A lightweight isolated execution environment in SmartNIC

Manami Mori
manami.mori@aist.go.jp
National Institute of Advanced Industrial Science and Technology / Tokyo Metropolitan University

Takaaki Fukai
takaaki.fukai@aist.go.jp
National Institute of Advanced Industrial Science and Technology

Ryuichi Sakamoto
r-sakamoto@gsic.titech.ac.jp
Tokyo Institute of Technology

Takahiro Hirofuchi
t.hirofuchi@aist.go.jp
National Institute of Advanced Industrial Science and Technology

Takuya Asaka
asaka@tmu.ac.jp
Tokyo Metropolitan University

Recently, an increase in data volumes from IoT services and big data has led to a significant increase in CPU workload for data processing. To process the data more efficiently and reduce the load, I/O devices with general-purpose processors, such as Arm processors, are introduced.

SmartNIC, which is one of these devices, has been introduced for various data processing tasks, such as the TCP stack. Opportunities to offload user application processes to general-purpose processors in multi-tenant environments are expected to increase in the future.

The problem with using SmartNIC in multi-tenants is the lack of secure isolation between offloaded tasks in the device. In the cloud system, some data processing tasks require high-level security. Therefore, virtual machines are used to isolate each tenant on the host CPU to prevent attacks from malicious tenants. On the other hand, the current offloading process runs on the SmartNIC without sufficient isolation.

To isolate processes, we can consider software-based virtualization and hardware-based approaches. However, current technology and state-of-the-art studies have problems in terms of performance, security, scalability, and adaptability to production-level SmartNICs. The virtual machine monitor provides the environment to run the operating systems separately. However, virtualization on the SmartNIC causes significant performance degression. Containers provide lightweight isolation, however that is insufficient against attacks that exploit vulnerabilities in the operating system.

S-NIC[3] is a state-of-the-art SmartNIC hardware design that provides strong isolation. S-NIC provides a partitioned set of hardware resources for each tenant and avoids sharing any hardware resources between tenants. Management cores can access the hardware resource for launching the function. After launch, the new function is isolated from the management cores. S-NIC works without virtualization overhead and it provides sufficient isolation for cloud systems. However, this approach cannot support dynamic resource scaling up and down, as S-NIC is unable to modify the environment's resources without tenant shutdown. Specifically, tasks in cloud systems often require resource adjustments according to the volume of requests. In addition, S-NIC needs hardware customization, including processors. This prohibits us from applying the design into the production-level SmartNIC which usually has hard-core processors such as the NVIDIA BlueField series and Zynq UltraScale+.

For sufficient isolation, high performance, and adaptability to production-level SmartNIC, we propose a lightweight task isolation method with minimal I/O virtualization by a hypervisor. In this method, the hypervisor splits devices logically instead of virtualizing them while it restricts access from each environment to other devices that are not assigned.

We are implementing a prototype of our proposal based on MilvusVisor[1] that is, for the Arm 64-bit architecture. We tested a SmartNIC implemented on an Alveo U25 FPGA board with Zynq Ultrascale+ ARM SoC and a TCP offload Engine[2]. For offloading tasks, FreeRTOS runs in the lightweight isolated environment.

Acknowledgments

This work was supported by JST, CREST Grant Number JPMJCR22M3, Japan.

References

[1] Manami Mori and Takaaki Fukai. 2022. MilvusVisor - A thin-hypervisor that runs on aarch64 CPUs. https://github.com/RIKEN-RCCS/MilvusVisor

[2] David Sidler et al. 2015. Scalable 10Gbps TCP/IP Stack Architecture for Reconfigurable Hardware. In *2015 IEEE 23rd Annual International Symposium on Field-Programmable Custom Computing Machines*. 36–43.

[3] Yang Zhou et al. 2024. SmartNIC Security Isolation in the Cloud with S-NIC. In *Proceedings of the Nineteenth European Conference on Computer Systems* (Athens, Greece) *(EuroSys '24)*. Association for Computing Machinery, New York, NY, USA, 851–869.