Categorizing container escape methodologies in multi-tenant environments

Rik Janssen

rik.janssen@os3.nl

Research Project 1
MSc Security and Network Engineering (SNE/OS3)
University of Amsterdam

Supervisor:
Max Hovens
KPMG Cyber Security Netherlands

February 6, 2018

Introduction

Operating-system-level virtualization

- Docker, LXC, FreeBSD Jails, etc.
- Container escapes
- Host system takeover
- Data leakage
- Multi-tenant environments

Introduction

Research Question

- How to systematically categorize vulnerabilities relating to multi-tenant environments that make use of operating-system-level virtualization?

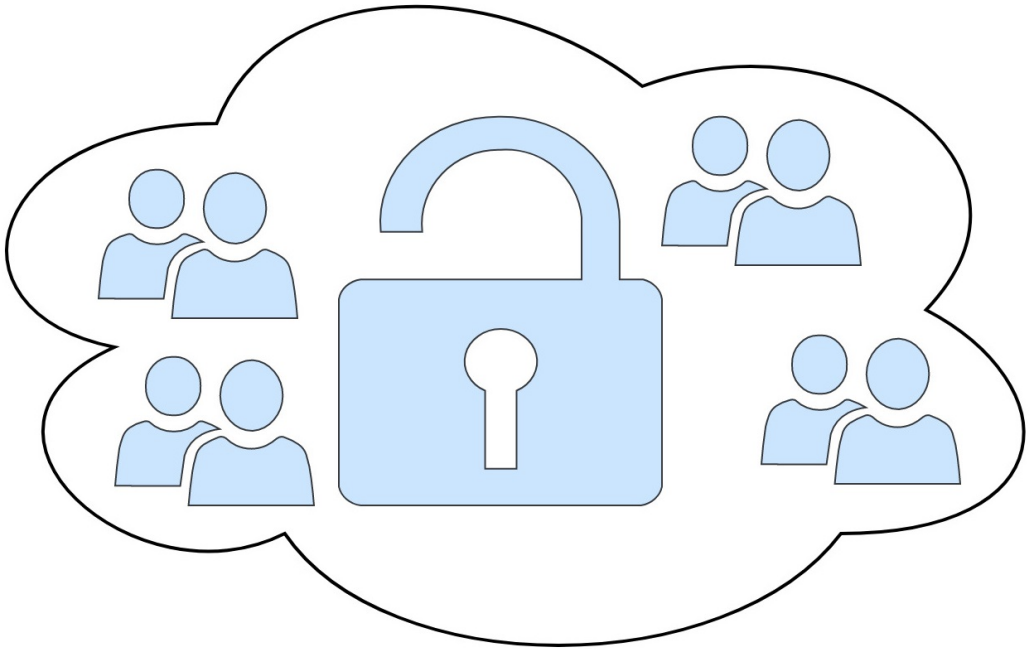

Figure 1: Public cloud

Related Work

Comparative study [1] of OS-level virtualization systems:

- defines attacker model;
- defines security requirements for isolation techniques;
- LXC, FreeBSD Jails, Solaris Zones, etc., but not Docker.

Later study [2] extends on previous study:

- Linux only (Docker);
- evaluates and extends security requirements;
- correlates CVEs to requirements.

Both focus on single container host.

Approach

Literature study

- Linux and Kubernetes
- Breakdown of components
- Mapping CVEs
- Extending requirements
- Affected workloads
- Mitigation techniques

