Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds

Thomas Ristenpart, Eran Tromer, Hovav Shacham, Stefan Savage

CCS'09

Presented by Bingzhe Liu

Achievement

- CCS'09

- Media Coverage
  - MIT Technology Review
  - The New York Times
  - Network World
  - …

- Citations
  - Cited by 1789 papers so far.
  - *Second most-cited security paper* of those published between 2008 and 2013.

Highlight

- Prove the existence of confidentiality breach within EC2
  - First work on cloud cartography.
  - Launch the attack against commercially available real cloud (Amazon EC2).
  - Claims up to 40% success in co-residence with target VM.
- Suggest mitigations

Cloud computing -- A new business model

- On-demand computing outsourcing
  - Large scale computer lease
  - Charge for the actual computation utilization

- Famous cloud computing platform in 2009
  - Amazon’s EC2 (Elastic Compute Cloud)
  - Microsoft’s Azure Service Platform
  - Rackspace’s Mosso

New threats in cloud

- Traditional system security mostly means keeping bad guys out.
  - The attacker needs to either compromise the auth/access control system, or impersonate existing users

- But clouds allow **co-residence**
  - Multiple independent users share the same physical infrastructure.
  - An attacker can legitimately be in the same physical machine as the target

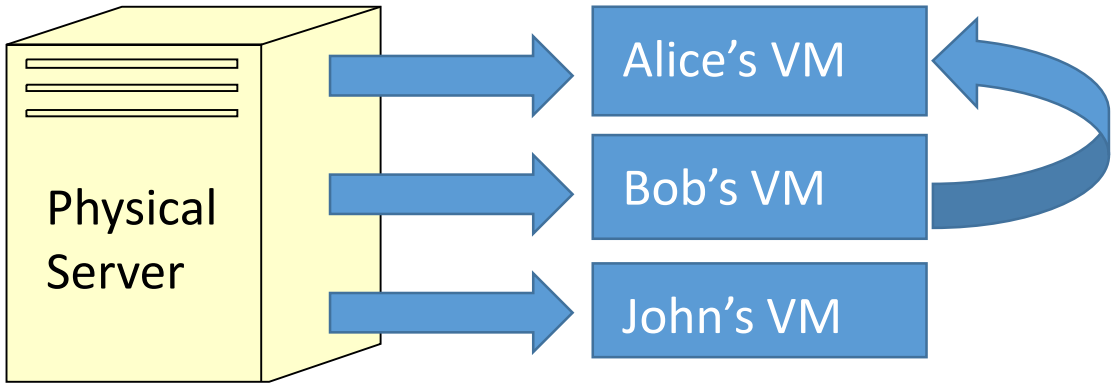

Case study: Amazon EC2

- Xen hypervisor, called Domain0, is used to manage guest images, physical resource provisioning, and access control rights.
- Dom0 routes packages and reports itself as a first hop.
- Consists of 2 regions (United States and Europe), each have 3 availability zones, 5 Linux instance types.
- Instances have a one-to-one mapping of internal IP addresses and external IP addresses, which are static.

- Note: Some of them are outdated.

Threat Model

- Attacker Model
  - Cloud infrastructure provider is trustworthy
  - Attacker is a malicious third party who can legitimately use cloud provider's service
- Targets
  - Known hosted service
  - Particular victim service that attacker interested

Challenges for attacker

1. Placement
   1) How to find out **WHERE** the target is located
   2) How to **CO-RESIDENT** with the target in the same physical host
2. Extraction
   3) How to **EXTEACT** confidential information via a cross-VM attack.

1. How to find out **WHERE** the target is located

Cloud Cartography

1. Figuring out Internal Addresses
   - WHOIS to survey external IPs associated with EC2
     - Three distinct IPs prefixed with /17, /18, /19
     - Total 57344 IPs
     - 14054 IPs with open port 80, 443
   - DNS lookup *within* EC2 mapped external to internal IPs

Cloud Cartography (cont'd)

2. Mapping the Internal IPs against zones
   - Different availability zones correspond to different /16 internal IP address ranges.

