

# Single Root I/O Virtualization

## Economical Optimization for Virtual Machines



The demand for server capacity is growing rapidly worldwide thanks to the ever-increasing amount of remote employment and engagement, teleconferencing, and robust desire for cloud-based content. These services rely on virtualization using virtual machines (VMs) and other technologies to streamline the apportionment of resources to meet the needs of the underlying structure. However, physical space and

budgets are both limited. As a result, efficient use of server hardware is always the primary goal. The Peripheral Component Interconnect Special Interest Group's (PCI-SIG) single root I/O virtualization (SR-IOV) specification works towards meeting this objective by providing a method of managing peripheral component interconnect express (PCIe) resources in a virtualized environment.

## Challenges

Current virtualization environments must overcome a number of architectural challenges with regard to overall network and input/output (I/O) performance:

- Software solutions for resource-sharing increase I/O overhead, introducing additional latency due to the need for an emulation layer
- Network throughput is reduced when using one physical device per virtual machine or when utilizing software-based sharing
- CPU and memory utilization increase as virtual machines are added due to the increased switching by the virtual switch of the managing hypervisor or virtual machine manager (VMM)
- Interrupt handling can bottleneck a CPU and effectively decrease the VM limit, or amount of VMs per physical server

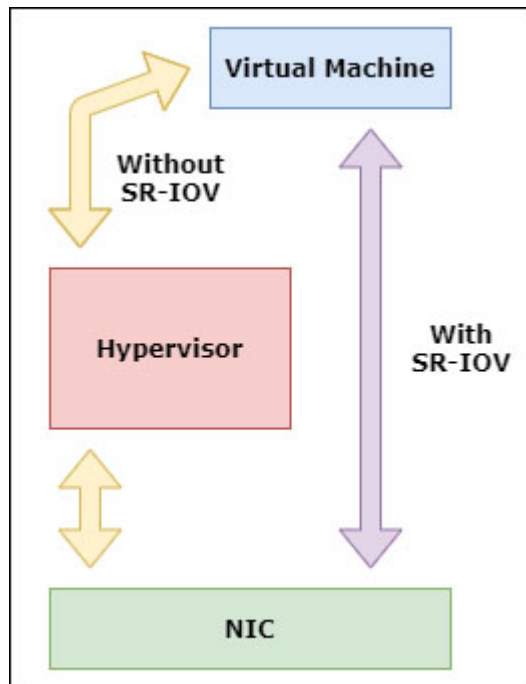
## Goals

There are some fundamental conditions for overcoming these and other challenges in the virtualized space which any solution must address:

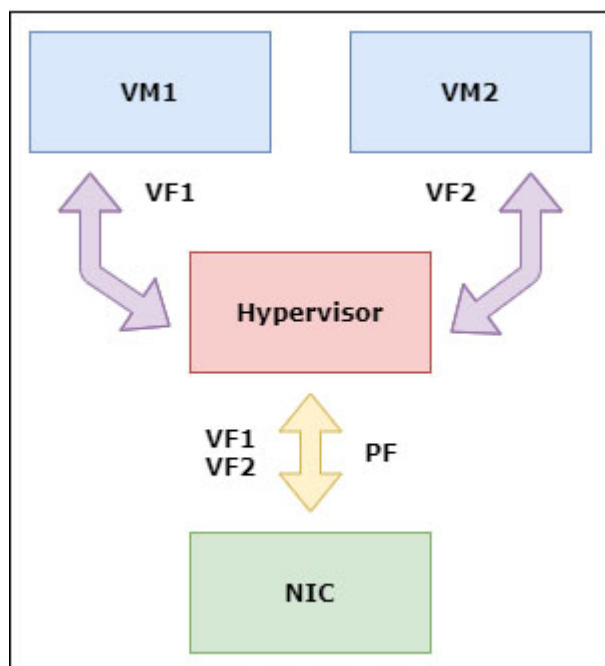
- There is always a desire for bare-metal or near native I/O performance, as you would have with an actual physical device
- Data movement must be standardized, which includes interrupts and memory functions, by bypassing the hypervisor
- VMs need to maintain isolation for security and data protection while also maintaining mobility for easier physical server relocation – this would also increase flexibility through VM independence

- Software-based overhead must be reduced especially with regard to CPU utilization in order to maintain processing cycles for useful work
- Scalability is a primary objective with the end goal of having a greater amount of VMs per physical server with maximum resource utilization, reducing the physical space requirement and overall cost

## SR-IOV Overview



*Data path with and without SR-IOV*



*VF & PF allocation*

The basis of SR-IOV is the interaction of physical functions (PF) and virtual functions (VF). A physical function would be the shared hardware PCIe device or resource, such as an Ethernet port on a singular network interface controller (NIC). A virtual function is a “lightweight” PCIe function with minimal overhead associated with that physical function. SR-IOV is able to share the PF I/O port without software emulation, reusing existing drivers and software – assuming the hardware device has support – through a mechanism to make the resource present as multiple devices for VF usage.

Each VF may have direct assignment to a specific VM, configured by the hypervisor, and may achieve near native performance with its own individual and unique data path. This partitions the PCIe function into various virtual interfaces, allowing multiple VMs to share a single hardware resource. This improves the efficiency of the PF port by maximizing throughput while reducing network traffic overhead by bypassing the software emulation layer in the hypervisor. Specialized PCIe devices may also have additional features for superior SR-IOV coordination.

## Summary

As the world adapts to changing conditions the expansion of server capability is a nature consequence to meet newfound demand. Rather than simply construct more infrastructure, it can be more economical to improve existing hardware while planning for the next generation by having or acquiring built-in support through innovative virtualization technologies. SR-IOV as a standard provides such an option by outlining an effective method of I/O resource sharing.

SR-IOV can achieve near native I/O performance for each VM by bypassing the software layer traditionally used to divide physical network resources. This comes with a reduction in network traffic latency and an improvement in network throughput. By reutilizing existing drivers and software while bypassing the hypervisor, one can reclaim critically limited CPU cycles and memory. This also ensures that physical devices see an optimal amount of utilization with minimal waste, especially desirable with scalable systems as you find within the cloud.

Also of importance in the market is the ability to move and aggregate hardware, including the need for live migration. This grants VMs the mobility to transfer between physical servers without minimal setup, making for smoother transitions. Additionally, the pass-through functionality allows continued isolation of VMs including for the purposes of data protection and system security. This increases server flexibility through VM independence. Coupled with the gains achieved by avoiding software-based overhead this significantly increases the amount of VMs achievable on a given physical server, reducing overall cost and the amount of physical space required while remaining scalable.

SR-IOV therefore offers a framework for improving server and network efficiency by improving upon the architecture of a virtualized environment. This not only increases the overall performance and capacity of the system, but also further reduces potential management struggles by offering increased flexibility. The ability to ramp up and relocate VMs is of critical importance as more of the world goes online, and SR-IOV facilitates this while maintaining a standard of security required by data-driven organizations in the 21st century.

## Our SSD Solution



Our ED1 Series with SR-IOV feature is a powerful, high performance SSD made for data center and hyperscale applications. It comes in M.2 and U.2 form factors. [Contact us](#) today to find out why Tier 1 cloud providers and Fortune 500 Enterprise companies adopted our SSDs.

Specifications and features are subject to change without prior notice. Images are samples only, not actual products. Please check with an authorized SSSTC representative for details.

## ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

© 2020 Solid State Storage Technology Corporation. All rights reserved.

**Learn more at [www.ssstc.com](http://www.ssstc.com)**