Results - Architecture overview

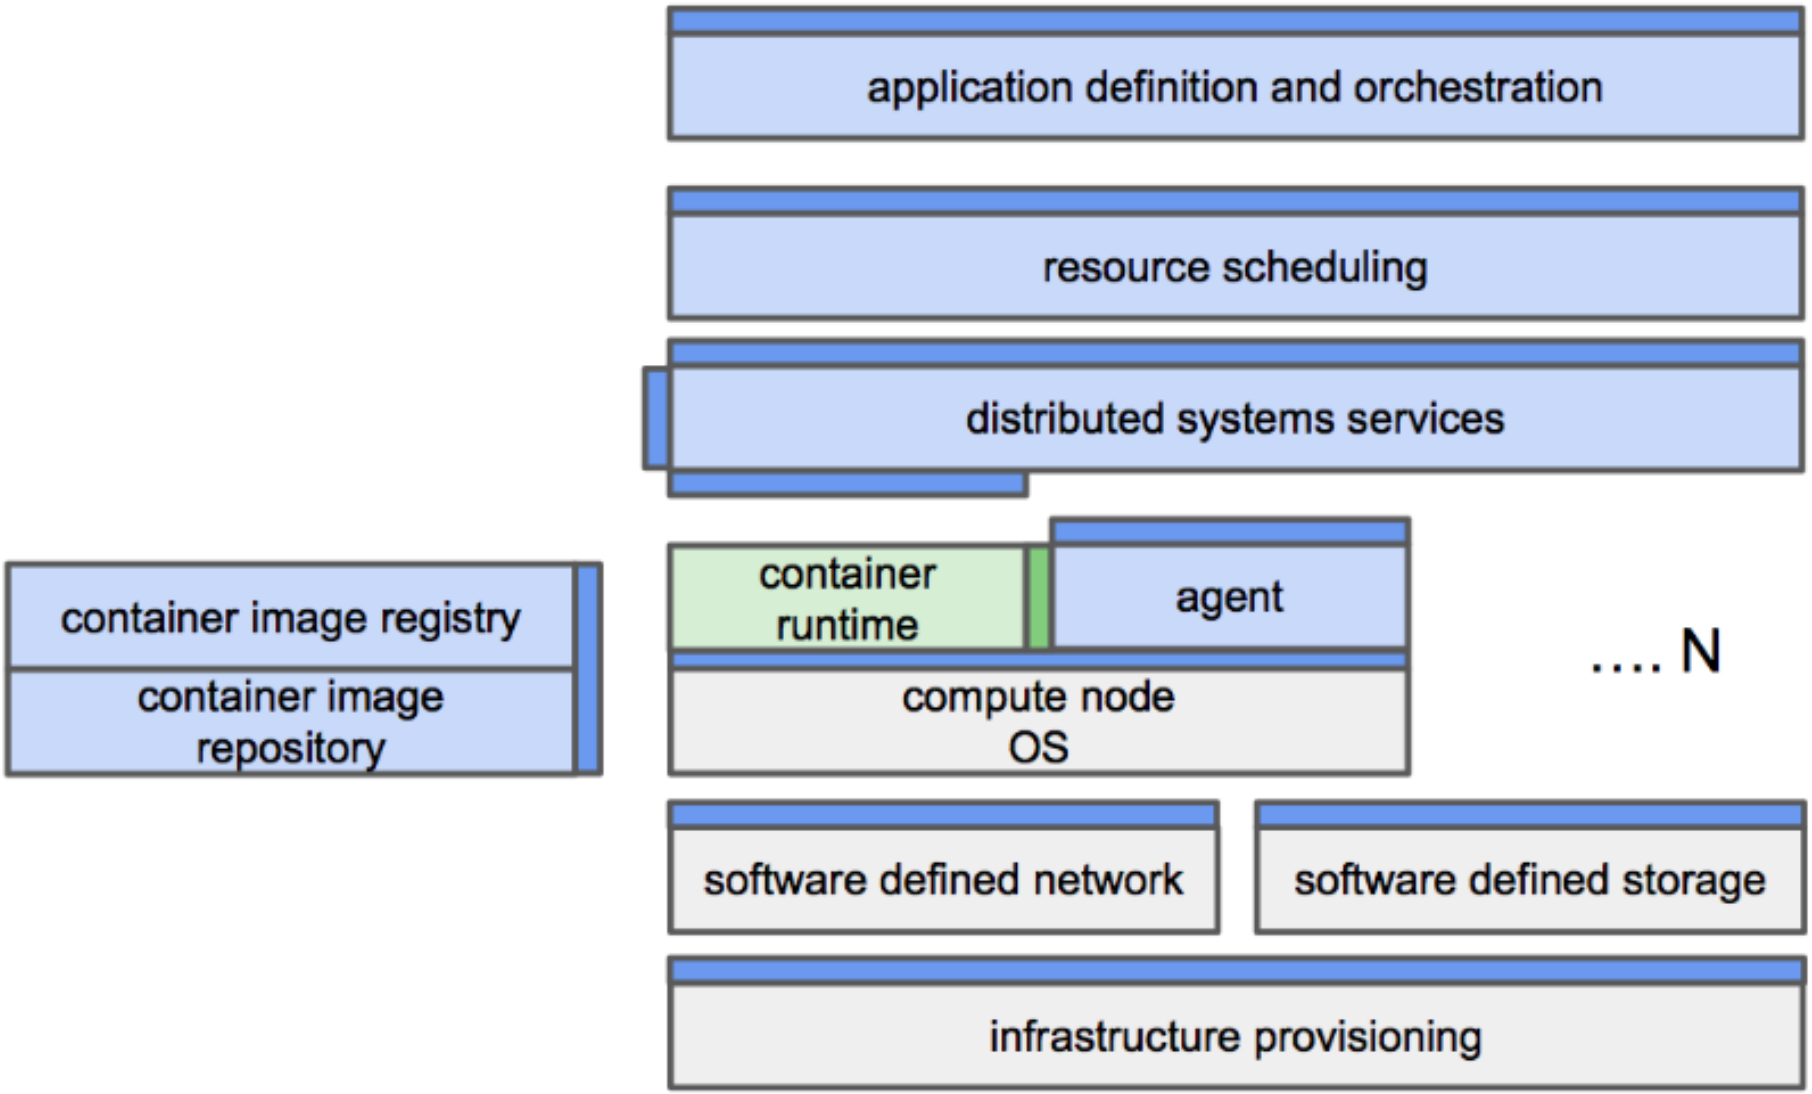

Figure 2: Cloud Native Computing Foundation (CNCF) Scope [3]

