enable- DSel=0 DScale=0 PME-
Capabilities: [90] MSI: Enable- Count=1/32 Maskable+ 64bit+
Address: 0000000000000000 Data: 0000
Masking: 00000000 Pending: 00000000
Capabilities: [b0] MSI-X: Enable+ Count=32 Masked-
Vector table: BAR=0 offset=00002000
PBA: BAR=0 offset=00003000
Capabilities: [100 v2] Advanced Error Reporting
UESta: DLP- SDES- TLP- FCP- CplttO- CplttAbrt- UnxCpltt- RxOF- MalfTLP- ECRC- UnsupReq- ACSViol-
UEMsk: DLP- SDES- TLP- FCP- CplttTO- CplttAbrt+ UnxCpltt- RxOF+ MalfTLP+ ECRC- UnsupReq- ACSViol-
UESvrt: DLP- SDES- TLP- FCP- CplttTO- CplttAbrt- UnxCpltt- RxOF- MalfTLP- ECRC- UnsupReq- ACSViol-
CESSta: RxErr- BadTLP- BadDLLP- Rollover- Timeout+ AdvNonFatalErr-
CEMsk: RxErr+ BadTLP+ BadDLLP+ Rollover+ Timeout+ AdvNonFatalErr+
AERCap: First Error Pointer: 00, ECRGenCap- ECRCGenEn- ECRCChkCap- ECRCChkEn-
MultHdrRecCap- MultHdrRecEn- TLPPrfxPres- HdrLogCap-
HeaderLog: 00000000 00000000 00000000 00000000
Capabilities: [200 v1] Latency Tolerance Reporting
Max snoop latency: 0ns
Max no snoop latency: 0ns
Capabilities: [300 v1] Secondary PCI Express
LnkCtl3: LnkEqIntrruptEn-, PerformEqu-
LaneErrStat: 0
Capabilities: [400 v1] L1 PM Substates
L1SubCap: PCI-PM L1.2+ PCI-PM L1.1- ASPM L1.2+ ASPM L1.1- L1 PM Substates+
PortCommonModeRestoreTime=60us PortPowerOnTime=100us
L1SubCtl1: PCI-PM L1.2- PCI-PM L1.1- ASPM L1.2- ASPM L1.1-
T.CommonMode=0us LTR1.2_Threshold=0ns
L1SubCtl2: T_PwrOn=10us
Kernel driver in use: nvme
Kernel modules: nvme
```

9.4 PCIe Device Detailed Info

You can view and capture the detailed information of the device by using the command line “`lspci -vvv -s XX:xx.x` (see screenshot below).

