Bolted: Tenant Controlled Security for Bare Metal Clouds

A. Mosayyebzadeh[1,2], A. Mohan[1,3], S. Tikale[1,2], M. Abdi[1,3], A. Raza[1,2], N. Schear[5], T. Hudson[4], C. Munson[5], L. Rudolph[4], G. Cooperman[3], P. Desnoyers[1,3], O. Krieger[1,2]

Introduction:

Public clouds can help tenants meet the elastic nature of their workloads but they do have huge security problems. Their enormous computing base i.e., the entire virtualization stack and cloud orchestration software, is an easy target for attackers. Even if we assume there is no external threat, the tenant still has to trust the provider who owns the host OS on the machines. Such a model drives away many security conscious clients.

Goals:

- In order to reduce the trusted computing base, we want to hand out bare metal nodes to tenants.
- We want the tenants to be able to easily get more nodes as their demand increases and then free up nodes when the demand decreases.
- We want tenants to carve out secure isolated enclaves with those bare metal nodes.

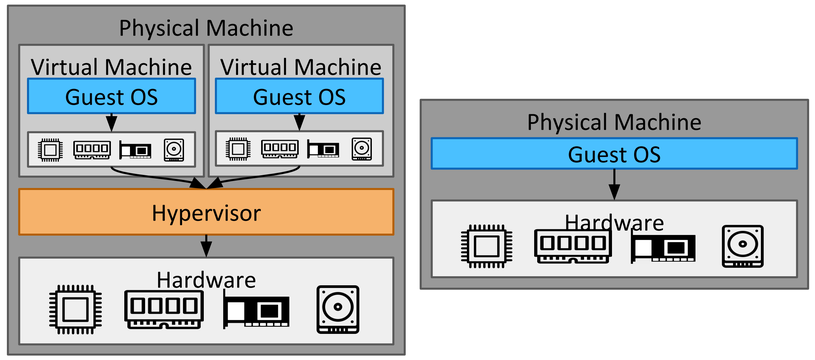

Bare metal nodes reduce trusted computing base

Challenges:

Bare metal nodes provide performance and privacy advantages over virtual machines, but the direct hardware access they give opens up new attack vectors that must be addressed.

- Malicious tenants can flash the firmware and leave persistent malware which can then infect tenants who use that node later
- If the cloud vendor somehow addresses the first challenge, tenants still have to trust the provider's ability to do so. This puts the trust back in the provider which we don't want.

Proposed System:

We have developed different micro services to realize the goals and address the challenges.

- To get isolated networks and allocate bare metal nodes, we use **HIL**
- To fast provision bare metal nodes, we use **BMI**
- To measure and attest firmware to ensure it has not been tampered with, we use **Keylime**

All these services are described below.

Services:

Hardware Isolation Layer (HIL):

- Exokernel for managing physical hardware
- Enables tenants to create isolated networks of machines
- Guarantees isolation in a multi-tenant environment

Bare Metal Imaging (BMI):

- Provides scalable system for network booting application disk images on physical hardware
- Reliably stores disk images inside CEPH object store

Keylime:

- Provides cloud scale TPM attestation infrastructure
- Automatically bootstraps an initial secret on nodes

Architecture:

