



Multi-Host Sharing of NVMe Drives and GPUs Using PCIe Fabrics

Abstract

This white paper discusses how PCIe fabrics can be used to create a flexible, low-latency, high-performance fabric interconnect to a shared pool of GPUs and NVMe SSDs while still supporting standard host operating system drivers. Using dynamic partitioning and multi-host single root I/O virtualization (SR-IOV) sharing techniques, GPU and NVMe resources can be "composed" or dynamically allocated to a specific host or set of hosts, allowing real-time allocation of resources to match workload requirements. Key concepts of PCIe fabrics and multi-host sharing of SR-IOV devices are also discussed.

Table of Contents

Abstract.....	1
1. Introduction.....	3
2. PCIe Hierarchy Restriction.....	4
3. PCIe Single Root Hierarchy Domain Restriction.....	6
4. PCIe Fabrics for Scaling.....	7
5. PCIe Fabrics for Multi-Host Sharing.....	9
6. Demonstration.....	10
7. Conclusion.....	14
8. Revision History.....	15
8.1. Revision 3.0.....	15
8.2. Revision 2.0.....	15
8.3. Revision 1.0.....	15
The Microchip Website.....	16
Product Change Notification Service.....	16
Customer Support.....	16
Microchip Devices Code Protection Feature.....	16
Legal Notice.....	17
Trademarks.....	17
Quality Management System.....	18
Worldwide Sales and Service.....	19

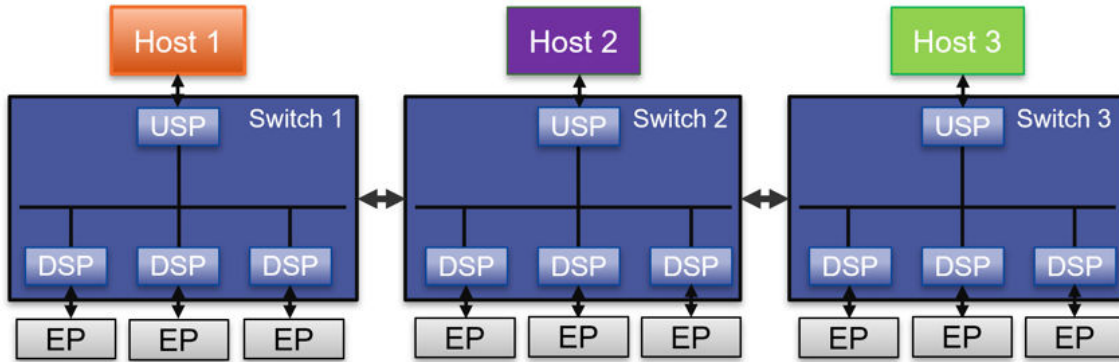
1. Introduction

As the use of GPUs for deep learning, artificial intelligence and machine learning increases, so does the demand for a more efficient deployment of GPU and NVMe resources. To effectively integrate these costly devices into systems, datacenter equipment designers require innovative technologies for efficiently sharing system resources among multiple hosts with a high-bandwidth, low-latency interconnection in disaggregated, composable architectures. PCIe is an effective, high-performance, and ubiquitous system interconnect, but there are limitations in the specification that limit its usage in such designs. This white paper discusses how PCIe fabrics are able to overcome the limitations of conventional PCIe tree-based systems to provide flexible, dynamic composition and sharing of system resources.

2. PCIe Hierarchy Restriction

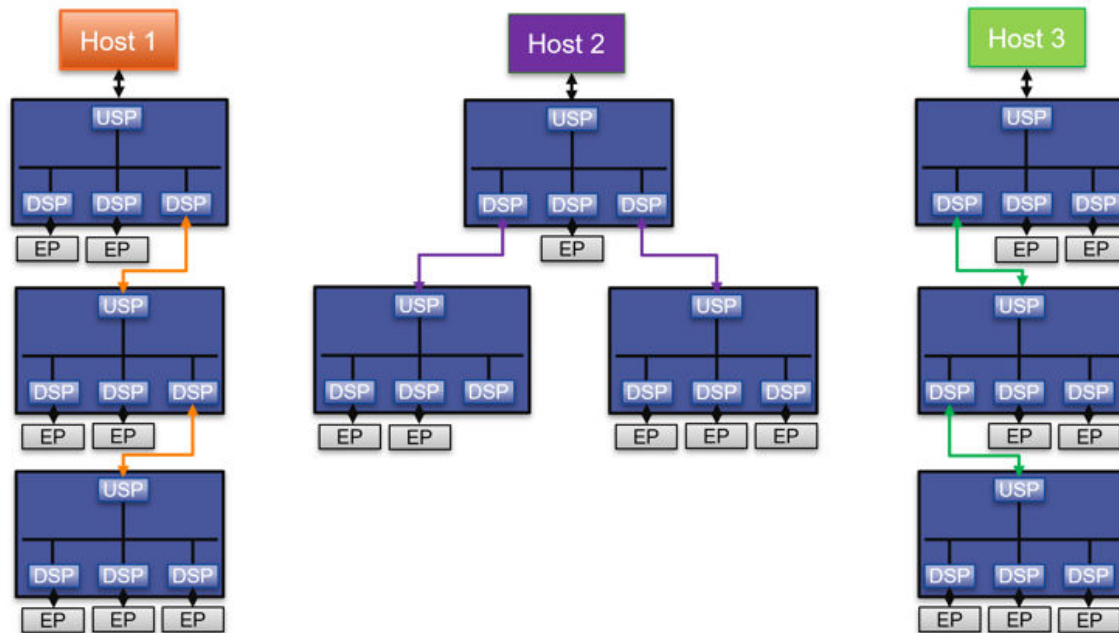
The PCIe standard defines a hierarchy domain as a rigid, hierarchical tree structure, which complicates the design of large scale systems involving multiple PCIe switches and host systems. For example, consider a system with three hosts and three switches, as shown in the following figure.

