

Toward Optimizing OS Image Deployment Using SmartNIC

Taiki Miyakawa
Shibaura Institute of Technology
Tokyo, Japan
ma23185@shibaura-it.ac.jp

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

Li Yanzhi
Shibaura Institute of Technology
Tokyo, Japan
nb22509@shibaura-it.ac.jp

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

Midori Sugaya
Shibaura Institute of Technology
Tokyo, Japan
doly@shibaura-it.ac.jp

ABSTRACT

In recent years, large-scale computer systems have been constructed to meet the growing demand for various information services. These systems often centrally manage OS images in network storage for efficiency. However, booting the OS image on the network storage has problems with storage performance and startup time. We propose an OS booting mechanism utilizing SmartNIC. In our proposed methods, the SmartNIC concurrently copies the OS image to local storage while processing I/O requests from the booting OS. Our approach leverages SmartNIC capabilities to optimize system boot and I/O performance.

INTRODUCTION

In recent years, the demand for various information services has led to the construction of large-scale computer systems. These systems manage OS images centrally using network storage for efficient management. When booting the OS from the OS image on the network storage, each computer must either attach the OS image via the network or copy it from the network storage to the local storage. Both methods have drawbacks: attaching network storage continuously degrades storage performance, while copying the entire image results in startup delays.

Previous research [1] addressed these issues with a hypervisor. In the proposed system, the OS boots from the network storage image while copying it to local storage. This process was made possible by I/O virtualization. Once copied, the hypervisor is removed, allowing direct access to local storage. However, this approach results in low I/O performance during the initial boot due to hypervisor intervention. This issue has become more significant with modern high-speed storage devices like NVMe SSDs.

To resolve this, we propose a SmartNIC-based method to boot the OS image on the network storage. This approach

aims to prevent I/O performance degradation by taking advantage of SmartNIC instead of virtualization techniques. While booting the OS, it copies the OS to local storage concurrently, as shown in Figure 1. The SmartNIC provides an NVMe SSD interface for the host's firmware, bootloader, and booting OS during initial startup. It processes I/O requests from the host while copying the OS image from the network storage to the local NVMe storage.

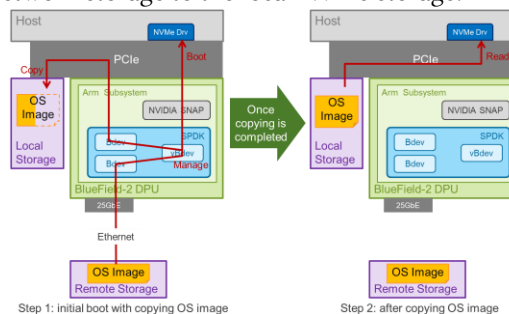


Figure 1: Proposed implementation of OS image deployment system

Our method requires SmartNIC to integrate three key functionalities: (1) providing a storage device interface accessible by the running OS, (2) reading and writing to local and network storage without OS intervention, and (3) dynamically switching between local and network storage based on I/O requests. Our presentation shows the design details and our current research status.

ACKNOWLEDGMENTS

This research was supported by JST CREST JPMJCR19K1.

REFERENCES

- [1] Y. Omote, et al. Improving Agility and Elasticity in Bare-metal Clouds. In Proc. of ASPLOS '15. New York, USA, 2015. 145–159. <https://doi.org/10.1145/2694344.269434>.