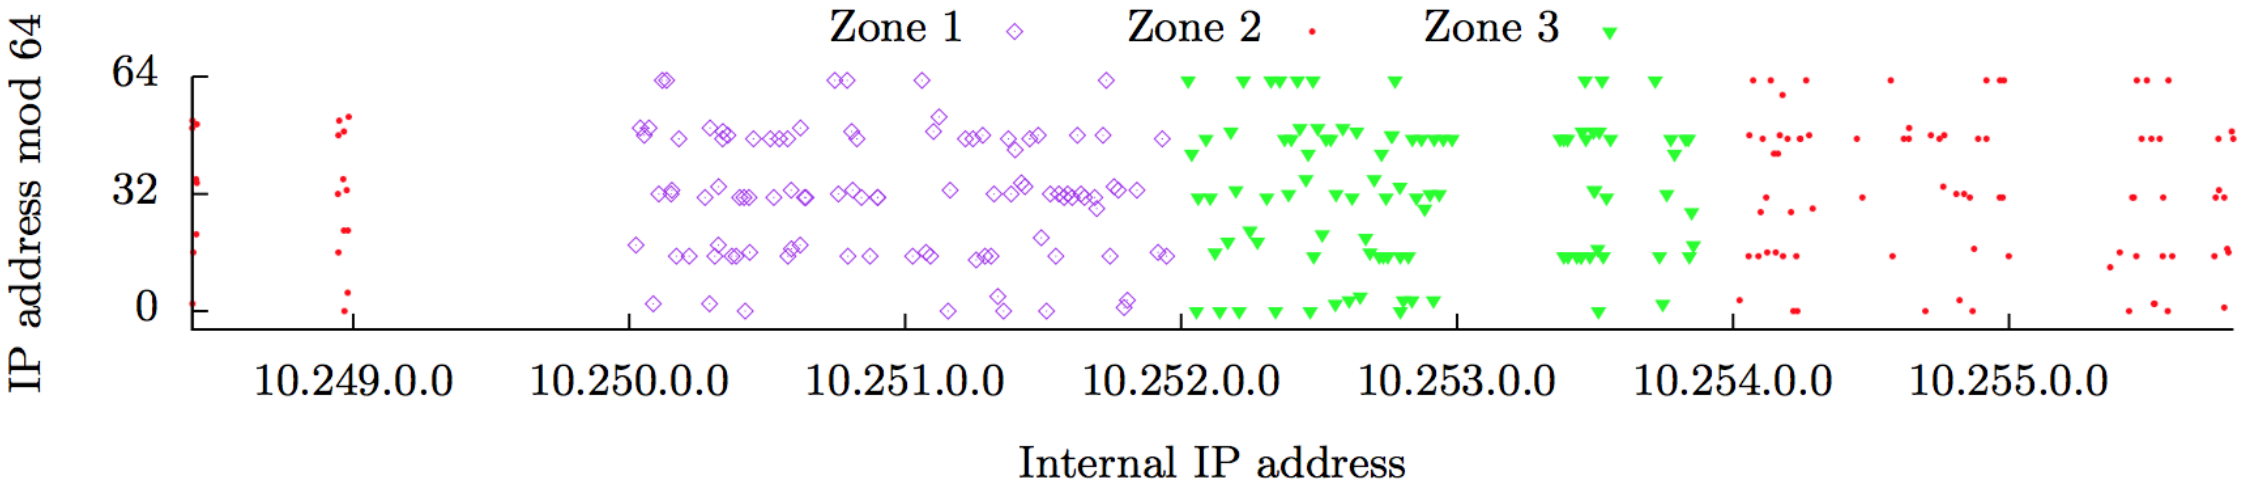

Cloud Cartography (cont'd)

3. Mapping the internal IPs in Zone 3 against instance types
   - Same instance types correspond loosely with similar IP address range regions.