Figure 2-1. Multi-Host Topology




To maintain a PCIe-specification compliant hierarchy, Host 1 must have a dedicated downstream port (DSP) in Switch 1 connected to a dedicated upstream port (USP) in Switch 2, and a dedicated DSP in Switch 2 connected to a dedicated USP in Switch 3. Similar requirements exist for Host 2 and Host 3, as shown in the following figure.

Figure 2-2. Hierarchy Requirements for Each Host



As a result, a simple, PCIe tree structure-based system requires three links between each switch dedicated to each host's PCIe topology. The inability to share these links between the hosts complicates the design and decreases system efficiency.

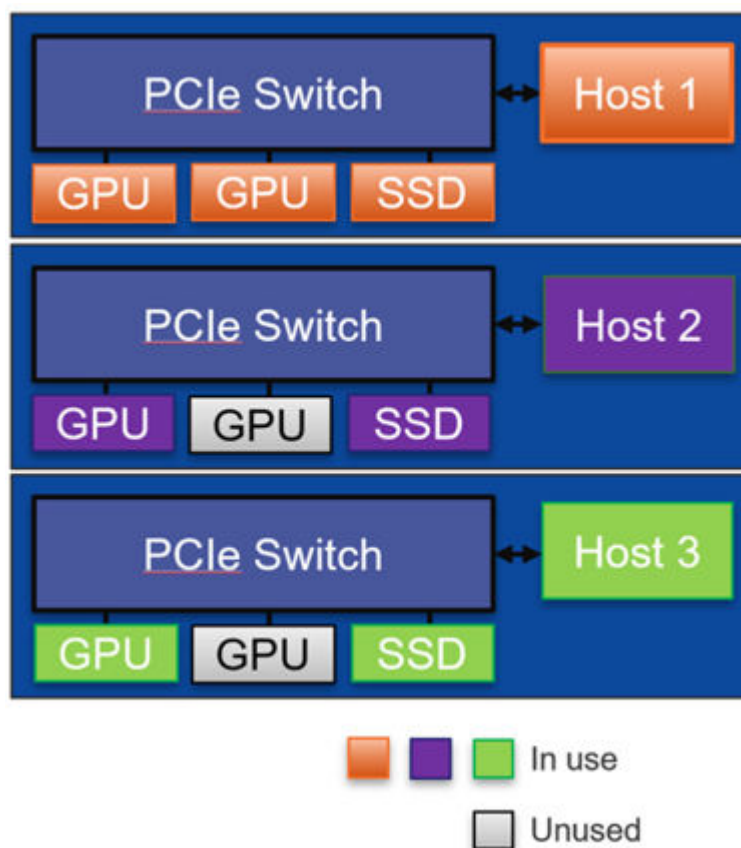
Figure 2-3. PCIe Link Required for Each Host



3. PCIe Single Root Hierarchy Domain Restriction

A typical PCIe specification-compliant hierarchy domain only contains one root port. Standards have been defined for extending a hierarchy domain to support multiple roots (Multi-Root I/O Virtualization and Sharing Specification Revision 1.0), but this is a complicated capability that has not been implemented by a major, modern host CPU. The result is that unused PCIe devices are stranded in the hierarchy domain of the host that owns them. An example of such a scenario is shown in the following figure.

Figure 3-1. Stranded, Unused PCIe EPs

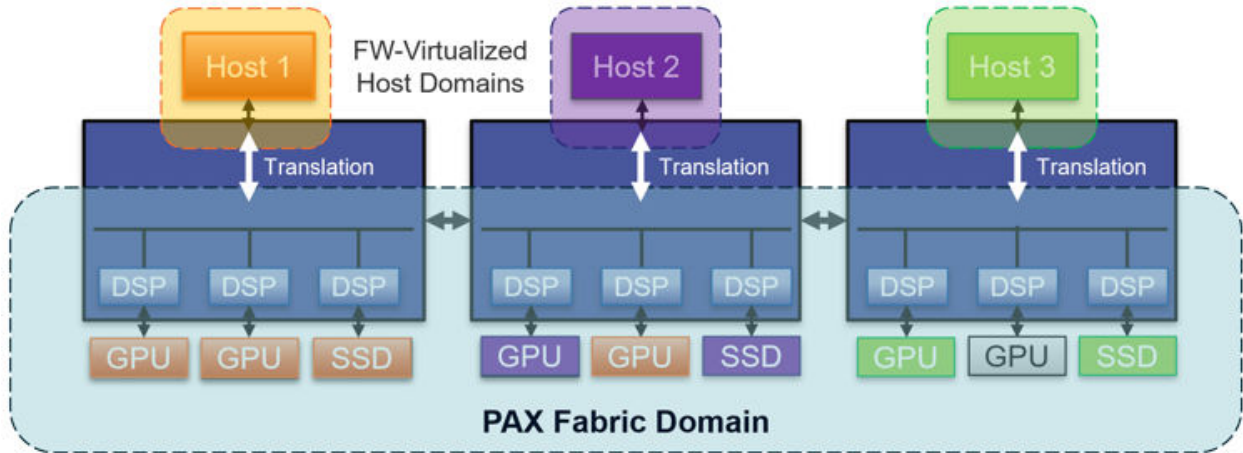


Host 1 has fully consumed its compute resources, but Hosts 2 and 3 are underutilizing theirs. It is impossible for Host 1 to access the unused compute resources because they are outside of its hierarchy domain. Device sharing can be implemented using non-transparent bridging (NTB), but that requires the implementation of complicated, non-standard drivers and software to be developed for each type of shared PCIe device.

