Nomad: Mitigating Arbitrary Cloud Side Channels via Provider-Assisted Migration

Soo-Jin Moon*, Vyas Sekar*, Michael K. Reiter ¶ *Carnegie Mellon University, ¶UNC - Chapel Hill

Contact: soojinm@andrew.cmu.edu

Project Silver

Motivation : Cross-VM side channels in clouds

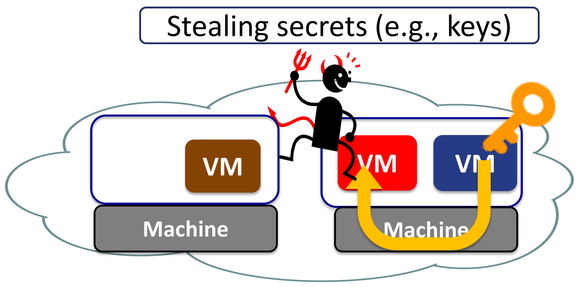

- Growing threat in multi-tenant clouds
- Any tenant is a potential threat
- Can exploit many different vectors (L2/L3 cache, storage, memory)

  e.g. ,Y. Zhang et al., CCS'12; T. Ristenpart el al., CCS'09; F. Liu et al.,Oakland15, and several more!

Current defenses: 1) Vector-specific 2) Need significant changes

Goals & Insights

#1: General
Applicable to broad spectrum of side channels → Minimize co-residency

#2: Immediately deployable
Minimal modifications to hardware, software & apps → Use VM Migration

Idea: Migration as a Provider-assisted Defense

Nomad Overview

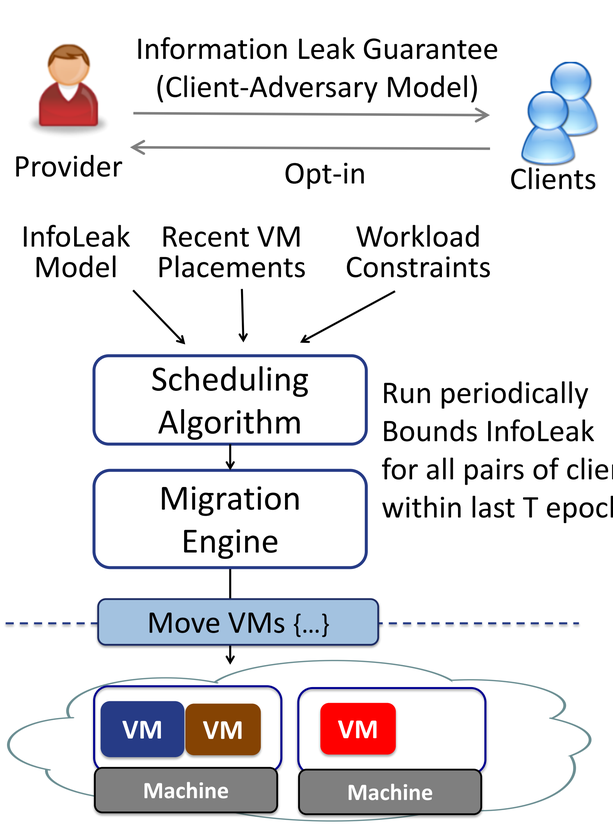

Run periodically
Bounds InfoLeak for all pairs of clients within last T epochs

Challenges & Solutions

C1: Logic
Formalize InfoLeak due to co-residency

Client VMs: Replicated? (R vs. NR)
Adversary VMs: Collaborating? (C vs. NC)

Adversary Dimension \ Client Dimension		
	<NR,NC> Least InfoLeak	<R,NC>
	<NR,C>	<R,C> Most InfoLeak

C2: Scalability
e.g., Can EC2 run this?

Scalable Greedy Algorithm
1) Prune search space
2) Incremental computation
3) Intra-epoch lazy evaluation

C3: Deployability:
Minimal changes?

~200 LOC of modifications in OpenStack

Key Results

Close to optimal InfoLeak

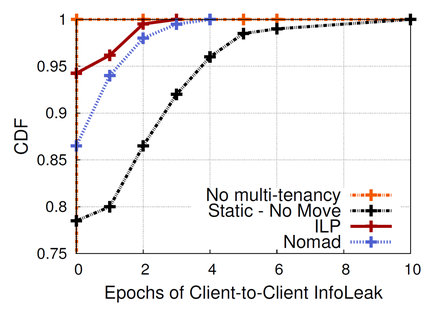

Scalable to large deployments

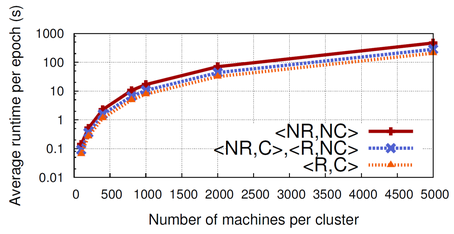

Cluster size of 40: >1 day to solve for ILP (Integer linear programming)

Minimal performance impact for cloud workloads

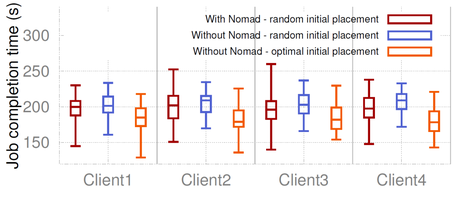

$$Norm.Throughput = \frac{T_{w/o} - T_w}{T_{w/o}} x\ 100$$

- 5th Norm. Throughput: **1.8%**
- 95th Norm. Throughput: **~0%**