Results - Architecture overview

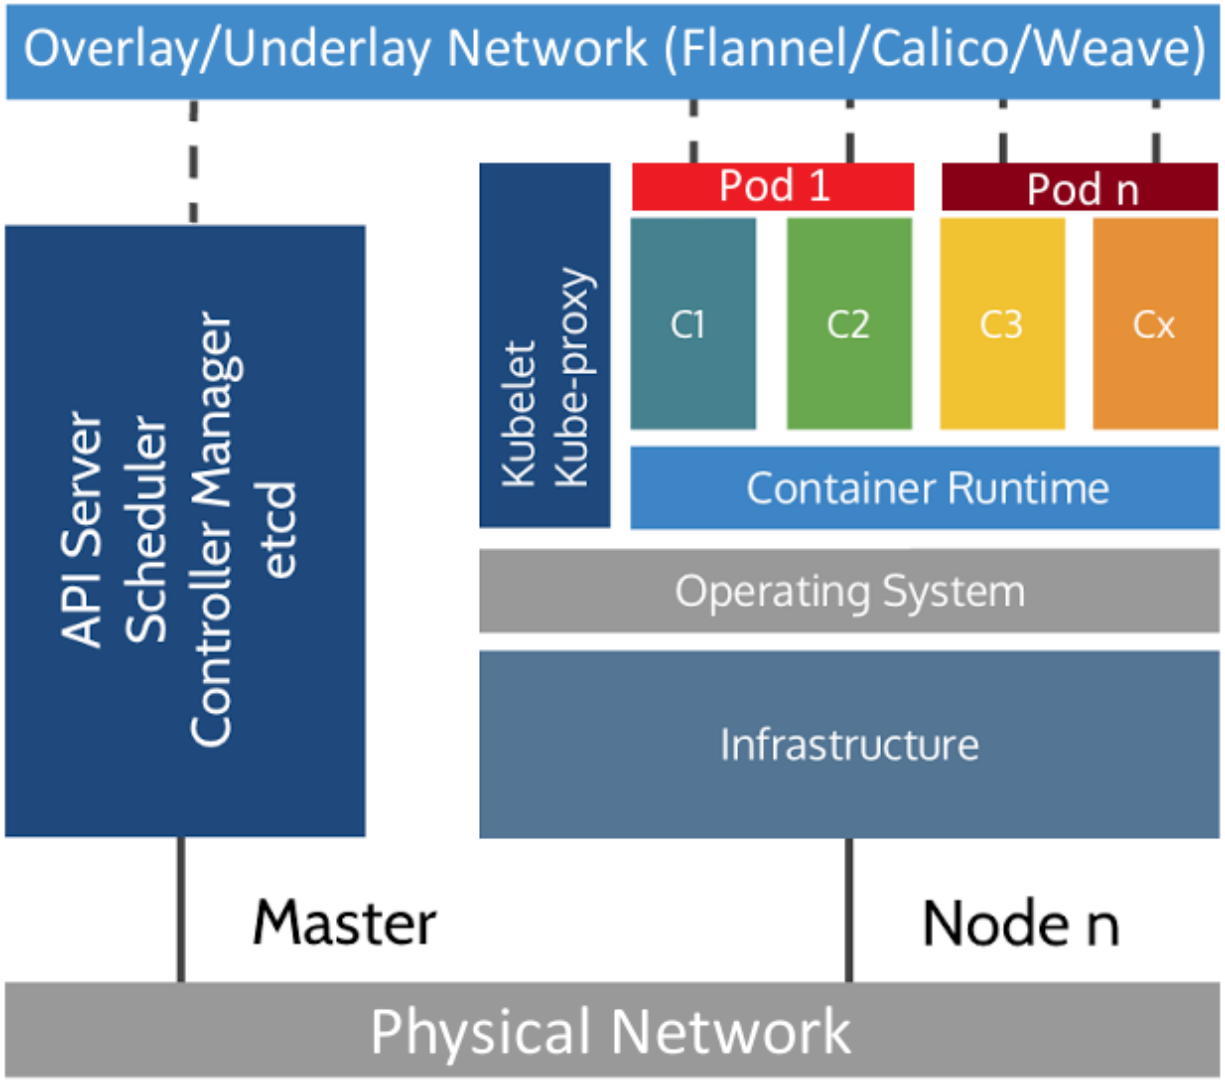

Figure 3: Kubernetes [4]

Results - Security model

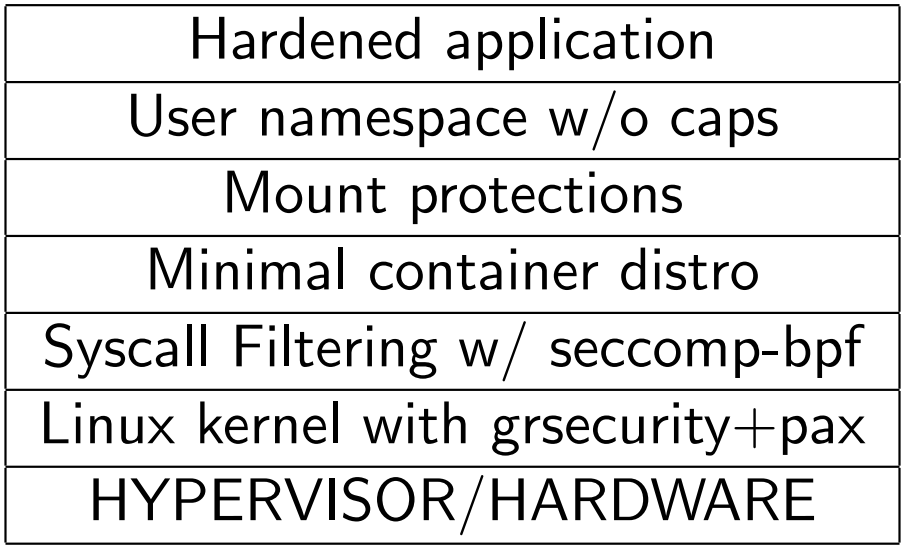

Figure 4: NCC Group's security model [5]

