NeSC: Self-Virtualizing Nested Storage Controller

Yonatan Gottesman Electrical Engineering Technion — Israel Institute of Technology yoni@tce.technion.ac.il Yoav Etsion Electrical Engineering and Computer Science Technion — Israel Institute of Technology yetsion@tce.technion.ac.il

Abstract—The emergence of high-speed, multi GB/s storage devices has shifted the performance bottleneck of storage virtualization to the software layers of the hypervisor. The hypervisor overheads can be avoided by allowing the virtual machine (VM) to directly access the storage device (a method known as direct device assignment), but this method voids all protection guarantees provided by filesystem permissions, since the device has no notion of client isolation. Recently, following the introduction of 10Gbs and higher networking interfaces, the PCIe specification was extended to include the SR-IOV specification for self-virtualizing devices, which allows a single physical device to present multiple virtual interfaces on the PCIe interconnect. Using SR-IOV, a hypervisor can directly assign a virtual PCIe device interface to each of its VMs. However, as networking interfaces simply multiplex packets sent from/to different clients, the specification does not dictate the semantics of a virtual storage device and how to maintain data isolation in a self-virtualizing device.

In this paper we present the self-virtualizing, nested storage controller (NeSC) architecture, which includes a filesystemagnostic protection mechanism that enables the physical device to export files as virtual PCIe storage devices. The protection mechanism maps file offsets to physical blocks and thereby offloads the hypervisor's storage layer functionality to hardware. Using NeSC, a hypervisor can securely expose its files as virtual PCIe devices and directly assign them to VMs. We have prototyped a 1GB/s NeSC controller using a Virtex-7 FPGA development board connected to the PCIe interconnect. Our evaluation of NeSC on a real system shows that NeSC virtual devices enable VMs to access their data with near-native performance (in terms of both throughput and latency).

I. INTRODUCTION

The prevalence of machine consolidation through the use of virtual machines (VMs) necessitates improvements in VM performance. On the architecture side, major processor vendors have introduced virtualization extensions to their instruction set architectures [1], [2], [3]. The emergence of high-speed networking interfaces also requires that VMs be granted direct access to physical devices, thereby eliminating the costly, software-based hypervisor device multiplexing.

Enabling untrusted VMs to directly access physical devices, however, compromises system security. To overcome the fundamental security issue, the PCIe specification was extended to support self-virtualizing devices through the Single-Root I/O Virtualization (SR-IOV) interface [4]). This method enables a physical device (*physical function* in SR-IOV parlance) to dynamically create virtual instances (*virtual functions*). Each virtual instance receives a separate address on the PCIe interconnect and can, therefore, be exclusively assigned to, and accessed by, a specific VM. This method thereby distinguishes between the physical device, managed by the hypervisor, and its virtual instances used by the VMs. Importantly, it is up to the physical device to interleave and execute requests issued to the virtual devices.

Self-virtualizing devices thus delegate the task of multiplexing VM requests from the software hypervisor to the device itself. The multiplexing policy, however, depends on the inherent semantics of the underlying device and must, naturally, isolate request streams sent by individual virtual devices (that represent client VMs). For some devices, such as networking interfaces, the physical device can simply interleave the packets sent by its virtual instances (while protecting the shared link state [5]). However, enforcing isolation is nontrivial when dealing with storage controllers/devices, which typically store a filesystem structure maintained by the hypervisor. The physical storage controller must therefore enforce the access permissions imposed by the filesystem it hosts. The introduction of next-generation, commercial PCIe SSDs that deliver multi-GB/s bandwidth [6], [7]) emphasizes the need for self-virtualizing storage technology.

In this paper we present NeSC, a self-virtualizing nested storage controller that enables hypervisors to expose files and persistent objects¹(or sets thereof) as virtual block devices that can be directly assigned to VMs.

NeSC implements the SR-IOV specification, which allows it to expose itself through multiple dynamically allocated PCIe addresses. This enables NeSC to present itself as multiple virtual storage devices (through the distinct PCIe addresses) that are directly assigned to VMs. By multiplexing the requests sent by the VMs to each virtual device, NeSC can enforce the access permissions set by the hypervisor and prevent VMs from accessing stored data for which they have no access permissions.

NeSC enforces isolation by associating each virtual NeSC device with a table that maps offsets in the virtual device to blocks on the physical device. This process follows the way filesystems map file offsets to disk blocks. VMs view virtual NeSC instances as regular PCIe storage controllers (block devices). Whenever the hypervisor wishes to grant a VM direct access to a file, it queries the filesystem for the file's logical-to-physical mapping and instantiates a virtual NeSC instance

¹We use the terms files and objects interchangeably in this paper.



Fig. 1. IO Virtualization techniques.

associated with the resulting mapping table. Each access by a VM to its virtual NeSC instance is then transparently mapped by NeSC to a physical disk block using the mapping table associated with the virtual device (e.g., the first block on the virtual device maps to offset zero in the mapped file).

We evaluate the performance benefits of NeSC using a real working prototype implemented on a Xilinx VC707 FPGA development board. Our evaluation shows that our NeSC prototype, which provides 800MB/s read bandwidth and almost 1GB/s write bandwidth, delivers $2.5 \times$ and $3 \times$ better read and write bandwidth, respectively, compared to a paravirtualized *virtio* [8] storage controller (the de facto standard for virtualizing storage in Linux hypervisors), and $4 \times$ and $6 \times$ better read and write bandwidth, respectively, compared to an emulated storage controller. We further show that these performance benefits are limited only by the bandwidth provided by our academic protoype. We expect that NeSC will greatly benefit commercial PCIe SSDs capable of delivering multi-GB/s of bandwidth.

Although this paper focuses on the performance benefits that NeSC provides for VMs, it is important to note that NeSC also provides secure and direct storage access for accelerators connected to the PCIe interconnet (e.g., GPGPUs, FPGAs). As virtual NeSC instances are directly accessible on the PCIe interconnect, they can be accessed directly by other PCIe devices (using direct device DMA), thereby removing the CPU from the accelerator-storage communication path.