4. PCIe Fabrics for Scaling


A PCIe fabric can be used to address issues with scaling a standard PCIe topology to multiple hosts and multiple switches. A PCIe fabric-based system implemented using Microchip's Switchtec™ technology PAX Advanced Fabric PCIe Switch is separated into two types of domains: a fabric domain containing all EPs and fabric links, and host domains for each connected host. Transactions from the host domains are translated to IDs and addresses in the fabric domain, and vice versa. Proprietary, non-hierarchical routing is used for traffic in the fabric domain. This allows the fabric links connecting the switches to be efficiently shared by all hosts in the system.

Figure 4-1. Types of Domains




Firmware running on an embedded processor in the fabric PCIe switch intercepts all configuration plane traffic from the host, including the PCIe enumeration process. It virtualizes a simple, PCIe specification-compliant switch with a configurable number of DSPs.

Figure 4-2. PCIe Switch Virtualization



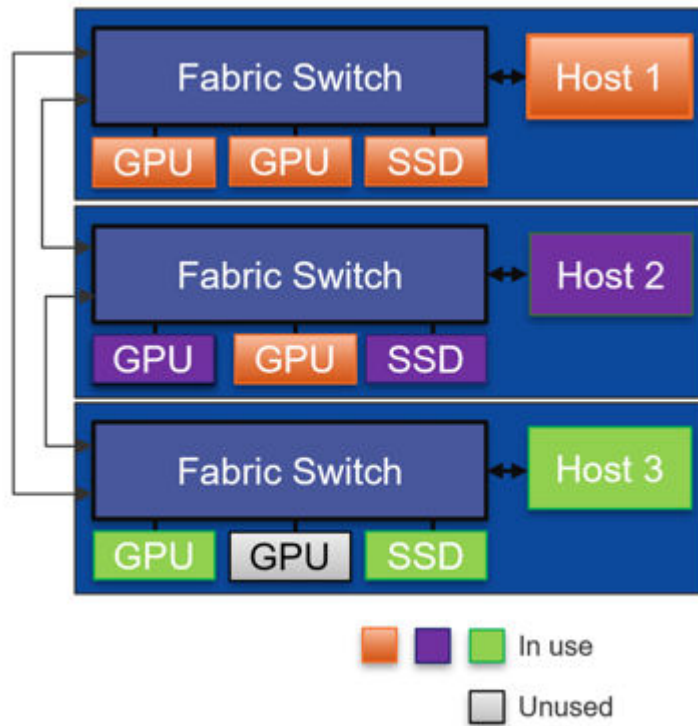
PCIe devices assigned to the host domain appear directly connected to the virtual PCIe switch. All control plane traffic will be routed to the switch firmware for processing, but data plane traffic will be routed directly to the PCIe EPs to ensure maximum system performance.

Figure 4-3. Traffic Routing in a PAX PCIe Fabric



Unused devices in other host domains are no longer stranded and can be dynamically assigned based on each host's resource requirements. Peer-to-peer traffic is supported within the fabric to enable modern AI/ML applications. These system capabilities are presented to hosts in a PCIe specification-compliant manner, so standard drivers can be used.

Figure 4-4. PCIe Fabric for EP Pooling




5. PCIe Fabrics for Multi-Host Sharing

Within the PCIe fabric, devices are assigned to hosts at a PCIe function-level granularity. As a result, multi-function devices can have individual functions assigned to different hosts, allowing for multi-host sharing. This capability can be used for devices with multiple functions and those that implement SR-IOV, a PCIe capability that allows a device to present multiple physical and Virtual Functions (VFs). Switch firmware can also be used to virtualize the configuration space of the VFs to alter their contents and capabilities. For example, the VFs of an SR-IOV capable device can be modified to present the VF as a standard, single function NVM device. This allows devices to be shared among hosts using standard, in-box drivers.

Not all multi-function and SR-IOV capable devices are designed to support this model for multi-host sharing. Each PCIe function must be designed to support standalone operation with minimal requirements for coordinating with other functions. PAX evaluation platforms are available to test EP compatibility with multi-host sharing across a fabric.