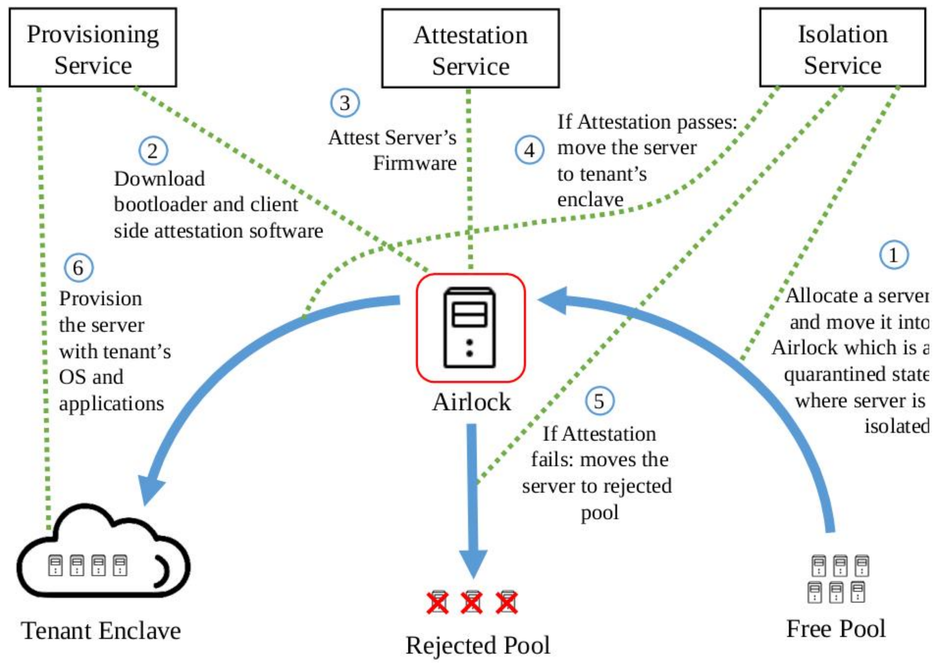

Experiments & Results:

- Provisioning Time

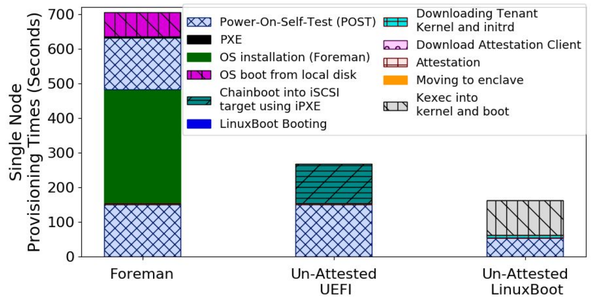

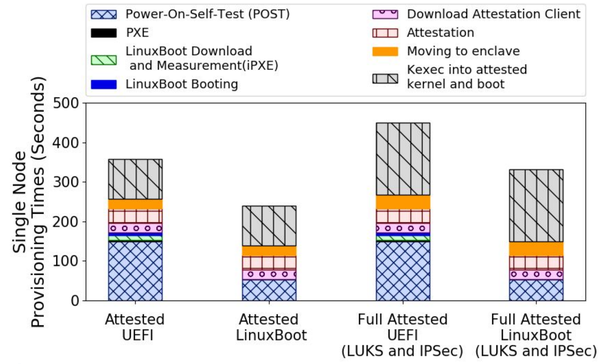

- Runtime:

1. MPI

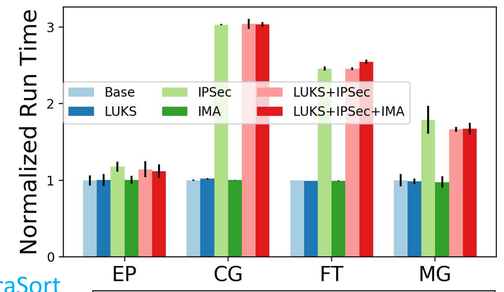

2. Spark TeraSort

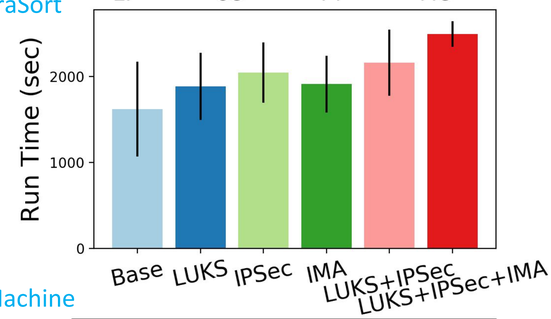

3. Virtual Machine

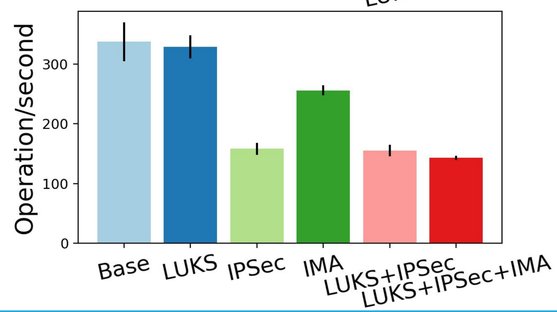

1. Mass Open Cloud 2. Boston University 4. Northeastern University 4. TwoSigma 5. MIT Lincoln Laboratory