In this paper we make the following contributions:

- We introduce NeSC, a self-virtualizing, nested storage controller that offloads file mapping and permission check operations to hardware and provides VMs with direct and secure storage access without hypervisor mediation.
- We propose a hardware implementation for NeSC, which is prototyped using a Xilinx VC707 (Virtex-7) FPGA development board.
- We evaluate the benefit of self-virtualizing storage using microbenchmarks and complete applications. Our evaluation shows that the NeSC prototype delivers 2.5× and 3× better read and write bandwidth, respectively, than

state-of-the-art software virtual storage.

The rest of this paper is organized as follows: Section II discusses the performance overheads of storage virtualization, and Section III discusses related work. We present the NeSC system design in Section IV and its architecture in Section V. We then present our evaluation methodology in Section VI, the evaluation in Section VII, and conclude in Section VIII.

II. ON THE PERFORMANCE OVERHEADS OF STORAGE VIRTUALIZATION

Hypervisors virtualize local storage resources by mapping guest storage devices onto files in their local filesystem, in a method commonly referred to as a nested filesystem [9]. As a result, they replicate the guest operating system's (OS) software layers that abstract and secure the storage devices. Notably, these software layers have been shown to present a performance bottleneck even when not replicated [10], due to the rapid increase in storage device bandwidth [6], [7]. Moreover, further performance degradation is caused by the method by which hypervisors virtualize storage devices and the resulting communication overheads between the guest OS and the underlying hypervisor. Consequently, the storage system is becoming a major bottleneck in modern virtualized environments. In this section we examine the sources of these overheads and outline how they can be mediated using a selfvirtualizing storage device.

Prevalent storage devices present the software stack with a raw array of logical block addresses (LBA), and it is up to the OS to provide a flexible method to partition the storage resources into logical objects, or files. In addition, the OS must enforce security policies to prevent applications from operating on data they are not allowed to operate on. The main software layer that provides these functionalities is the filesystem, which combines both allocation and mapping strategies to construct logical objects and map them to physical blocks (for brevity, we focus the discussion on these two functionalities and ignore the plethora of other goals set by different filesystems). In addition, the filesystem layer also maintains protection and access permissions. Besides the filesystem, another common



Fig. 2. The performance benefit of direct device assignment over virtio for high-speed storage devices. Fast devices were emulated using an in-memory disk (ramdisk) whose bandwidth, due to the overheads of the software layers, peaks at 3.6GB/s.

layer is the block layer, which caches disk blocks and abstracts the subtleties of different storage devices from the filesystem layer.

When an application accesses a file, the OS uses the filesystem layer to check the access permissions and map the file offset to an LBA on the storage device. It then accesses its block layer, which retrieves the block either from its caches or from the physical device. In a VM, this process is replicated since the storage device viewed by the guest OS is actually a virtual device that is mapped by the hypervisor to a file on the host's storage system. Consequently, the hypervisor invokes its own filesystem and block layers to retrieve the data to the guest OS.

Figure 1 illustrates the three most common methods by which hypervisors virtualize storage devices:

- Full device emulation [11] (Figure 1a) In this method, the host emulates a known device that the guest already has a driver for. The host traps device accesses by the VM and converts them to operations on real hardware. The emulated device is represented as a file on the host filesystem, and whenever the guest tries to access a virtual LBA on the device, the host converts the virtual LBA to an offset in the file.
- 2) Paravirtualization [12], [8] (Figure 1b) This method, commonly referred to as virtio after its Linux implementation, eliminates the need for the hypervisor to emulate a complete physical device and enables the guest VM to directly request a virtual LBA from the hypervisor, thereby improving performance. This is the most common storage virtualization method used in modern hypervisors.
- 3) Direct device assignment [13] (Figure 1c) This method allows the guest VM to directly interact with the physical device without hypervisor mediation. Consequently, it delivers the best storage performance to the VM. However, since storage devices do not enforce isolation among clients, it does not allow multiple VMs to share a physical device.

Figure 2 quantifies the potential bandwidth speedup of



Fig. 3. Exporting files as virtual devices using NeSC.

direct device assignment over the common virtio interface for high-speed storage devices. We have emulated such devices by throttling the bandwidth of an in-memory storage device (ramdisk). Notably, due to OS overhead incurred by its software layers, the ramdisk bandwidth peaks at 3.6GB/s.

The figure shows the raw write speedups obtained using direct device assignment over virtio for different device speeds, as observed by a guest VM application. Notably, we see that compared to the state-of-the-art virtio method, direct device assignment roughly doubles the storage bandwidth provided to virtual machines for modern, multi GB/s storage devices. The reason for these speedups is that as device bandwidth increases, the software overheads associated with virtualizing a storage device become a performance bottleneck. Using the direct device assignment method eliminates both the virtualization overheads as well as the overheads incurred by the replication of the software layers in the hypervisor and the guest OS.

The potential performance benefits of direct device assignment motivate the incorporation of protection and isolation facilities into the storage device. These facilities will enable multiple guest VMs to share a directly accessed physical device without compromising data protection.

This paper presents the nested storage controller (NeSC), which enables multiple VMs to concurrently access files on the host's filesystem without compromising storage security (when NeSC manages a single disk, it can be viewed simply as a PCIe SSD). Figure 3 illustrates how NeSC provides VMs with secure access to a shared physical device. NeSC leverages the SR-IOV features of the PCIe gen3 [4] to export host files as multiple virtual devices on the PCIe address space. Each virtual device is associated with a collection of non-contiguous blocks on the physical device, which represent a file, and is prevented from accessing other blocks. VMs can therefore directly access the virtual device and bypass the virtualization protocol and hypervisor software (notably, NeSC is compatible with the modern NVMe specification [14]). The rest of this paper describes the design, architecture, and evaluation of NeSC.

III. RELATED WORK

NeSC was motivated by the necessity to redesign the hardware/software storage stack for modern, high-speed, selfvirtualizing storage devices. In this section we examine related research efforts.