Results - Attack surface via linking

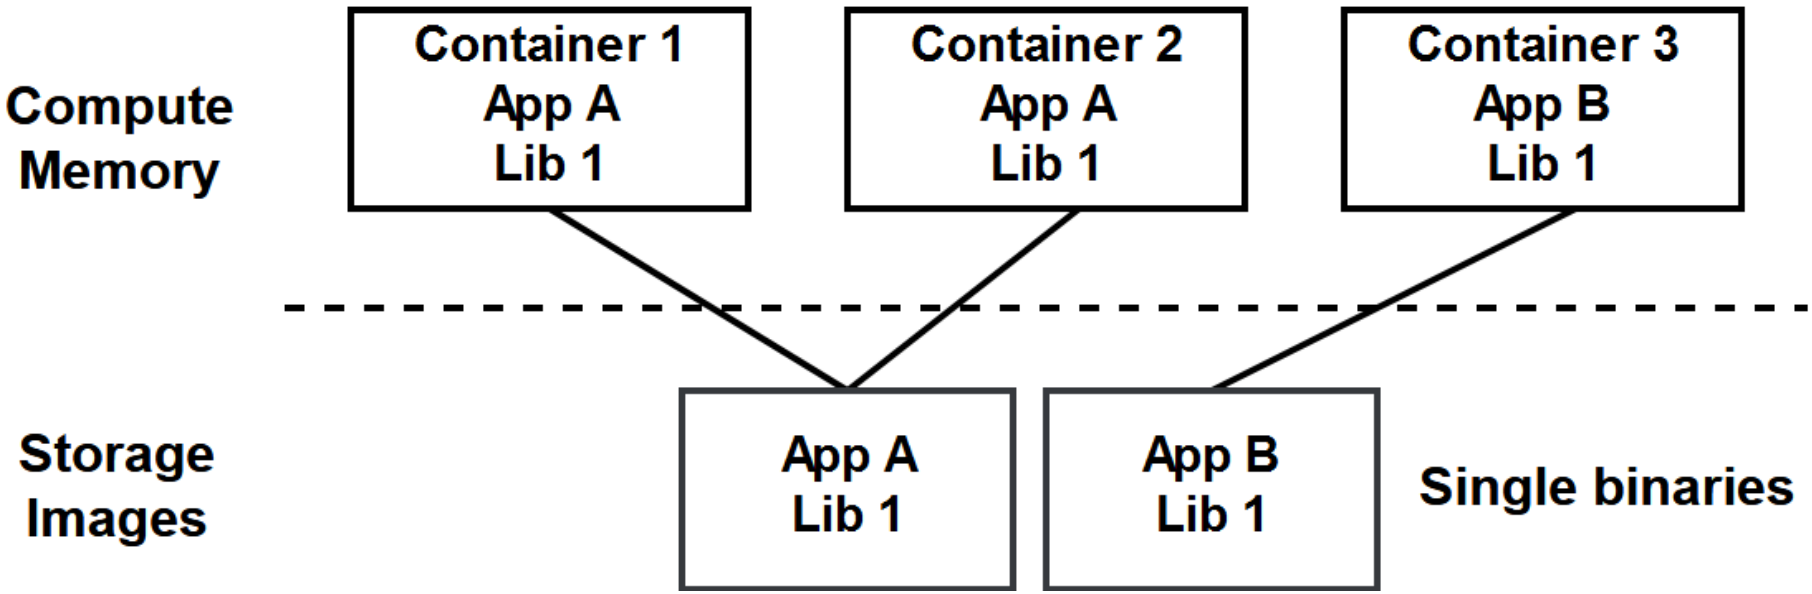

Figure 5: Static linking

Discussion - Attack surface via linking

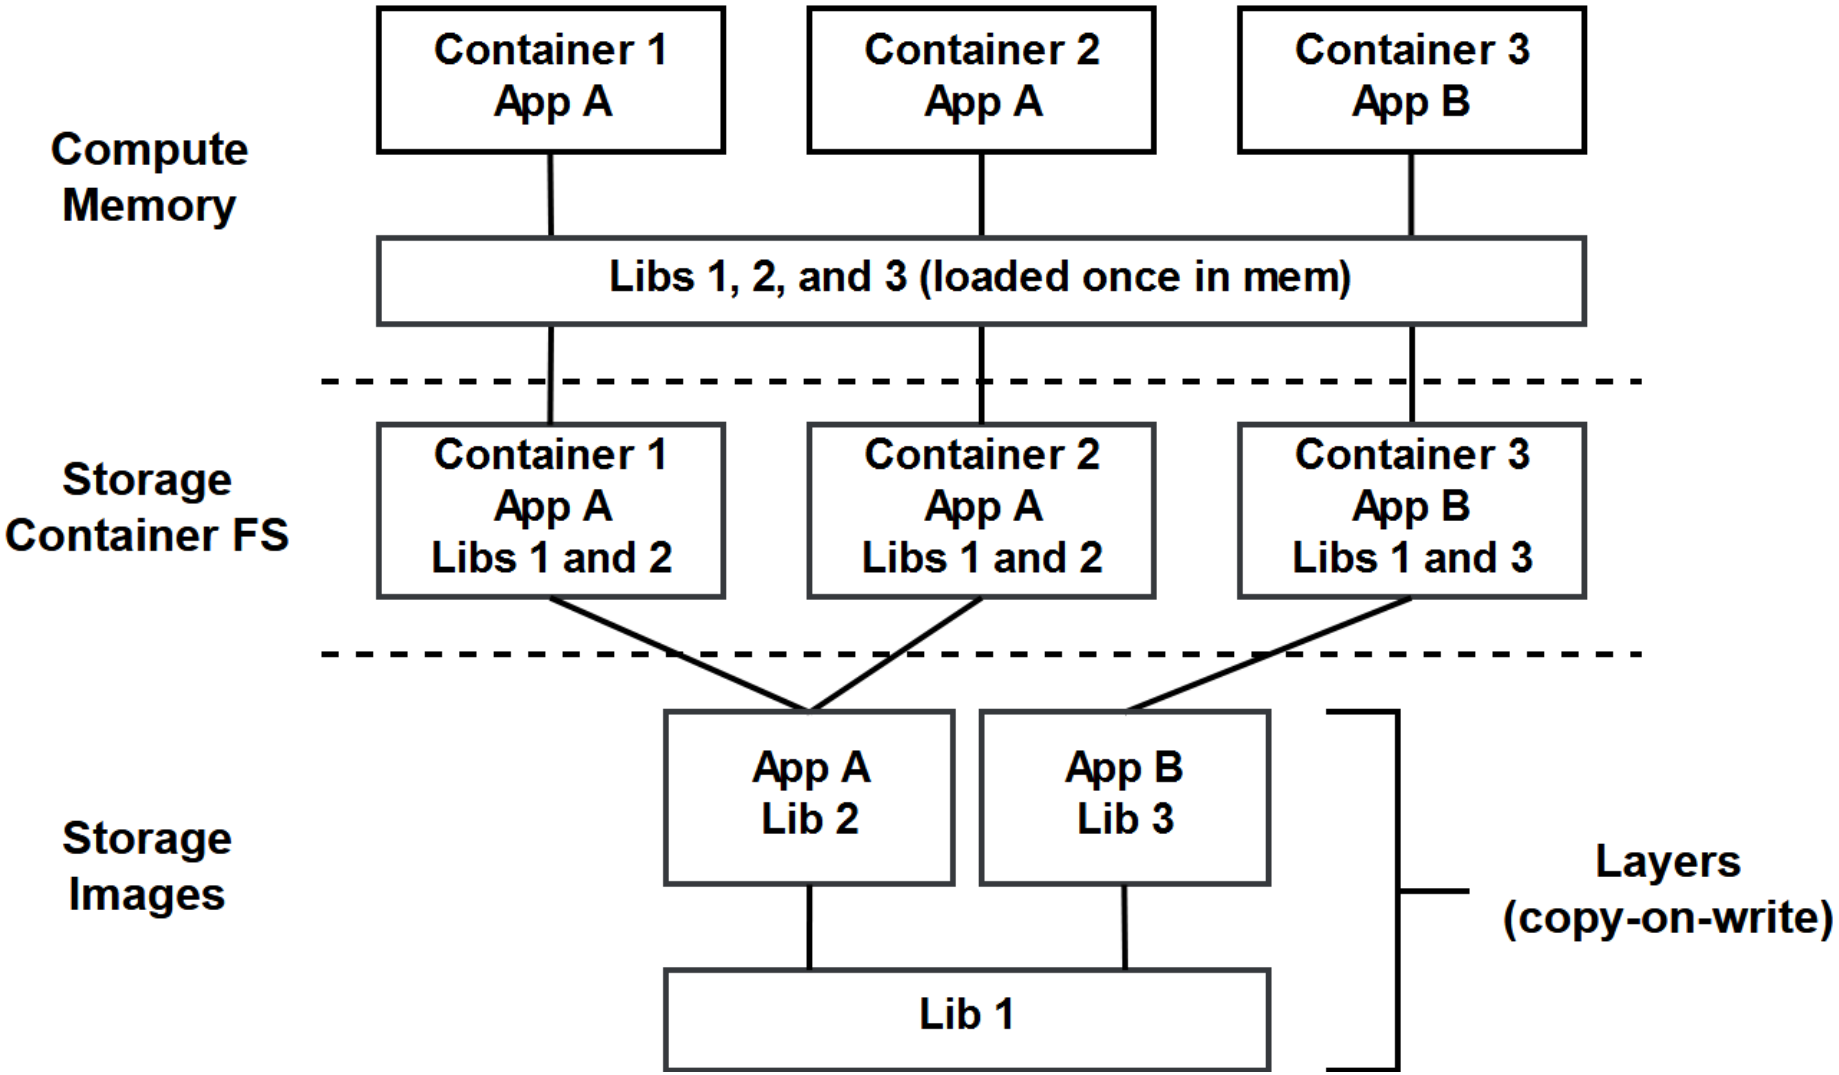

Figure 6: Dynamic linking

Discussion - CVEs

Etcd

- Distributed key-value store
- Single source of configuration truth
- Authorization Modules (e.g. 'ABAC')
- One CVE found (CVE-2015-5305)

OpenShift

- Red Hat
- Integrated platform (Atomic OS + Docker + Kubernetes)

Results - CVE-2017-5754: 'Meltdown'

Side-channel attack [6] utilizing out-of-order execution and CPU cache timings to read **all** physical memory.

- Hardware-level vulnerability
- Possible from any user space process
- Containers specifically mentioned in regards to information leakage
- Intel covered in paper
- Mitigate using kernel page-table isolation patches

Conclusion

Research question

How to systematically categorize vulnerabilities relating to multi-tenant environments that make use of operating-system-level virtualization?