Figure 5-1. Multi-Host Sharing of SR-IOV VFs

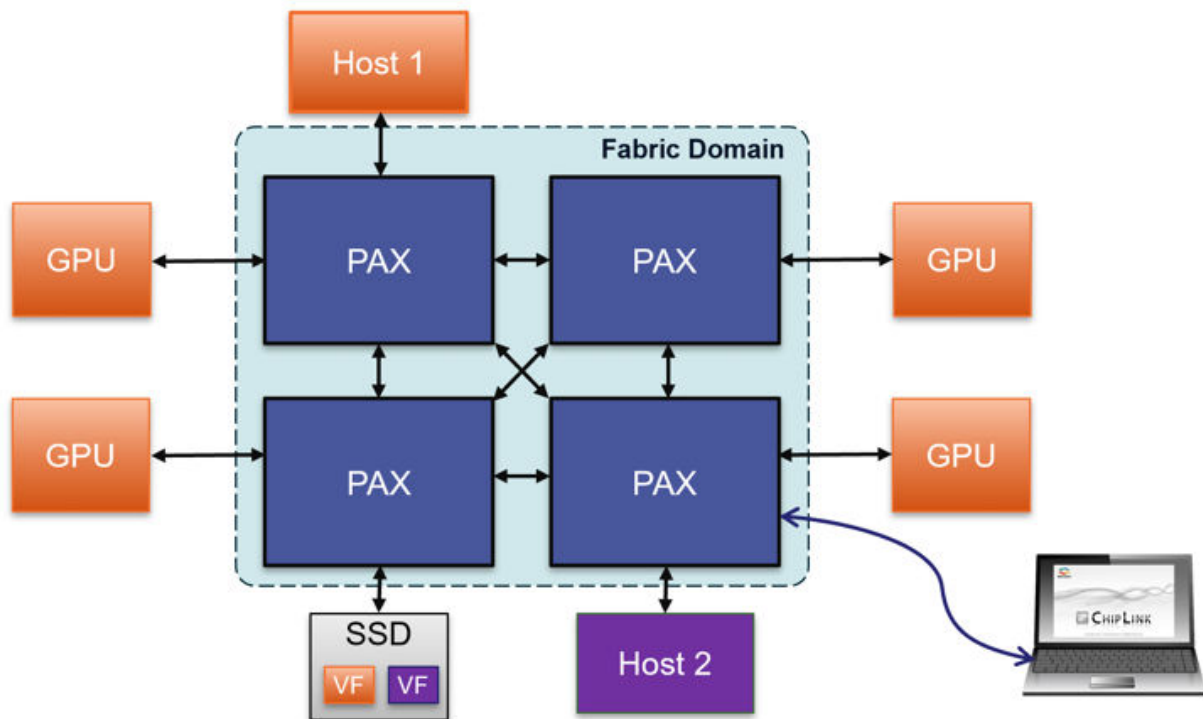


6. Demonstration

Microchip has successfully proven the concepts presented in this paper with a real-world example. We have implemented a proof-of-concept system that demonstrates dynamic assignment of GPUs and multi-host sharing of SR-IOV SSDs. The hosts tested are running Windows Server 2016 and Ubuntu Server 16.04 LTS, neither of which require any custom drivers or software to facilitate the multi-host sharing. The hosts run traffic representative of actual AI/ML workloads, including Nvidia’s CUDA peer-to-peer traffic benchmarking utility, p2pBandwidthLatencyTest, and training the cifar10 image classification Tensorflow model.

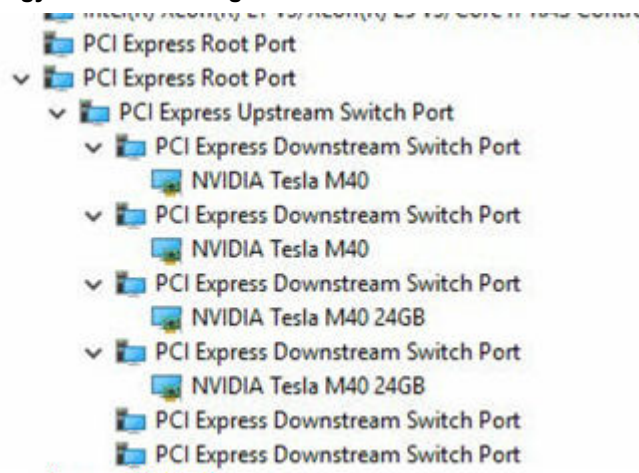
The demo system comprises of four PAX fabric PCIe switches, four Nvidia Tesla GPGPUs, one Samsung PM1725a NVM device with SR-IOV support, and two Supermicro servers interconnected as illustrated below. The embedded switch firmware handles the low-level configuration and management of the switch hardware, so the demo system is managed from Microchip’s debug and diagnostics utility, ChipLink, connected over a simple UART management interface.

Figure 6-1. Demo System Block Diagram



Initially, all GPUs are assigned to Host 1, which is running Windows Server 2016, to increase the performance of the AI training. The virtual switch presented to the host can be seen in the Windows Device Manager tool. All GPUs appear as though they are directly connected to the virtual switch; the fabric links and complexities of the physical topology are obscured from the host.

Figure 6-2. Virtual Switch Topology in Device Manager



The p2pBandwidthLatencyTest utility is used to measure the bandwidth of GPU-to-GPU transfers across the PCIe fabric.

Figure 6-3. p2pBandwidthLatencyTest Performance Excerpt from Host 1

```
P2P Connectivity Matrix
D\D      0      1      2      3
0         1      1      1      1
1         1      1      1      1
2         1      1      1      1
3         1      1      1      1
```

```
Unidirectional P2P=Enabled Bandwidth Matrix (GB/s)
D\D      0      1      2      3
0  210.61  12.96  12.54  12.53
1   12.52 211.35  13.08  13.06
2   12.52  12.52 212.61  13.05
3   13.06  13.06  12.54 211.36
```

```
Bidirectional P2P=Enabled Bandwidth Matrix (GB/s)
D\D      0      1      2      3
0  213.51  24.81  24.77  24.72
1   24.73 213.55  24.73  25.74
2   24.53  24.57 214.73  24.86
3   24.83  25.72  24.73 214.58
```

With the GPUs discovered and the standard Nvidia drivers loaded, an AI workload is started. In this example, the user is training the cifar10 image classification algorithm. The training algorithm will be distributed across all four GPUs.