The Moneta [15] and Moneta-D [16] projects both try to optimize the I/O stack by introducing new hardware/software interfaces. The Moneta project introduces a storage array architecture for high-speed non-volatile memories. Moneta-D extends the design to allow applications to directly access storage by introducing per-block capabilities [17]. When an application directly accesses the storage, permission checks are done by the hardware to preserve protection policies dictated by the OS. The design, however, requires applications to use a user space library to map file offsets to the physical blocks on the device. In contrast, NeSC enforces file protection by delegating file mapping to the device level rather than to the application level.

Willow [18] examines offloading computation to small storage processing units (SPU) that reside on the SSD. The SPUs run a small OS (SPU-OS) that enforces data protection. Willow implements a custom software protocol that maintains file mapping coherence between the main OS and the SPU-OS instances.

NVMe [14] is a new protocol for accessing high-speed storage devices. Typically implemented over PCIe, NVMe defines an abstract concept of *address spaces* through which applications and VMs can access subsets of the target storage device. The protocol, however, does not specify how address spaces are defined, how they are maintained, and what they represent. NeSC therefore complements the abstract NVMe address spaces and enables the protocol to support protected, self-virtualizing storage devices.

MultiLanes [19] mediates the contention on a hypervisor's shared I/O stack by presenting each VM with a dedicated, virtualized I/O stack. Nevertheless, the underlying device accesses issued by each VM must still be mediated, in software, by the hypervisor.

Peter et al. [20] argue that the SR-IOV concept should be extended and that the operating system should serve as a system's control plain and enable applications to directly access devices. The NeSC storage controller is in line with their view.

Finally, FlashMap [21] provides a unified translation layer for main memory and SSDs, which enables applications to map files to their memory address space. Specifically, FlashMap unifies three orthogonal translation layers used in such scenarios: the page tables, the OS file mapping, and the flash translation layer. FlashMap, however, only supports memory-mapped SSD content, and it does not address how the unified translation would support virtualized environments.

In summary, NeSC's self-virtualizing design is unique in that it delegates all protection checks and filesystem mappings to the hardware. Furthermore, its filesystem-agnostic design preserves the hypervisor's flexibility to select and manage the host filesystem.

IV. DESIGN

In this section we overview the NeSC design principles and basic functionality.

A. The NeSC interface

NeSC implements the SR-IOV specification [4], which allows an I/O device to be shared by multiple VMs. An SR-IOV device can present itself as multiple standard PCIe devices. A single physical function (PF) represents the main devices with all its features. Additionally, the device may dynamically expose multiple virtual functions (VFs), which typically present to clients (e.g., VMs, accelerators) a subset of features supported by the main device. Importantly, each VF has a unique PCIe address and can receive direct I/O requests from its associated client without hypervisor or OS intervention. Since both the PF and the VFs effectively represent different facets of a single physical device, which multiplexes and processes the communication of all the clients with the different facets, it is thus up to the physical device to determine the behavior of the device exported as the PF and that of a virtual device exposed as a VF.

The physical function: The NeSC PF exports a full-featured PCIe storage controller. It is mapped only to the hypervisor address space and allows it to fully manage the physical device. The hypervisor uses the PF interface to (1) manage the underlying NeSC storage and thereby its filesystem; and (2) manage the NeSC virtual devices by controlling the creation and deletion of VFs and the subsets of storage they are allowed to access.

Virtual functions: Each NeSC VF is viewed by the system as a PCIe device that provides a complete block device interface. A VF can thus be mapped to a VM's address space and viewed by it as a fully fledged storage device, which can be programmed by the VM's driver to issue read/write commands to the storage device. Technically, the only difference between the PF and VF interface is that a VF is not allowed to create nested VFs (although, in principle, such a mechanism can be implemented to support nested virtualization).

The PF and VF differ semantically in that VFs can only access the subset of the storage they are associated with. When a VM accesses a VF, all of its read/write block requests are translated to the subset of blocks assigned to that VM as described below.

B. Virtual-to-physical block mapping

Decoupling the PF from the VFs enables the hypervisor to logically manage the physical device with its own filesystem, and expose files (or collections thereof) to client VMs as raw virtual devices through VFs (as illustrated in Figure 3). Each VF is thus associated with a mapping table that translates client requests to physical blocks. Specifically, since client VMs view VFs as regular block devices, they send requests pertaining to LBAs on the virtual device. NeSC refers to client LBAs as virtual LBAs (vLBA), to host LBAs as physical LBAs (pBLA), and the translation process is referred to as a vLBA-to-pLBA translation.



Fig. 4. The extent tree structure used for address translations.

The vLBA-to-pLBA mapping is performed using per-VF mapping tables. The mapping tables are designed as *extent trees*, a method inspired by modern UNIX filesystems. Traditionally, UNIX-derived filesystems used per-file direct and indirect block mapping tables to map offsets in a file to their corresponding data block. But tracking individual blocks incurs large spatial and latency overheads when dealing with large files. Modern UNIX filesystems (e.g., ext4 [22], btrfs [23], xfs [24]) therefore group contiguous physical blocks into *extents* and construct extent trees, which consist of variants of B-trees [25], to spatially map offsets in the device to extents. Each file is associated with an extent tree (pointed to by the file's *inode*) that maps file offsets to physical blocks.

The key benefit of extent trees is that their depth is not fixed but rather depends on the mapping itself. In ext4, for example, a 100MB file can be allocated using a single extent, thereby obviating the need to create the indirect mappings for each individual block. Extents thus improve performance and also reduce management overheads.

Figure 4a illustrates the NeSC VF extent tree (which resembles the ext4 extent tree format). Each node in the tree comprises either a group of node pointers, which point to the next level of nodes, or a group of extent pointers, which are the tree leaves that point to the physical location of the extents. The header of each node indicates whether it contains node indices or extent pointers.

Figure 4b illustrates the content of the entries in each type of node in the extent tree. Each *extent pointer* entry consists of the first logical block it represents, a pointer to the first physical block of its extent, and the size of the extent. Each *node pointer* entry comprises the first logical block it represents, the number of (non-contiguous) logical blocks it covers, and a pointer to its array of child nodes. If memory becomes tight, the hypervisor can prune parts of the extent tree and mark the pruned sections by storing NULL in their respective *Next Node Pointer*. When NeSC needs to access a pruned subtree, it interrupts the host to regenerate the mappings.