```

root@localhost:~# lspci -vvv -s 0b:00.0
+03.1-[09-0f]----00.0-[0a-0f]---+00.0-[0b]----00.0 Phison Electronics Corporation E16 PCIe4 NVMe Controller
|
+01.0-[0c]----00.0 Phison Electronics Corporation E16 PCIe4 NVMe Controller
[root@localhost ~]# lspci -vvv -s 0b:00.0
0b:00.0 Non-Volatile memory controller: Phison Electronics Corporation E16 PCIe4 NVMe Controller (rev 01) (prog-if 02 [NVM Express])
Subsystem: Phison Electronics Corporation E16 PCIe4 NVMe Controller
Physical Slot: 16
Control: I/O- Mem+ BusMaster+ SpecCycle- MemWINV- VGASnoop- ParErr- Stepping- SERR- FastB2B- DisINTx+
Status: Cap+ 66MHz- UDF- FastB2B- ParErr- DEVSEL=fast >TAbort- <TAbort- <MAbort- >SERR- <PERR- INTx-
Latency: 0, Cache Line Size: 64 bytes
Interrupt: pin A routed to IRQ 62
NUMA node: 0
Region 0: Memory at f6c00000 (64-bit, non-prefetchable) [size=16K]
Capabilities: [80] Express (v2) Endpoint, MSI 00
  DevCap: MaxPayload 256 bytes, PhantFunc 0, Latency L0s unlimited, L1 unlimited
    ExtTag+ AttnBtn- AttnInd- PwrInd- RBE+ FLReset+ SlotPowerLimit 25.000W
  DevCtl: Report errors: Correctable+ Non-Fatal+ Fatal+ Unsupported+
    RixOrd+ ExtTag- PhantFunc- AuxPwr- NoSnoop+ FLReset-
    MaxPayload 256 bytes, MaxReadReq 512 bytes
  DevSta: CorrErr- UncorrErr- FatalErr- UnsuppReq- AuxPwr- TransPend-
  LnkCap: Port #1, Speed 16GT/s, Width x4, ASPM not supported, Exit Latency L0s unlimited, L1 unlimited
    ClockPM- Surprise- LLActRep- BwNot- ASPMOptComp+
  LnkCtl: ASPM Disabled; RC6 64 bytes Disabled- CommClk-
    ExtSynch- ClockPM- AutWidDis- BWInt- AutBWInt-
  LnkSta: Speed 16GT/s, Width x4, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
  DevCap2: Completion Timeout: Range ABCD, TimeoutDis+, LTR+, OBFF Not Supported
  DevCtl2: Completion Timeout: 50us to 50ms, TimeoutDis-, LTR-, OBFF Disabled
  LnkCtl2: Target Link Speed: 16GT/s, EnterCompliance- SpeedDis-
    Transmit Margin: Normal Operating Range, EnterModifiedCompliance- ComplianceSOS-
    Compliance De-emphasis: -6dB
  LnkSta2: Current De-emphasis Level: -6dB, EqualizationComplete+, EqualizationPhase1+
    EqualizationPhase2+, EqualizationPhase3+, LinkEqualizationRequest-
  Capabilities: [d0] MSI-X: Enable+ Count=9 Masked-
    Vector table: BAR=0 offset=00002000
    PBA: BAR=0 offset=00003000
  Capabilities: [e0] MSI: Enable- Count=1/8 Maskable- 64bit+
    Address: 0000000000000000 Data: 0000
  Capabilities: [f8] Power Management version 3
    Flags: PMEClk- DSI- D1- D2- AuxCurrent=0mA PME(D0-,D1-,D2-,D3hot-,D3cold-)
    Status: D0 NoSoftRst+ PME-Enable- DSel=0 DScale=0 PME-
  Capabilities: [100 v1] Latency Tolerance Reporting
    Max snoop latency: 0ns
    Max no snoop latency: 0ns
  Capabilities: [110 v1] L1 PM Substates
    L1SubCap: PCI-PM L1.2+ PCI-PM L1.1+ ASPM L1.2+ ASPM L1.1+ L1_PM_Substates+
    PortCommonModeRestoreTime=10us PortTPowerOnTime=300us
  Capabilities: [128 v1] Alternative Routing-ID Interpretation (ARI)
    ARICap: MFVC- ACS-, Next Function: 0
    ARICtl: MFVC- ACS-, Function Group: 0
  Capabilities: [1e0 v1] #25
  Capabilities: [200 v2] Advanced Error Reporting
    UESSta: DLP- SDES- TLP- FCP- CmpItTO- CmpItAbrt- UnxCmpIt- RxOF- MalfTLP- ECRC- UnsupReq- ACSViol-
    UEMsk: DLP- SDES- TLP- FCP- CmpItTO- CmpItAbrt- UnxCmpIt- RxOF- MalfTLP- ECRC- UnsupReq- ACSViol-
    UESvrt: DLP+ SDES- TLP- FCP+ CmpItTO- CmpItAbrt- UnxCmpIt- RxOF- MalfTLP+ ECRC- UnsupReq- ACSViol-
    CESta: RxErr- BadTLP- BadDLLP- Rollover- Timeout- NonFatalErr-
    CEMsk: RxErr- BadTLP- BadDLLP- Rollover- Timeout- NonFatalErr+
    AERCap: First Error Pointer: 00, GenCap- CGenEn- ChkCap+ ChkEn-
  Capabilities: [300 v1] #19
  Capabilities: [340 v1] #26
  Capabilities: [378 v1] #27
  Kernel driver in use: nvme
  Kernel modules: nvme
[root@localhost ~]#

```

9.5 Smartctl

Smartctl (**Self-Monitoring, Analysis and Reporting Technology**) is a command line utility or a tool that performs SMART tasks such as printing the SMART self-test and error logs, enabling, and disabling SMART automatic testing, and initiating device self-tests.

Smartctl command allows you to check for errors and extract info regarding the disks that are used, see example below

```
[root@localhost etc]# smartctl -a /dev/nvme0n1
smartctl 7.0 2018-12-30 r4883 [x86_64-linux-3.10.0-1062.1.2.el7.x86_64] (local build)
Copyright (C) 2002-18, Bruce Allen, Christian Franke, www.smartmontools.org