Figure 6-4. cifar10 Multi-GPU Training Output Excerpt

```
Administrator: Command Prompt - python cifar10_multi_gpu_train.py
2018-04-06 18:37:52.698874: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1030] Found device 3 with properties:
name: Tesla M40 24GB major: 5 minor: 2 memoryClockRate(GHz): 1.112
pciBusID: 0000:07:00.0
totalMemory: 22.43GiB freeMemory: 22.18GiB
2018-04-06 18:37:52.701446: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1045] Device peer to peer matrix
2018-04-06 18:37:52.702693: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1051] DMA: 0 1 2 3
2018-04-06 18:37:52.703145: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1061] 0:  Y Y Y Y
2018-04-06 18:37:52.704059: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1061] 1:  Y Y Y Y
2018-04-06 18:37:52.704960: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1061] 2:  Y Y Y Y
2018-04-06 18:37:52.705856: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1061] 3:  Y Y Y Y
2018-04-06 18:37:52.706817: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1120] Creating TensorFlow device (/device:GPU:0) -> (device: 0, name: Tesla M40, pci bus id: 0000:04:00.
0, compute capability: 5.2)
2018-04-06 18:37:52.707744: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1120] Creating TensorFlow device (/device:GPU:1) -> (device: 1, name: Tesla M40, pci bus id: 0000:05:00.
0, compute capability: 5.2)
2018-04-06 18:37:52.708655: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1120] Creating TensorFlow device (/device:GPU:2) -> (device: 2, name: Tesla M40 24GB, pci bus id: 0000:0
6:00.0, compute capability: 5.2)
2018-04-06 18:37:52.709626: I C:\tf_jenkins\home\workspace\rel-win\M\windows-gpu\PY\36\tensorflow\core\common_runtime\gp
\gpu_device.cc:1120] Creating TensorFlow device (/device:GPU:3) -> (device: 3, name: Tesla M40 24GB, pci bus id: 0000:0
7:00.0, compute capability: 5.2)
```

When the workload completes, the user can release two of the GPUs back into the fabric pool. Now, we'll assign one of the SR-IOV drive's VFs to each host. On Host 1, the drive appears as a "Standard NVM Express Controller" using the standard, in-box Windows driver for NVM devices.

Figure 6-5. SR-IOV NVM VF Virtualized as Standard NVM Controller


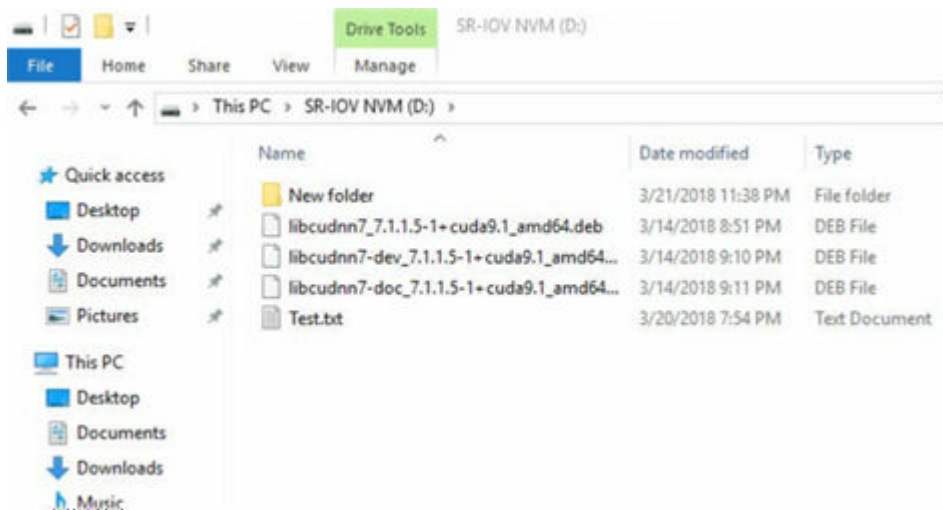


Figure 6-6. SR-IOV NVM VF Partition



Host 2 is running Ubuntu Server 16.04 LTS, and its PCIe topology is also a simple switch with locally attached GPUs and a standard NVM device.

Figure 6-7. Virtual Switch Topology in Linux

```
00.0-[03-07]--+00.0-[04]----00.0 NVIDIA Corporation GM200GL [Tesla M40]
      +-01.0-[05]----00.0 NVIDIA Corporation GM200GL [Tesla M40]
      +-02.0-[06]----00.0 Samsung Electronics Co Ltd Device a822
      \-03.0-[07]--
```

The p2pBandwidthLatencyTest utility produces similar performance results on Host 2.

Figure 6-8. p2pBandwidthLatencyTest Performance Excerpt from Host 2

```
P2P Connectivity Matrix
D\D      0      1
0         1      1
1         1      1

Unidirectional P2P=Enabled Bandwidth Matrix (GB/s)
D\D      0      1
0 214.21 12.27
1 13.06 212.99

Bidirectional P2P=Enabled Bandwidth Matrix (GB/s)
D\D      0      1
0 214.53 24.33
1 24.83 215.55
```

The NVM VF is presented to Host 2 as a standard NVM device, so its standard, in-box drivers can be used.