Each NeSC's VF is associated with an extent tree, which is stored in host memory, and the NeSC architecture (described in Section V) stores the pointer to the root of each VF's extent tree. Whenever a VF is accessed, its extent tree is traversed using DMA accesses from the device to host memory. To mitigate the DMA latency, extents are cached on the NeSC device (more on this in Section V).

Importantly, this use of software-defined, hardwaretraversed per-VF extent trees eliminates the need to enforce protection and isolation in the hypervisor software layers, as discussed in Section II. Instead, this task is offloaded to hardware and thereby mitigates one of the key performance bottlenecks of virtualized storage [9].

The per-VF extent tree model also enables the hypervisor to decouple the virtual device's size from its physical layout. This allows the hypervisor to initialize virtual devices whose logical size is larger than their allocated physical space, and to allocate further physical space when needed. This enables the hypervisor to maintain a compact physical representation of the stored data.

Finally, the NeSC design also enables multiple VFs to share an extent tree and thereby files. Nevertheless, NeSC only guarantees the consistency of the extent tree for shared files; it is up to the client VMs to address data synchronization and consistency issues.

C. Operational flow

Creating a new virtual disk: When creating a new virtual device, the hypervisor first creates the extent tree that will map the virtual device to physical blocks. Since most modern filesystems use extent trees to map files to their physical layout, this stage typically consists of translating the filesystem's own per-file extent tree to the NeSC tree format. The hypervisor does not fully preallocate all the physical blocks needed for the new virtual device (i.e., lazy allocation).

The hypervisor then creates a new NeSC VF through the PF. It initializes the VF configuration registers (e.g., the size of the virtual device), writes the extent tree to memory and



Fig. 5. Read and write flow in NeSC.

sets the VF's base extent tree configuration register to point to the in-memory extent tree.

Following the two previous steps, the VF is ready and the virtual device is configured. The hypervisor then connects the VM to the new virtual machine, either by executing a new machine or by notifying an existing machine that a new device is available for its use (virtual device hotplug).

Read flow: The read flow is depicted in Figure 5a. The VM's NeSC block driver is attached to the VF and can issue read requests to the device. Large requests are broken down by the driver to scatter-gather lists of smaller chunks. Our NeSC implementation operates at 1KB block granularity (which is, for example, the smallest block size supported by ext4), so the chunks sent by the block driver are broken down by NeSC to 1KB blocks. The device then translates each 1024 byte request address through the extent tree mappings of that virtual function and creates a new queue of physical requests to read from physical storage and DMA back to the host memory.

The translation of a vLBA to a pLBA can fail in one of two cases: 1) The pLBA was never allocated due to lazy allocation. Following the POSIX standard, which dictates that unmapped areas inside a file (holes in the file) should read as zeros, NeSC transparently DMAs zeros to the destination buffer in host memory; or 2) The pLBA was allocated, but its mapping was pruned from the extent tree due to memory pressure. NeSC identifies this case when it encounters a NULL node pointer when traversing the tree, and resolves it by sending an interrupt to the host and requesting that the hypervisor regenerates the pruned mappings (e.g., by re-reading the filesystem-specific mappings from disk).

Write flow: The write flow is shown in Figure 5b. Similar to read requests, write requests are broken down by NeSC to 1KB chunks.

For each chunk, the device tries to perform vLBA-to-pLBA mapping using the VF's extent tree. If the translation succeeds, the data is written to the persistent physical storage.

Just like in the read case, the translation can fail due to unallocated blocks or pruned tree nodes. Both cases require NeSC to interrupt the hypervisor, which will request the filesystem for the mappings and rebuild the tree. The filesystem might have to allocate new blocks if this is the first time the vLBA is accessed. Once the hypervisor finishes handling the request, it signals NeSC to restart the extent tree lookup, which is now guaranteed to succeed. If, however, the hypervisor cannot allocate more space for the VF (not shown in Figure 5b) due to a lack of physical storage or exhausted storage quotas, it signals an error to the PF which, in turn, triggers the VF to send a write failure interrupt to the requesting VM (we note that it is possible to optimize this process by delegating the interrupt directly to the VM [26]).

D. Other design issues

Nested filesystems: A client VM will often manage its own filesystem inside its nested storage device. Given that a virtual

device is stored as a file in the hypervisor's filesystem, this use case is commonly referred to as a nested filesystem.

Modern filesystems frequently use journaling for improved fault tolerance. This causes a well-known inefficiency in nested filesystems known as nested journaling [9]. The inefficiency is caused as both the internal and external filesystems redundantly log the internal filesystem's data and meta-data updates. The common solution to this inefficiency is to tune the hypervisor's filesystem to only log meta-data changes for the file at hand and let the VM handle its internal filesystem's data integrity independently.

NeSC naturally lends itself to this common solution. Since NeSC VM clients directly access their data, the hypervisor's filesystem is not aware of the internal filesystem updates, whose integrity it handled by the VM, and only tracks its own meta-data updates.

Direct storage accesses from accelerators: A straightforward extension of NeSC is to export data to accelerators in the system. Traditionally, when an accelerator on the system needs to access storage, it must use the host OS as an intermediary and thereby waste CPU cycles and energy.

While not implemented in the prototype, we note that NeSC can be easily extended to enable direct accelerator-storage communications. This can be easily achieved by modifying the VF request-response interface, which is suitable for block devices, to a direct device-to-device DMA interface (in which offset 0 in the device matches offset 0 in the file, and so on).

This simple change to the VF interface enables accelerators to directly access storage using DMA without interrupting the main processor and the OS. Such a design corresponds with the data-centric server concept presented by Ahn et al. [27].

Quality of Service (QoS): Modern hypervisors support different QoS policies for each emulated disk image. Because the hypervisor mediates every disk access, it can easily enforce the QoS policy for each device. NeSC can be extended to enforce the hypervisor's QoS policy by modifying its DMA engine to support different priorities for each VF. Similar techniques are widely used by NICs that support SR-IOV.

