Low-Overhead Privilege Escalation Detection Mechanism for High-Performance Computers Using SmartNIC

Wassapon Watanakeesuntorn, Kohei Taniguchi, Junya Yamamoto,
Keichi Takahashi, Hirotake Abe, Arata Endo, Chonho Lee, Susumu Date

THE UNIVERSITY OF OSAKA
D3 CENTER DIGITAL DESIGN, DATABILITY, AND DECISION INTELLIGENCE

Background and Motivation

- HPC systems enforce job isolation using user authentication and process-level privilege management.
- Privilege escalation attacks allow attackers to gain administrative privileges, potentially leading to severe cybersecurity threats, including lateral movement within the system.
- Continuously verifying user privileges with Effective User ID (EUID) introduces monitoring overhead on the compute node.

Propose a privilege escalation detection mechanism using a Smart Network Interface Card (SmartNIC)

- NVIDIA BlueField-2 DPU is used as SmartNIC
- Offload memory monitoring from the CPU and detect privilege escalation caused by malicious jobs.
- Suspend malicious job and disconnect the compute node from the network

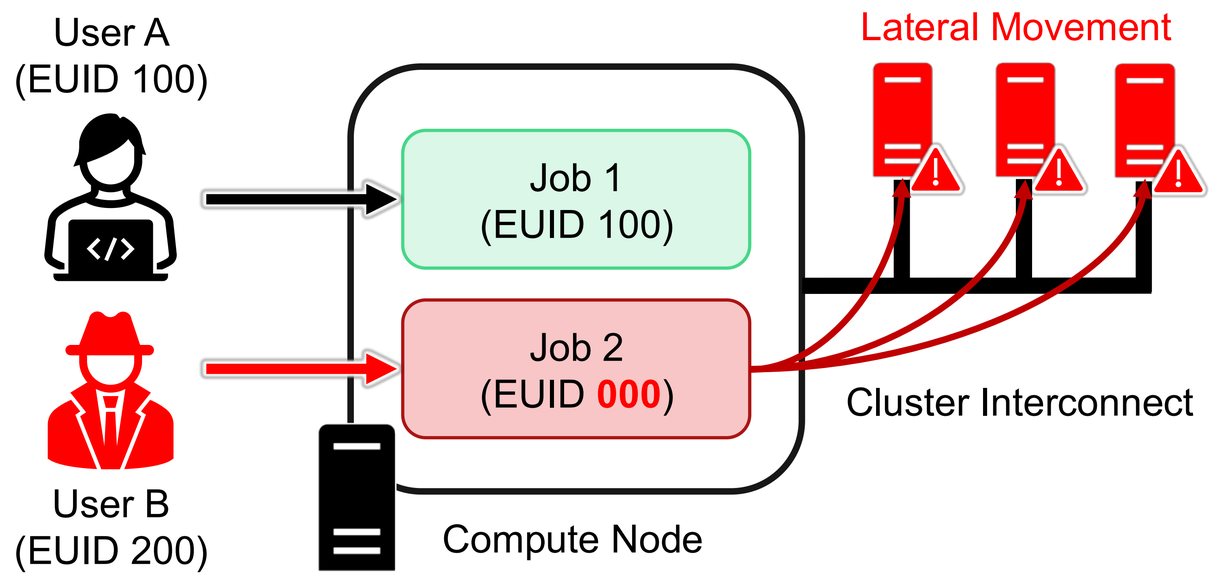

Implementation

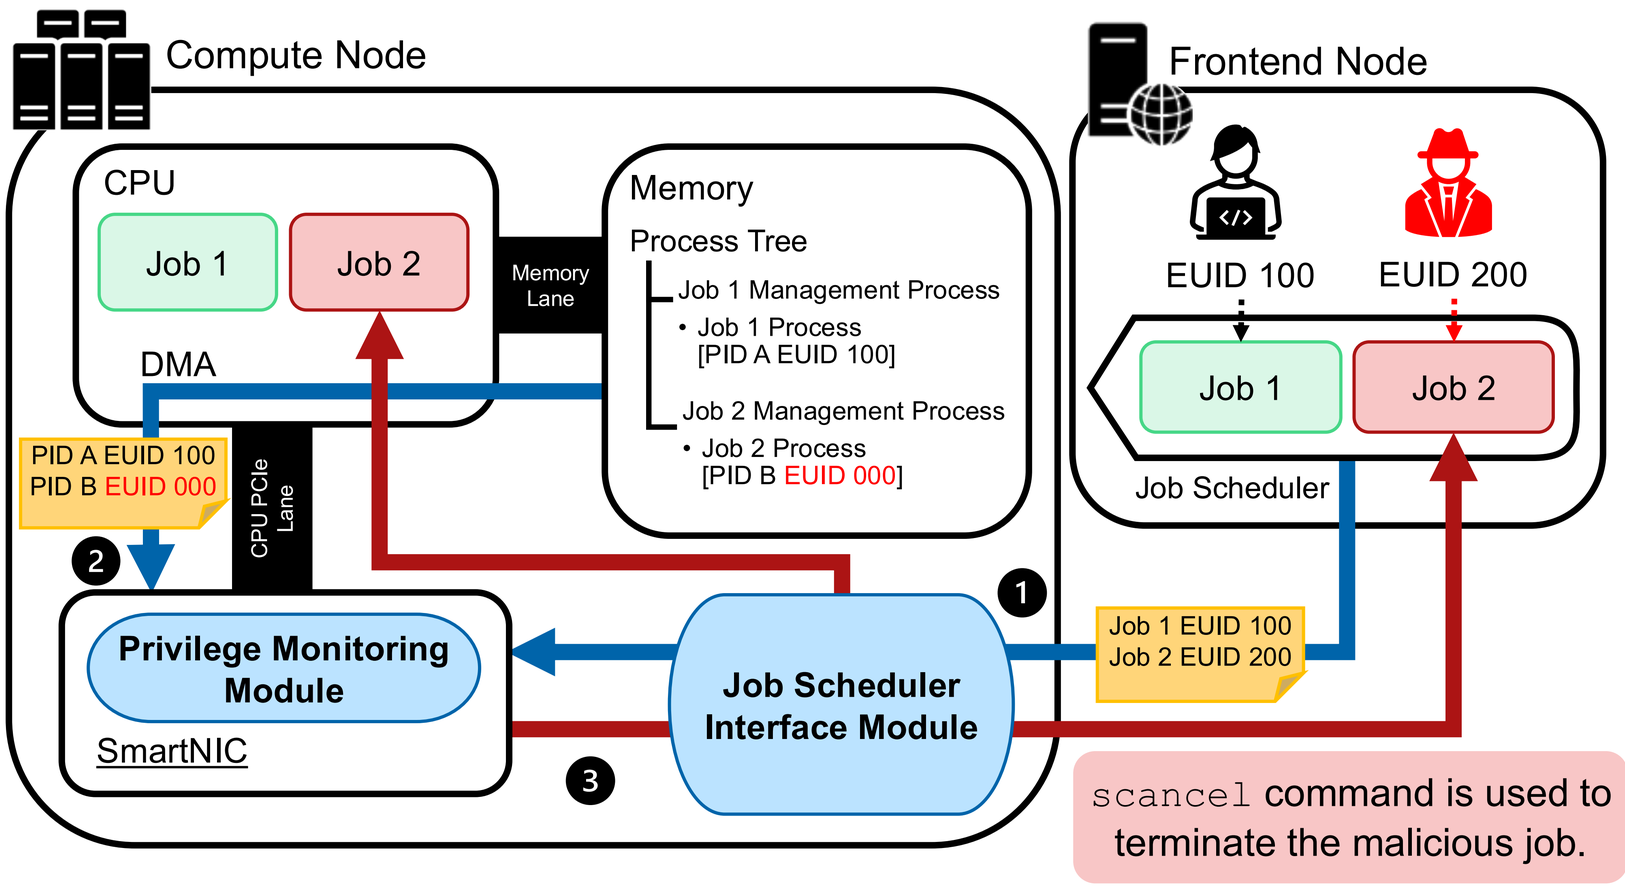

1. Job Scheduler Interface Module collects metadata from the Slurm scheduler, including the job ID and the EUID of the submitting user.
2. Privilege Monitoring Module, running on the SmartNIC, accesses host memory via DMA to retrieve the EUIDs of running processes.
3. If the SmartNIC detects an EUID mismatch between the user information and the process, it terminates the malicious job and disconnects compute node from the network.