Figure 6-9. SR-IOV NVM VF in Device Listing and Mounted Volume Contents

```
ubuntu@bbjapps-ubuntu1604-se:~$ ls /dev
autofs          hidrau1        loop3          nvme0n1p2     sg0            tty2           tty4           tty6           ttyS20        userio
bbjapps-ubuntu1604-se-sg hidrau2        loop4          port          sg1            tty20          tty40          tty60          ttyS21        vcs
block          hidrau3        loop5          ppp           sg2            tty21          tty41          tty61          ttyS22        vcs1
log            hpet          loop6          psaux         snapshot      tty22          tty42          tty62          ttyS23        vcs2
brfs-control   klogd         loop7          ptmx          snapshot      tty23          tty43          tty63          ttyS24        vcs3
bus            hwrng         loop-control  ptp0          sad           tty24          tty44          tty7           ttyS25        vcs4
char           i2c-0         nmap          ptp1          stderr        tty25          tty45          tty8           ttyS26        vcs5
console        i2c-1         ncclog        pl            stdin         tty26          tty46          tty9           ttyS27        vcs6
core           i2c-2         nca           random        stdout        tty27          tty47          ttypprintk    ttyS28        vcsa
cpm            i2c-3         memory_bandwidth rfcill       tty           tty28          tty48          tty50          ttyS29        vcsa1
cpu_dma_latency i2c-4         nqsusd        rtc           tty0          tty29          tty49          tty51          ttyS3         vcsa2
cuse          i2c-5         sel           rtc0          tty1          tty3           tty5           ttyS10         ttyS30        vcsa3
disk          i2c-6         network_latency sda           tty10         tty30          tty50          ttyS11         ttyS31        vcsa4
dm-0          initctl       network_throughput sda1         tty11         tty31          tty51          ttyS12         ttyS4         vcsa5
dm-1          lvm           null          sda2         tty12         tty32          tty52          ttyS13         ttyS5         vcsa6
dri           kmsg         nvidia0       sda5         tty13         tty33          tty53          ttyS14         ttyS6         vfo
cryptfs       kvm           nvidia1       sdb           tty14         tty34          tty54          ttyS15         ttyS7         wpa_arbiter
fb0           lightdm       nvidiaactl    sdb1         tty15         tty35          tty55          ttyS16         ttyS8         vhci
fd            log           nvidia-uvm    sdb2         tty16         tty36          tty56          ttyS17         ttyS9         vhost-net
full          loop0         nvme0         sdb3         tty17         tty37          tty57          ttyS18         uhid         zero
fuse          loop1         nvme0n1       sdc           tty18         tty38          tty58          ttyS19         uinput
hidrau0       loop2         nvme0n1p1     sdc1         tty19         tty39          tty59          ttyS2          urandom

ubuntu@bbjapps-ubuntu1604-se:~$ sudo mount /dev/nvme0n1p2 /mnt
ubuntu@bbjapps-ubuntu1604-se:~$ ls /mnt
libcuda7_7.1.1.5-1+cuda9.1_and64.deb  libcuda7-doc_7.1.1.5-1+cuda9.1_and64.deb  $RECYCLE_BIN  Test.txt
libcuda7-dev_7.1.1.5-1+cuda9.1_and64.deb  New folder  System Volume Information
```

As demonstrated above, the virtual PCIe switch and all dynamic assignment operations are presented to the host as fully PCIe specification-compliant, enabling the hosts to use standard, in-box drivers. The embedded switch firmware provides a simple management interface so the PCIe fabric can be configured and managed by a light-weight external processor. Device peer-to-peer transactions are enabled by default and require no additional configuration or management from an external fabric manager.

7. Conclusion

PCIe fabrics provide the high-bandwidth, low-latency interconnections required to implement the next generation of disaggregated, composable architectures. System designers can provision or share system resources in scalable, efficient ways using standard in-box drivers.

Microchip's Switchtec technology PAX Advanced Fabric PCIe switches enable new architectures for next-generation solutions. They provide a scalable, low-latency, high-bandwidth, and cost-effective option for system design. With its virtualization of switch topologies and multi-host sharing of PCIe devices using standard, in-box drivers, PAX significantly reduces time to market and storage costs and simplifies system development. For more information, visit [Switchtec PAX Advanced Fabric PCIe Switches](#) or email: ClientEngagementStaff@microchip.com.

8. Revision History

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the most current publication.

8.1 Revision 3.0

Revision 3.0 was published in October 2020. The template was updated to the Microchip style, and it was assigned literature number DS00003702A.

8.2 Revision 2.0

Revision 2.0 was published in October 2019. It is the first publication of this document.

8.3 Revision 1.0

Revision 1.0 was an internal publication.

The Microchip Website

Microchip provides online support via our website at www.microchip.com/. This website is used to make files and information easily available to customers. Some of the content available includes:

- **Product Support** – Data sheets and errata, application notes and sample programs, design resources, user's guides and hardware support documents, latest software releases and archived software
- **General Technical Support** – Frequently Asked Questions (FAQs), technical support requests, online discussion groups, Microchip design partner program member listing
- **Business of Microchip** – Product selector and ordering guides, latest Microchip press releases, listing of seminars and events, listings of Microchip sales offices, distributors and factory representatives

Product Change Notification Service

Microchip's product change notification service helps keep customers current on Microchip products. Subscribers will receive email notification whenever there are changes, updates, revisions or errata related to a specified product family or development tool of interest.

To register, go to www.microchip.com/pcn and follow the registration instructions.

Customer Support

Users of Microchip products can receive assistance through several channels:

- Distributor or Representative
- Local Sales Office
- Embedded Solutions Engineer (ESE)
- Technical Support

Customers should contact their distributor, representative or ESE for support. Local sales offices are also available to help customers. A listing of sales offices and locations is included in this document.

Technical support is available through the website at: www.microchip.com/support

Microchip Devices Code Protection Feature

Note the following details of the code protection feature on Microchip devices:

- Microchip products meet the specifications contained in their particular Microchip Data Sheet.
- Microchip believes that its family of products is secure when used in the intended manner and under normal conditions.
- There are dishonest and possibly illegal methods being used in attempts to breach the code protection features of the Microchip devices. We believe that these methods require using the Microchip products in a manner outside the operating specifications contained in Microchip's Data Sheets. Attempts to breach these code protection features, most likely, cannot be accomplished without violating Microchip's intellectual property rights.
- Microchip is willing to work with any customer who is concerned about the integrity of its code.
- Neither Microchip nor any other semiconductor manufacturer can guarantee the security of its code. Code protection does not mean that we are guaranteeing the product is "unbreakable." Code protection is constantly evolving. We at Microchip are committed to continuously improving the code protection features of our products. Attempts to break Microchip's code protection feature may be a violation of the Digital Millennium Copyright Act. If such acts allow unauthorized access to your software or other copyrighted work, you may have a right to sue for relief under that Act.