Unified buffer cache: Different VMs often share storage blocks with one another. To minimize buffer duplication across VMs' buffer caches, common hypervisors use their own buffer cache as a unified buffer cache by forcing VMs to minimize their individual buffer caches (specifically, balloon drivers in each VM reclaim physical pages; the resulting memory pressure forces the VM to free buffered pages). As NeSC enables VMs to directly access storage, it prevents the hypervisor from maintaining a unified buffer cache. Instead, NeSC relies on memory deduplication to eliminate buffer duplication across guest VMs'. Notably, memory deduplication is supported by all modern hypervisors (e.g., TPS in VMware, KSM in kvm and "Memory CoW" in Xen).

V. ARCHITECTURE

As an SR-IOV device, NeSC presents itself as multiple devices on the PCIe interconnect. The architecture must maintain a separate context for each PCIe device (PF and VFs), but a



Fig. 6. An outline of the NeSC architecture. The PF and VFs each have a set of control registers and request/response queues. The core NeSC microarchitecture multiplexes the activity of the individual VFs.

key tenet of the microarchitecture is multiplexing traffic from the different devices.

Figure 6 outlines the core design of the NeSC architecture. For each function (PF and VF alike), NeSC maintains a set of control registers and request/response queues. All requests sent to the different functions are multiplexed through the address translation and request service units. Similarly, all traffic between the host and the device is multiplexed through a single DMA engine.

Every request received by the NeSC device is labeled with the ID of the NeSC function to which it was sent. This enables the multiplexed facilities to associate requests with a specific function (according to the SR-IOV specification, the ID of the PF is 0). PCIe addressing is hierarchical, and each entity on the interconnect is identified by a triplet: the PCIe *bus ID*, the ID of the device on the bus (*device ID*), and the function inside the device (*function ID*). In PCIe parlance, addresses are referred to as *bus:device:function* triplets or *BDF* for short. Since PCIe addresses of the different NeSC functions share the same bus and device IDs, associating each request with its originating function ID is sufficient for accurate bookkeeping. The BDF triplet is originated by the PCIe interface of NeSC and is unforgeable by a virtual machine.

Control registers

Communication with a PCIe device is handled through *base* address registers (BARs). BARs have a predetermined size, and are mapped to the system's logical address space when the PCIe interconnect is scanned (typically by the system's firmware). Since each BAR is assigned a logical bus address, the hypervisor can map the BARs directly to clients' virtual address spaces, and to its own.

Although NeSC assigns a BAR for each function (physical and virtual alike), the control registers for all functions can be mapped to a single physical structure inside the device. The



Fig. 7. High-level view of the NeSC virtual function multiplexer design. The microarchitecture multiplexes requests from the different virtual functions.

PCIe specifications make it possible to define control registers as offsets into a device's BAR. When a certain offset inside a BAR is read from or written to, the PCIe device internally translates the offset to an address in a shared structure. Since our prototype can support up to 64 VFs (and the additional PF), the NeSC prototype uses a single 130KB SRAM array (2048B per function).

The set of VF-specific control registers are listed as follows (the PCIe specification mandates a number of additional standard control registers, omitted for brevity).

- 1) **ExtentTreeRoot** (8B) This register contains the base address (in host memory) of the root node of the extent tree associated with a VF. It is set by the hypervisor when a new VF is created.
- 2) **MissAddress, MissSize** (8B, 4B) These two registers are set by NeSC when a translation of a write request sent to the VF misses in the extent tree. After setting these values, NeSC sends an interrupt to the hypervisor so it will allocate additional physical space to the file associated with the VF.
- 3) RewalkTree (4B) This register is used by the hypervisor to signal NeSC that the physical storage allocation has succeeded. When the hypervisor writes a value of 1 to a VF's RewalkTree register, NeSC reissues the stalled write requests to the extent tree walk unit.

In addition to the NeSC-specific control registers listed above, each VF also exposes a set of registers for controlling a DMA ring buffer [28], which is the de facto standard for communicating with devices. We omit those for brevity.

A. The NeSC virtual function multiplexer

In this section we describe the implementation of the NeSC virtual function multiplexer, which is depicted in Figure 7. The main role of the virtual function multiplexer is to process access requests from the multiple VFs, translate their addresses to physical storage blocks, and respond to the client VM. Access requests are sent by the VM's block driver. The driver typically breaks large requests into a sequence of smaller 4KB

requests that match the system's page size. Requests sent from different client VMs are stored in per-client request queues. NeSC dequeues client requests in a round-robin manner in order to prevent client starvation (a full study of client scheduling algorithms and different quality-of-service guarantees is beyond the scope of this paper).

Client requests address their target storage blocks using vLBAs, or logical block addresses as viewed by the virtual device. These addresses must be translated to pLBAs using the per-VF extent tree. The client requests are therefore pushed to a shared vLBA queue for translation, along with the root pointer to their VF's extent tree.

NeSC includes a dedicated *translation unit* (discussed below) that dequeues pending requests from the vLBA queue, translates their vLBAs to pLBAs and pushes the translated requests to a pLBA queue. A *data transfer* unit then manages direct accesses to the persistent storage device. For write requests, the data transfer unit simply writes the data to storage and generates an acknowledge response message that will be sent to the client. For read requests, the data transfer unit reads the target data from the physical storage and prepares a response message. The message is then DMAed to the client VM's destination buffer in host memory.

While we only discuss the processing of client VM requests, NeSC also includes a dedicated out-of-band (OOB) channel to process hypervisor requests sent through the PF. The OOB channel is needed so that VF write requests whose translation is blocked will not block PF requests. The addition of the OOB is, however, simple since PF requests use pLBAs and need not be translated. The OOB channel, therefore, bypasses the NeSC flow preceding the pLBA queue and does not affect the the address translation unit.

B. The vLBA-to-pLBA translation unit

The translation unit translates the vLBA addresses used in client VM requests to pLBA addresses, which can be used to access the physical storage. The translation uses the extent tree associated with a request's originating VF.



Fig. 8. The vLBA-to-pLBA translation unit.

Figure 8 illustrates the translation unit and its components. These include a *block walk unit* that traverses the extent tree and a *block translation lookaside buffer* (BTLB) that caches recent vLBA mappings. Furthermore, the unit may send interrupts to the hypervisor when a mapping is not found (e.g., when a client VM writes to an unallocated portion of its virtual device), and it is charged with observing the VFs' *RewalkTree* registers through which the hypervisor signals that the mapping was fixed and the extent tree can be re-examined.