=== START OF INFORMATION SECTION ===
Model Number:          Sabrent Rocket 4.0 2TB
Serial Number:         7F600703139D02311888
Firmware Version:      RKT401.2
PCI Vendor/Subsystem ID: 0x1987
IEEE OUI Identifier:   0x6479a7
Total NVM Capacity:    2,000,398,934,016 [2.00 TB]
Unallocated NVM Capacity: 0
Controller ID:         1
Number of Namespaces: 1
Namespace 1 Size/Capacity: 2,000,398,934,016 [2.00 TB]
Namespace 1 Formatted LBA Size: 512
Namespace 1 IEEE EUI-64: 6479a7 31f3082e70
Local Time is:         Mon Aug 10 21:17:56 2020 EDT
Firmware Updates (0x12): 1 Slot, no Reset required
Optional Admin Commands (0x0017): Security Format Fmw_DL Self_Test
Optional NVM Commands (0x005d): Comp DS Mngmt Wr_Zero Sav/Sel_Feat Timestmp
Maximum Data Transfer Size: 512 Pages
Warning Comp. Temp. Threshold: 90 Celsius
Critical Comp. Temp. Threshold: 95 Celsius

Supported Power States
St Op   Max   Active   Idle   RL RT WL WT  Ent_Lat  Ex_Lat
0 +     9.78W -       -     0 0 0 0      0        0
1 +     6.75W -       -     1 1 1 1      0        0
2 +     5.23W -       -     2 2 2 2      0        0
3 -     0.0490W -     -     3 3 3 3    2000    2000
4 -     0.0018W -     -     4 4 4 4   25000   25000

Supported LBA Sizes (NSID 0x1)
Id Fmt  Data  Metadt  Rel_Perf
0 +     512    0      2
1 -    4096    0      1

=== START OF SMART DATA SECTION ===
SMART overall-health self-assessment test result: PASSED

SMART/Health Information (NVMe Log 0x02)
Critical Warning:          0x00
Temperature:              28 Celsius
Available Spare:          100%
Available Spare Threshold: 5%
Percentage Used:          0%
Data Units Read:          13,608,598 [6.96 TB]
Data Units Written:        12,949,704 [6.63 TB]
Host Read Commands:       163,551,249
Host Write Commands:      211,987,584
Controller Busy Time:     81
Power Cycles:              204
Power On Hours:           296
Unsafe Shutdowns:        63
Media and Data Integrity Errors: 0
Error Information Log Entries: 186
Warning Comp. Temperature Time: 0
Critical Comp. Temperature Time: 0