Legal Notice

Information contained in this publication is provided for the sole purpose of designing with and using Microchip products. Information regarding device applications and the like is provided only for your convenience and may be superseded by updates. It is your responsibility to ensure that your application meets with your specifications.

THIS INFORMATION IS PROVIDED BY MICROCHIP "AS IS". MICROCHIP MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WHETHER EXPRESS OR IMPLIED, WRITTEN OR ORAL, STATUTORY OR OTHERWISE, RELATED TO THE INFORMATION INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE OR WARRANTIES RELATED TO ITS CONDITION, QUALITY, OR PERFORMANCE.

IN NO EVENT WILL MICROCHIP BE LIABLE FOR ANY INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL OR CONSEQUENTIAL LOSS, DAMAGE, COST OR EXPENSE OF ANY KIND WHATSOEVER RELATED TO THE INFORMATION OR ITS USE, HOWEVER CAUSED, EVEN IF MICROCHIP HAS BEEN ADVISED OF THE POSSIBILITY OR THE DAMAGES ARE FORESEEABLE. TO THE FULLEST EXTENT ALLOWED BY LAW, MICROCHIP'S TOTAL LIABILITY ON ALL CLAIMS IN ANY WAY RELATED TO THE INFORMATION OR ITS USE WILL NOT EXCEED THE AMOUNT OF FEES, IF ANY, THAT YOU HAVE PAID DIRECTLY TO MICROCHIP FOR THE INFORMATION. Use of Microchip devices in life support and/or safety applications is entirely at the buyer's risk, and the buyer agrees to defend, indemnify and hold harmless Microchip from any and all damages, claims, suits, or expenses resulting from such use. No licenses are conveyed, implicitly or otherwise, under any Microchip intellectual property rights unless otherwise stated.

Trademarks

The Microchip name and logo, the Microchip logo, Adaptec, AnyRate, AVR, AVR logo, AVR Freaks, BesTime, BitCloud, chipKIT, chipKIT logo, CryptoMemory, CryptoRF, dsPIC, FlashFlex, flexPWR, HELDO, IGLOO, JukeBlox, KeeLoq, Klear, LANCheck, LinkMD, maXStylus, maXTouch, MediaLB, megaAVR, Microsemi, Microsemi logo, MOST, MOST logo, MPLAB, OptoLyzer, PackeTime, PIC, picoPower, PICSTART, PIC32 logo, PolarFire, Prochip Designer, QTouch, SAM-BA, SenGenuity, SpyNIC, SST, SST Logo, SuperFlash, Symmetricom, SyncServer, Tachyon, TimeSource, tinyAVR, UNI/O, Vectron, and XMEGA are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

AgileSwitch, APT, ClockWorks, The Embedded Control Solutions Company, EtherSynch, FlashTec, Hyper Speed Control, HyperLight Load, IntelliMOS, Libero, motorBench, mTouch, Powermite 3, Precision Edge, ProASIC, ProASIC Plus, ProASIC Plus logo, Quiet-Wire, SmartFusion, SyncWorld, Temux, TimeCesium, TimeHub, TimePictra, TimeProvider, WinPath, and ZL are registered trademarks of Microchip Technology Incorporated in the U.S.A.

Adjacent Key Suppression, AKS, Analog-for-the-Digital Age, Any Capacitor, AnyIn, AnyOut, Augmented Switching, BlueSky, BodyCom, CodeGuard, CryptoAuthentication, CryptoAutomotive, CryptoCompanion, CryptoController, dsPICDEM, dsPICDEM.net, Dynamic Average Matching, DAM, ECAN, Espresso T1S, EtherGREEN, IdealBridge, In-Circuit Serial Programming, ICSP, INICnet, Intelligent Paralleling, Inter-Chip Connectivity, JitterBlocker, maxCrypto, maxView, memBrain, Mindi, MiWi, MPASM, MPF, MPLAB Certified logo, MPLIB, MPLINK, MultiTRAK, NetDetach, Omniscient Code Generation, PICDEM, PICDEM.net, PICKit, PICtail, PowerSmart, PureSilicon, QMatrix, REAL ICE, Ripple Blocker, RTAX, RTG4, SAM-ICE, Serial Quad I/O, simpleMAP, SimpliPHY, SmartBuffer, SMART-I.S., storClad, SQL, SuperSwitcher, SuperSwitcher II, Switchtec, SynchroPHY, Total Endurance, TSHARC, USBCheck, VariSense, VectorBlox, VeriPHY, ViewSpan, WiperLock, XpressConnect, and ZENA are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries.

SQTP is a service mark of Microchip Technology Incorporated in the U.S.A.

The Adaptec logo, Frequency on Demand, Silicon Storage Technology, and Symmcom are registered trademarks of Microchip Technology Inc. in other countries.

GestIC is a registered trademark of Microchip Technology Germany II GmbH & Co. KG, a subsidiary of Microchip Technology Inc., in other countries.

All other trademarks mentioned herein are property of their respective companies.

© 2020, Microchip Technology Incorporated, Printed in the U.S.A., All Rights Reserved.

ISBN: 978-1-5224-6960-5

Quality Management System

For information regarding Microchip's Quality Management Systems, please visit www.microchip.com/quality.