Block walk unit: This unit executes the block walk procedure by traversing the extent tree. When a client VM request arrives at the unit, the root node of the associated extent tree is DMAed and the unit tries to match the vLBA to the offsets listed in the root node's entries (be they node pointers or extent pointers). The matched entry is used as the next node in the traversal, and the process continues recursively until an extent is matched. For read operations, an unmatched vLBA means that the client VM is trying to read from an unmapped file offset (a hole in the file) and zeros must be returned. If a match is not found on a write operation, the unit sets the VF's *MissAddress* and *MissSize* control registers, interrupts the host, and waits until the hypervisor signals that the missing blocks were allocated (using the *RewalkTree* control register).

The block walk unit is designed for throughput. Since the main performance bottleneck of the unit is the DMA transaction of the next level in the tree, the unit can overlap two translation processes to (almost) hide the DMA latency.

Block Translation lookaside buffer (BTLB): Given that storage access exhibits spatial locality, and extents typically span more than one block, the translation unit maintains a small cache of the last 8 extents used in translation. Specifically, this enables the BTLB to maintain at least the last mapping for each of the last 8 VFs it serviced.

Before the translation unit begins the block walk, it checks the BTLB to see whether the currently translated vLBA matches one of the extents in the BTLB. If a match is found, the unit skips the block walk and creates a new pLBA for the output queue. On a miss, the vLBA will be passed to the block walk unit and, after the translation completes, the mapping will be inserted to the BTLB (evicting the oldest entry). Finally, the BTLB cache must not prevent the hypervisor from executing traditional storage optimizations (e.g., block deduplication). Consequently, NeSC enables the PF (representing the hypervisor) to flush the BTLB cache in order to preserve meta-data consistency.

In summary, the NeSC implementation effectively balances the need to multiplex different VFs with the application of traditional optimizations such as caching and latency hiding. The following sections describe our evaluation of the proposed NeSC design.

VI. METHODOLOGY

Experimental system: Table I describes our experimental system, which consists of a Supermicro X9DRG-QF server equipped with a Xilinx VC707 evaluation board.

We implemented NeSC using a VC707's Virtex-7 FPGA. Since this revision of the Virtex7 FPGA only supports PCIe Gen2, we had to emulate the self-virtualizing features rather than use the SR-IOV protocol extension. We have emulated the SR-IOV functionalities by dividing the device's memorymapped BAR to 4KB pages. The first page exports the NeSC PF, and subsequent pages export complete VF interfaces. Upon QEMU startup, the hypervisor maps one of the VF interfaces (offset in the BAR) into the physical address space of the VM. A multiplexer in the device examines the address to which each PCIe packet (TLP) was sent and queues it to the appropriate VF's command queue. For example, a read TLP that was sent to address 4244 in the NeSC device would have been routed by the multiplexer to offset 128 in the first VF.

Furthermore, since the emulated VFs are not recognized by the IOMMU, VMs cannot DMA data directly to the VC707 board. Instead, the hypervisor allocates trampoline buffers for each VM, and VMs have to copy data to/from the trampoline buffers before/after initiating a DMA operation.

We note that both SR-IOV emulation and trampoline buffers actually impede the performance of NeSC and thereby provide a pessimistic device model. Using a true PCIe gen3 device would improve NeSC's performance.

Host machine		
Machine model	Supermicro X9DRG-QF	
Processor	Dual socket Intel(R) Xeon(R) CPU E5-	
	2665 @ 2.40GHz (Sandybridge)	
Memory	64GB DDR3 1333 MHz	
Operating system	Ubuntu 12.04.5 LTS (kernel 3.5.0)	
Virtualized system		
Virtual machine	QEMU version 2.2.0 with KVM	
monitor		
Guest OS	Linux 3.13	
Guest RAM	128MB	
Filesystem on NeSC	ext4	
volume		
Prototyping platform		
Model	Xilinx VC707 Evaluation board	
FPGA	Virtex-7 (XC7VX485T-2FFG1761C)	
RAM	1GB DDR3 800MHz	
Host I/O	PCI Express x8 gen2	
ΤΔΒΙΕΙ		

EXPERIMENTAL PLATFORM

Microbenchmark		
GNU dd [29]		read/write files using different operational
		parameters.
Macrobenchmarks		
Sysbench	File	A sequence of random file operations
I/O [30]		
Postmark [31]		Mail server simulation
MySQL [32]		Relational database server serving the
		SysBench OLTP workload

TABLE II Benchmarks

We further note that we do not emulate a specific access latency technology for the emulated storage device. Instead, we simply use direct DRAM read and write latencies.

We used the QEMU/KVM virtual machine monitor. Since the VC707 board only has 1GB of RAM, we could only emulate 1GB storage on the NeSC device. In order to prevent the entire simulated storage device from being cached in RAM, we limited the VM's RAM to 128MB. We validated that this limitation does not induce swapping in any of the benchmarks.

Finally, in order for NeSC to truly function as a selfvirtualizing device, we implemented the VF guest driver, which is a simple block device driver, and the hypervisor PF driver, which acts as both a block device driver and as the NeSC management driver for creating and deleting VFs.

Benchmarks: Table II lists the benchmarks used to evaluate NeSC. We first evaluated read/write performance metrics (e.g., bandwidth, latency) using the *dd* Unix utility. In addition, we used a common set of macrobenchmarks that stress the storage system. All applications used the virtual device through an underlying ext4 filesystem.

VII. EVALUATION

This section reports the evaluation results of the NeSC prototype. We examine the performance benefits of enabling VMs to directly access NeSC, and compare its performance



Fig. 9. Raw access latency observed for read (top) and write (bottom) operations for different block sizes.

to those of virtio and device emulation. Our baseline storage device is the NeSC PF, which presents the hypervisor a raw storage device with no file mapping capabilities. The baseline measurement is done by the hypervisor without any virtualization layer.