Conclusion

Research question

How to systematically categorize vulnerabilities relating to multi-tenant environments that make use of operating-system-level virtualization?

Answer (preliminary)

- Application architecture is critical due to tight coupling with the underlying infrastructure.
- Security boundary is moved to the infrastructure (i.e. hardware) level.
- Utilize integrated approach (i.e. don't use upstream).

Future work

- Compliance correlation (e.g. ISO, HIPAA, NEN)
- Other OS-level virtualization platforms (e.g. Windows, *BSD, Nomad, Mesos, etc.)
- Other virtualization techniques (e.g. full and/or para)
- Automated workload classification and orchestration
- Istio platform (alpha)

References I

[1] E. Reshetova, J. Karhunen, T. Nyman, and N. Asokan, "Security of os-level virtualization technologies," in *Nordic Conference on Secure IT Systems*, pp. 77–93, Springer, 2014.

[2] S. Laurén, M. R. Memarian, M. Conti, and V. Leppänen, "Analysis of security in modern container platforms," in *Research Advances in Cloud Computing*, pp. 351–369, Springer, 2017.

[3] The Linux Foundation, "Cloud Native Computing Foundation (CNCF) Charter," 2017.

[Online]. Available: `https://www.cncf.io/about/charter`. [Accessed: Jan. 15, 2018].

[4] I. Gunaratne, "A Reference Architecture for Deploying WSO2 Middleware on Kubernetes," Jan. 2017.

[Online]. Available: `https://medium.com/containermind/a-reference-architecture-for-deploying-wso2-middleware-on-kuberr` [Accessed: Jan. 17, 2018].

References II

[5] A. Grattafiori, "Def con 23 - aaron grattafiori - linux containers: Future or fantasy? - 101 track." Youtube. 2015. Available: https://www.youtube.com/watch?v=iN6QbszB1R8 [Accessed: Jan. 3, 2018].

[6] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg, "Meltdown." arXiv:1801.01207, Jan. 2018.