Worldwide Sales and Service

AMERICAS	ASIA/PACIFIC	ASIA/PACIFIC	EUROPE
<p>Corporate Office 2355 West Chandler Blvd. Chandler, AZ 85224-6199 Tel: 480-792-7200 Tel: 480-792-7277 Technical Support: www.microchip.com/support Web Address: www.microchip.com</p> <p>Atlanta Duluth, GA Tel: 678-957-9614 Fax: 678-957-1455</p> <p>Austin, TX Tel: 512-257-3370</p> <p>Boston Westborough, MA Tel: 774-760-0087 Fax: 774-760-0088</p> <p>Chicago Itasca, IL Tel: 630-285-0071 Fax: 630-285-0075</p> <p>Dallas Addison, TX Tel: 972-818-7423 Fax: 972-818-2924</p> <p>Detroit Novi, MI Tel: 248-848-4000</p> <p>Houston, TX Tel: 281-894-5983</p> <p>Indianapolis Noblesville, IN Tel: 317-773-8323 Fax: 317-773-5453 Tel: 317-536-2380</p> <p>Los Angeles Mission Viejo, CA Tel: 949-462-9523 Fax: 949-462-9608 Tel: 951-273-7800</p> <p>Raleigh, NC Tel: 919-844-7510</p> <p>New York, NY Tel: 631-435-6000</p> <p>San Jose, CA Tel: 408-735-9110 Tel: 408-436-4270</p> <p>Canada - Toronto Tel: 905-695-1980 Fax: 905-695-2078</p>	<p>Australia - Sydney Tel: 61-2-9868-6733</p> <p>China - Beijing Tel: 86-10-8569-7000</p> <p>China - Chengdu Tel: 86-28-8665-5511</p> <p>China - Chongqing Tel: 86-23-8980-9588</p> <p>China - Dongguan Tel: 86-769-8702-9880</p> <p>China - Guangzhou Tel: 86-20-8755-8029</p> <p>China - Hangzhou Tel: 86-571-8792-8115</p> <p>China - Hong Kong SAR Tel: 852-2943-5100</p> <p>China - Nanjing Tel: 86-25-8473-2460</p> <p>China - Qingdao Tel: 86-532-8502-7355</p> <p>China - Shanghai Tel: 86-21-3326-8000</p> <p>China - Shenyang Tel: 86-24-2334-2829</p> <p>China - Shenzhen Tel: 86-755-8864-2200</p> <p>China - Suzhou Tel: 86-186-6233-1526</p> <p>China - Wuhan Tel: 86-27-5980-5300</p> <p>China - Xian Tel: 86-29-8833-7252</p> <p>China - Xiamen Tel: 86-592-2388138</p> <p>China - Zhuhai Tel: 86-756-3210040</p>	<p>India - Bangalore Tel: 91-80-3090-4444</p> <p>India - New Delhi Tel: 91-11-4160-8631</p> <p>India - Pune Tel: 91-20-4121-0141</p> <p>Japan - Osaka Tel: 81-6-6152-7160</p> <p>Japan - Tokyo Tel: 81-3-6880-3770</p> <p>Korea - Daegu Tel: 82-53-744-4301</p> <p>Korea - Seoul Tel: 82-2-554-7200</p> <p>Malaysia - Kuala Lumpur Tel: 60-3-7651-7906</p> <p>Malaysia - Penang Tel: 60-4-227-8870</p> <p>Philippines - Manila Tel: 63-2-634-9065</p> <p>Singapore Tel: 65-6334-8870</p> <p>Taiwan - Hsin Chu Tel: 886-3-577-8366</p> <p>Taiwan - Kaohsiung Tel: 886-7-213-7830</p> <p>Taiwan - Taipei Tel: 886-2-2508-8600</p> <p>Thailand - Bangkok Tel: 66-2-694-1351</p> <p>Vietnam - Ho Chi Minh Tel: 84-28-5448-2100</p>	<p>Austria - Wels Tel: 43-7242-2244-39 Fax: 43-7242-2244-393</p> <p>Denmark - Copenhagen Tel: 45-4485-5910 Fax: 45-4485-2829</p> <p>Finland - Espoo Tel: 358-9-4520-820</p> <p>France - Paris Tel: 33-1-69-53-63-20 Fax: 33-1-69-30-90-79</p> <p>Germany - Garching Tel: 49-8931-9700</p> <p>Germany - Haan Tel: 49-2129-3766400</p> <p>Germany - Heilbronn Tel: 49-7131-72400</p> <p>Germany - Karlsruhe Tel: 49-721-625370</p> <p>Germany - Munich Tel: 49-89-627-144-0 Fax: 49-89-627-144-44</p> <p>Germany - Rosenheim Tel: 49-8031-354-560</p> <p>Israel - Ra'anana Tel: 972-9-744-7705</p> <p>Italy - Milan Tel: 39-0331-742611 Fax: 39-0331-466781</p> <p>Italy - Padova Tel: 39-049-7625286</p> <p>Netherlands - Drunen Tel: 31-416-690399 Fax: 31-416-690340</p> <p>Norway - Trondheim Tel: 47-72884388</p> <p>Poland - Warsaw Tel: 48-22-3325737</p> <p>Romania - Bucharest Tel: 40-21-407-87-50</p> <p>Spain - Madrid Tel: 34-91-708-08-90 Fax: 34-91-708-08-91</p> <p>Sweden - Gothenberg Tel: 46-31-704-60-40</p> <p>Sweden - Stockholm Tel: 46-8-5090-4654</p> <p>UK - Wokingham Tel: 44-118-921-5800 Fax: 44-118-921-5820</p>