A. Raw device performance

We begin by examining the raw device performance observed by the guest VM when accessing a virtual NeSC device. A file on the hypervisor's filesystem is used to create a VF, which is then mapped to the guest VM. These results are compared to mapping the PF itself to the guest VM using either virtio and device emulation. The baseline (marked *Host* in the figures) is the performance observed when the hypervisor directly accesses the PF block device (without virtualization). In all configurations, we examine the performance of reads and writes to the raw virtual device, without creating a filesystem.

When a guest accesses the device using virtio or device emulation, each access request is processed by the guest I/O stack and storage device driver, delivered to the hypervisor, and is then processed by the hypervisor's I/O stack and its own device driver. When directly assigning a NeSC VF to the guest, requests pass down the guest I/O stack directly to the device. In contrast, in the baseline evaluation (i.e., *Host* in the figure) the requests only pass down the hypervisor I/O stack.

The performance itself was measured using dd [29] for different block sizes.

NeSC latency: Figure 9 shows the latency observed for read (top) and write (bottom) operations, using request sizes varying from 512B to 32KB. The figure shows that the latency obtained by NeSC for both read and write is similar to



Fig. 10. Raw bandwidth observed for read (top) and write (bottom) operations for different block sizes.

that obtained by the host when directly accessing the PF. Furthermore, the NeSC latency is over $6 \times$ faster than virtio and over $20 \times$ faster than device emulation for accesses smaller than 4KB.

NeSC bandwidth: Figure 10 shows the bandwidth delivered for read (top) and write (bottom) operations, using request sizes varying from 512B to 32KB.

We begin by examining the read bandwidth. The figure shows that for reads smaller than 16KB, NeSC obtained bandwidth close to that of the baseline and outperforms virtio by over $2.5\times$. Furthermore, for very large block sizes (over 2MB), the bandwidths delivered by NeSC and virtio converge. This is because such large accesses ameliorate the overheads incurred by VM/hypervisor traps (vmenter/vmexit on Intel platforms).

When examining the write bandwidth, we observe that NeSC delivers performance similar to the baseline for all block sizes. This result is better than that achieved for read bandwidth, for which NeSC is $\sim 10\%$ slower than the baseline for blocks of size 32KB and larger. Moreover, NeSC's write bandwidth is consistently and substantially better than virtio and emulation, peaking at over $3\times$ for 32KB block sizes.

Filesystem overheads: We next examine the overheads incurred by filesystem translations. To this end, we compare the access latency observed by the guest VM when accessing the raw device with that observed when accessing an ext4 filesystem on the virtual device.

Figure 11 shows the write latency when accessing the device with and without a filesystem, for NeSC and virtio virtualization (for brevity, we omit the results obtained using an emulated device since they were much worse than virtio).



(a) Applications speedups when using NeSC over device emulation. (b) Application speedups when using NeSC over virtio.

Fig. 12. Application speedups over other storage virtualization methods.

We only show results for writes because those are more prohibitive for NeSC, since writes may require the VF to request extent allocations from the OS's filesystem.

The figure demonstrates the performance benefits of NeSC. While the filesystem overhead consistently increases NeSC's write latency by $\sim 40\mu s$, the latency obtained using NeSC is similar to that of a raw virtio device. Using virtio with a filesystem incurs an extra $\sim 170\mu s$, which is over $4 \times$ slower than NeSC with a filesystem for writes smaller than 8KB.

We note that the similar latencies observed when using a filesystem with NeSC and when using a raw virtio device indicate that NeSC effectively eliminates the hypervisor's filesystem overheads.

B. Application performance

We finally examine how the raw performance delivered by NeSC translates to application level speedups. To this end, we compare the performance of the applications (listed in Figure II) when running in a guest Linux VM whose storage device is mapped to a NeSC VF, to a virtio device and to a fully emulated device.

The virtual storage device is mapped to the VM in the following manner. The hypervisor creates an ext4 filesystem on the raw device using the NeSC PF. The virtual storage device is stored as an image file (with ext4 filesystem) on the hypervisor's filesystem, and the hypervisor maps the file to the VM using either of the mapping facilities: virtio, emulation or a NeSC VF.

Figure 12 shows the application speedups obtained using a NeSC VF device over device emulation and virtio. For the MySQL OLTP benchmark, the figure shows that mapping the virtual disk using a NeSC virtual device improves performance by $1.2 \times$ over both device emulation (Figure 12a) and virtio (Figure 12b).

The performance improvements provided by NeSC are even more substantial for Postmark and Sysbench File I/O. Postmark enjoys more than $6\times$ speedup over device emulation and more than $2\times$ over virtio. Sysbench File I/O, on the other hand, enjoys a dramatic $13\times$ speedup over device emulation, and $1.3\times$ over virtio.

In summary, the evaluation of the NeSC prototype demonstrates its substantial performance benefits over state-of-the-art storage virtualization methods. The benefits are consistent for both read and write microbenchmarks and for common storage benchmarks.

VIII. CONCLUSIONS

The emergence of multi-GB/s storage devices has shifted storage virtualization bottlenecks from the storage devices to the software layers. System designers must therefore delegate the virtualization overheads in order to enable virtualized environments to benefit from high-speed storage.

In this paper we presented the *nested*, *self-virtualizing storage controller* (NeSC), which enables files stored on the hypervisor-managed filesystem to be directly mapped to guest VMs. NeSC delegates filesystem functionality to the storage device by incorporating mapping facilities that translate file offsets to disk blocks. This enables NeSC to leverage the selfvirtualization SR-IOV protocol and thereby expose files as virtual devices on the PCIe interconnect, which can be mapped to guest VMs.

We prototyped NeSC using a Virtex-7 FPGA and evaluated its performance benefits on a real system. Our comparison of NeSC to the leading *virtio* storage virtualization method, shows that NeSC effectively eliminates the hypervisor's filesystem overheads, as filesystem accesses to a NeSC virtual device incur the same latency as accesses to a raw virtio device. Furthermore, we have shown that NeSC speeds up common storage benchmarks by $1.2 \times -2 \times$.

ACKNOWLEDGMENT

We thank the anonymous referees for their valuable comments. We also thank Lluis Vilanova for his insightful suggestions. This research was funded by the Israel Science Foundation (ISF grant 769/12) and the Israel Ministry of Science, Technology, and Space. Y. Etsion was supported by the Center for Computer Engineering at Technion.