Error Information (NVMe Log 0x01, max 63 entries)
No Errors Logged
```


10 Benchmark Performance (Read & Write)

Performance read and write test results of FOUR M.2 modules mounted on the OSS card carrier.


- Computer and OS: SuperMicro X12DPG-QT6 running Windows 10 Pro
- Benchmark Software: CrystallDisk Mark 8.0.1 x6 2007-2021
- M.2 Module: 4 Sabrent Rocket 4.0 1TB

10.1 Disk Management Tool

Windows 10 Disk Management is a built-in utility that allows you to see and manage any internal and external hard drives connected to your computer.


1 st Test	SEQ READ	SEQ WRITE
Volume: Striped (all four M.2). NTFS	11 GB /s	12 GB /s

Striped Volume: A striped volume (RAID 0) combines areas of free space from multiple hard disks (anywhere from 2 to 32) into one logical volume.




2 nd Test	SEQ READ	SEQ WRITE
Spanned Volume (all four M.2), NTFS	5 GB /s	4 GB /s

Spanned Volume: A spanned volume combines areas of unallocated space from multiple disks into one logical volume.



3 rd Test	SEQ READ	SEQ WRITE
Mirrored Volume (You can only setup two M2 devices for Mirrored Volume). NTFS	9 GB /s	4 GB /s


Mirrored Volume: is a fault-tolerant dynamic volume. It provides data redundancy by using two copies of volume or copying data stored on the volume.



4 th Test	SEQ READ	SEQ WRITE
Testing single M.2, installed directly on the Supermicro motherboard. NTFS	5.009 GB /s	4.251 GB /s



5 th Test	SEQ READ	SEQ WRITE
Testing single M.2, installed directly on the PCIe carrier card	5.006 GB /s	4.244 GB /s



11 FAQ's

Q: Do I need to plug in the external AUX power cable to the mezzanine card (OSS-568)?

A: Yes, if the x1 card edge connector on the mezzanine card is not plugged-in to a card slot. The mezzanine card requires power.

Q: Is it OK to have both external Aux power cable and the x1 card edge connected? Will this damage the card?

A: Yes, it is OK. It will cause no electrical failure and will not damage the card.

Q: Can I mix and match different M2 brand / vendor?

A: Not recommended, as this can affect the performance of the card and some M2 media may not initialize properly.

Q: Do I need to install driver for the PCIe M.2 carrier card on Windows OS (i.e., Windows 10 and or Server)?

A: No, driver is automatically loaded by the OS when the PCIe carrier card is detected.

Q: Can I hot-plug or remove the canister (containing M.2 module) while the computer is ON (Windows)?

A: Yes, you can hot-plug the canister if there is no activity running (i.e., moving files in and out).

Q: Do I need to install driver for the M.2 module / media?

A: There are few M.2 module / media will require driver to achieve full performance and for proper operation. Please consult the manual associated with the M.2 module that you are using.

Q: Will the M.2 PCIe carrier card works on Linux based OS?

A: Yes, the card will work on Ubuntu and Centos.

Q: Only two out of four M.2 modules are detected. What should I check?

A: Check the data cables, make sure both data cables are connected between the primary card and the mezzanine card.

Q: The OS is not recognizing the other two M.2 modules that are mounted on the mezzanine card. What should I check?

A: Check the external power cable is attached to the mezzanine card or the x1 card edge connector is plugged-in to a PCIe slot.

Q: Can I remove the PCIe card carrier from the PCIe card slot while computer is running?

A: No, do not remove the PCIe card carries while the computer is ON. You must power down the computer first before removing it from the motherboard.

12 Contacting Technical Support

Our support department can be reached by fax at (858) 530-2733 or by phone at (858) 530-2511. Support is available Monday through Friday, 8:00 AM to 5:00 PM PT. When contacting One Stop Systems Technical Support, please be sure to include the following information:

- | | |
|------------------|--|
| 1) Name | 7) Serial Number |
| 2) Company Name | 8) Computer Make |
| 3) Phone Number | 9) Computer Model |
| 4) Fax Number | 10) Operating System and Version |
| 5) Email Address | 11) Make/Model of PCI cards in expansion chassis |
| 6) Model Number | 12) Detailed description of the problem |

You can also visit our web site at: <https://www.onestopsystems.com/support-0>

To submit a support ticket or case, use our OSS Online Support portal: <https://onestopsystems.desk.com/customer/portal/emails/new>

For a quick response, use the Technical Support and RMA Request Form available in the Support Section of the website. Simply complete the form with all required information. Please make sure that your problem description is sufficiently detailed to help us understand your problem. For example: Do not say "Won't boot up." Do say "Tried all the steps in the Troubleshooting Section and it still won't boot up."

For faster diagnosis of your problem, please run the two utility programs described in the following sections and include the diagnostic files they generate with your email.

11 Returning Merchandise to One Stop Systems

If factory service is required, you must contact OSS Service Representative to obtain a Return Merchandise Authorization (RMA) number. Put this number and your return address on the shipping label when you return the item(s) for service. **One Stop Systems will return any product that is not accompanied by an RMA number.** Please note that One Stop Systems WILL NOT accept COD packages, so be sure to return the product freight and duties paid.

Ship the well-packaged product to the address below:

RMA # _____
One Stop Systems
2235 Enterprise Street, Suite#110
Escondido, CA 92029
USA

It is not required, though highly recommended, that you keep the packaging from the original shipment of your One Stop Systems product. However, if you return a product to One Stop Systems for warranty repair/ replacement or take advantage of the 30-day money back guarantee, you will need to package the product in a manner like the way it was received from our plant. One Stop Systems cannot be responsible for any physical damage to the product or component pieces of the product (such as the host or expansion interfaces for the PCIe expansion chassis) that are damaged due to inadequate packing. Physical damage sustained in such a situation will be repaired at the owner's expense in accordance with Out of Warranty Procedures. Please, protect your investment, a bit more padding in a good box will go a long way to ensuring the device is returned to use in the same condition you shipped it in. Please call for an RMA number first.

12 Shipping / transporting the card

Use appropriate packaging materials.



IMPORTANT

PCIe cards should be removed (or not to be installed) prior to shipping to avoid or prevent possible damage, failure to do so, will void the warranty of the unit.

11 APPENDIX A Compliance

FCC

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the service personnel will be required to correct the interference at his own expense. This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received including interference that may cause undesired operation. Changes or modifications not expressly approved by the party responsible for compliance could void the service personnel's authority to operate the equipment.



NOTE

The assembler of a personal computer system may be required to test the system and/or make necessary modifications if a system is found to cause harmful interferences or to be noncompliant with the appropriate standards for its intended use.

Industry Canada

This Class A digital apparatus complies with Canadian ICES-003. Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada

CE

The product(s) described in this manual complies with all applicable European Union (CE) directives. One Stop Systems will not retest or recertify systems or components that have been reconfigured by customers



2235 Enterprise Street, Suite#110, Escondido CA 92029

Toll-Free : +1(800)285-8900 US • Main: +1 (760) 745-9883 • Fax: +1 (760) 745-9824

www.onestopsystems.com