Evaluation and Results

Evaluation Environment

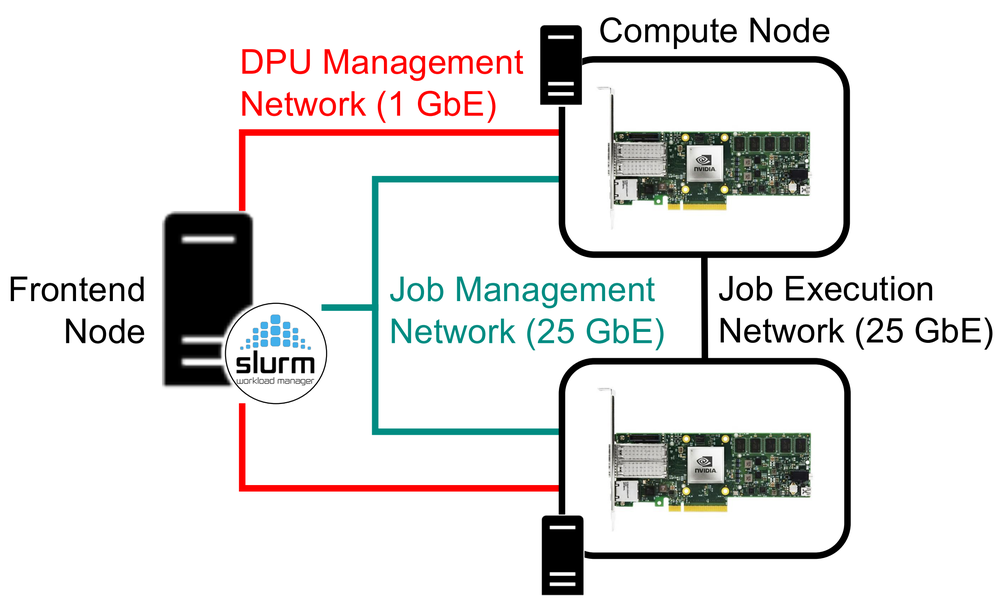

SmartNIC can detect privilege escalation within 141 ms and introduced only a negligible 0.62% overhead.

Breakdown of Detection Mechanism Runtime

The processes are arranged as a chain-like tree, in which each process has at most one child, and the leaf process performs the privilege escalation.

	Privilege Escalation Detection	Job Termination	Network Isolation
1 Process	127 ms	142 ms	263 ms
100 Processes	141 ms	128 ms	262 ms

Overhead Evaluation with Himeno Benchmark

- Himeno Benchmark is a memory-intensive benchmark.
- This evaluation aims to assess the trade-off between SmartNIC-based monitoring with DMA and host-based monitoring using one CPU core.

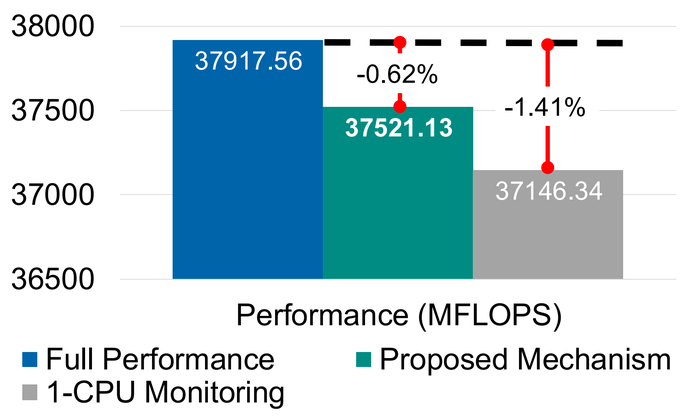

Acknowledgement: This poster is based on results obtained from the project, “Research and Development Project of the Enhanced infrastructures for Post-5G Information and Communication Systems” (JPNP20017), commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

CloudEdge Fusion