REFERENCES

- G. J. Popek and R. P. Goldberg, "Formal requirements for virtualizable third generation architectures," *Comm. ACM*, vol. 17, no. 7, pp. 412–421, 1974.
- [2] Intel, Intel 64 and IA-32 Architectures Software Developer Manual, Dec 2015.
- [3] ARM, ARMv8 Architecture Reference Manual, Jan 2016.
- [4] PCI-SIG, "PCI I/O virtualization." http://pcisig.com/specifications/iov/.

- [5] I. Smolyar, M. Ben-Yehuda, and D. Tsafrir, "Securing self-virtualizing Ethernet devices," in USENIX Security Symp., pp. 335–350, Aug 2015.
- [6] Intel, "Product brief: Intel solid state drive data center p3608 series," 2015.
- [7] Seagate, "Press release: Seagate demonstrates fastest-ever SSD flash drive." www.seagate.com/em/en/about-seagate/news/seagatedemonstrates-fastest-ever-ssd-flash-drive-pr/, Mar 2016.
- [8] R. Russell, "virtio: towards a de-facto standard for virtual I/O devices," Operating Systems Review, vol. 42, no. 5, pp. 95–103, 2008.
- [9] D. Le, H. Huang, and H. Wang, "Understanding performance implications of nested file systems in a virtualized environment," in USENIX Conf. on File and Storage Technologies (FAST), 2012.
- [10] Y. J. Yu, D. I. Shin, W. Shin, N. Y. Song, J. W. Choi, H. S. Kim, H. Eom, and H. Y. Yeom, "Optimizing the block I/O subsystem for fast storage devices," ACM Trans. on Computer Systems, vol. 32, pp. 1–48, June 2014.
- [11] J. Sugerman, G. Venkitachalam, and B.-H. Lim, "Virtualizing i/o devices on vmware workstation's hosted virtual machine monitor.," in *Usenix Annual Technical Conf.*, pp. 1–14, 2001.
- [12] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, "Xen and the art of virtualization," in ACM Symp. on Operating Systems Principles (SOSP), pp. 164–177, 2003.
- [13] H. Raj and K. Schwan, "High performance and scalable i/o virtualization via self-virtualized devices," in *Intl. Symp. on High performance Distributed Computing (HPDC)*, pp. 179–188, 2007.
- [14] NVM Express, NVM Express Specification, Oct 2015. Revision 1.2a.
- [15] A. M. Caulfield, A. De, J. Coburn, T. I. Mollow, R. K. Gupta, and S. Swanson, "Moneta: A high-performance storage array architecture for next-generation, non-volatile memories," in *Intl. Symp. on Microarchitecture (MICRO)*, pp. 385–395, 2010.
- [16] A. M. Caulfield, T. I. Mollov, L. A. Eisner, A. De, J. Coburn, and S. Swanson, "Providing safe, user space access to fast, solid state disks," in *Intl. Conf. on Arch. Support for Programming Languages & Operating Systems (ASPLOS)*, pp. 387–400, 2012.
- [17] H. M. Levy, Capability-Based Computer Systems. Digital Press, 1984.
- [18] S. Seshadri, M. Gahagan, S. Bhaskaran, T. Bunker, A. De, Y. Jin, Y. Liu, and S. Swanson, "Willow: A user-programmable SSD," in *Symp. on Operating Systems Design & Impl. (OSDI)*, pp. 67–80, 2014.
- [19] J. Kang, B. Zhang, T. Wo, C. Hu, and J. Huai, "Multilanes: providing virtualized storage for os-level virtualization on many cores," in USENIX Conf. on File and Storage Technologies (FAST), pp. 317–329, 2014.
- [20] S. Peter, J. Li, I. Zhang, D. R. Ports, D. Woos, A. Krishnamurthy, T. Anderson, and T. Roscoe, "Arrakis: The operating system is the control plane," ACM Trans. on Computer Systems, vol. 33, no. 4, p. 11, 2016.
- [21] J. Huang, A. Badam, M. Qureshi, and K. Schwan, "Unified address translation for memory-mapped SSDs with FlashMap," in *Intl. Symp.* on Computer Architecture (ISCA), pp. 580–591, 2015.
- [22] A. Mathur, M. Cao, S. Bhattacharya, A. Dilger, A. Tomas, and L. Vivier, "The new ext4 filesystem: current status and future plans," in *Ottawa Linux Symp.*, Jun 2007.
- [23] O. Rodeh, J. Bacik, and C. Mason, "BTRFS: The linux B-Tree filesystem," ACM Trans. on Storage, vol. 9, pp. 1–32, Aug. 2013.
- [24] A. Sweeney, D. Doucette, W. Hu, C. Anderson, M. Nishimoto, and G. Peck, "Scalability in the XFS file system," in Usenix Annual Technical Conf., 1996.
- [25] D. Comer, "Ubiquitous B-Tree," ACM Computing Surveys, vol. 11, pp. 121–137, June 1979.
- [26] A. Gordon, N. Amit, N. Har'El, M. Ben-Yehuda, A. Landau, A. Schuster, and D. Tsafrir, "ELI: Bare-metal performance for I/O virtualization," in *Intl. Conf. on Arch. Support for Programming Languages & Operating Systems (ASPLOS)*, 2012.
- [27] J. Ahn, D. Kwon, Y. Kim, M. Ajdari, J. Lee, and J. Kim, "DCS: a fast and scalable device-centric server architecture," in *Intl. Symp. on Microarchitecture (MICRO)*, 2015.
- [28] R. Love, *Linux Kernel Development*. Addison-Wesley Professional, 3rd ed., Jul 2010.
- [29] "Gnu coreutils." http://www.gnu.org/software/coreutils/coreutils.html.
- [30] A. Kopytov, "SysBench: a system performance benchmark." http://sysbench.sourceforge.net, 2004.
- [31] J. Katcher, "Postmark: A new file system benchmark," Tech. Rep. TR3022, Network Appliance, 1997.
- [32] "MySQL." https://www.mysql.com.