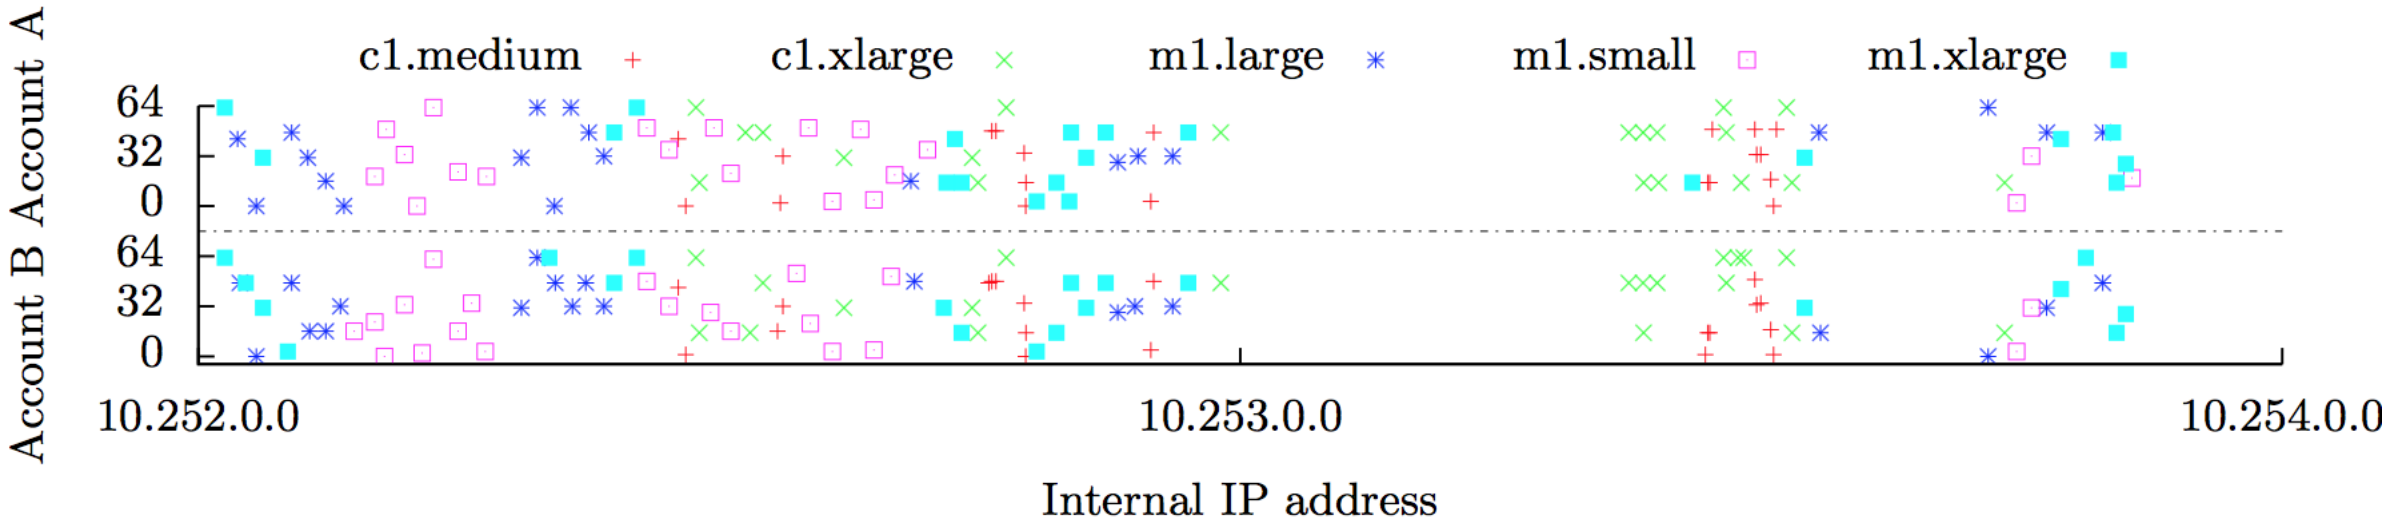

Mitigations

- Mapping
  - Use a randomized scheme to allocate IP addresses instead of static
  - Block some scanning tools/activities (nmap,traceroute)

2. How to **CO-RESIDENT** with the target in the same physical host

Determine co-residence

- Strategy1: Network-based co-residence checks
  - Instances are likely co-residence if they have:
    - Matching Dom0 IP address
    - Small packet round-trip times(RTT)
    - Numerically close internal IP addresses (within 7)
  - Verified via a hard-disk-based covert channel
  - Conclusion of test: Effective false positive rate of ZERO for the co-resident checks.

- Strategy2: Side-channels (in section 8)

VM Placement in EC2

- Strategy 1: Brute-forcing placement
  - A success rate of 8.4%
- Strategy 2: Abusing placement locality
  - With the observation of **strong placement locality**: two instances running at roughly the same time are often assigned to the same host
  - How:
    - Attacker knows when the target instances will be launched
    - Inference availability zone and instance type from its IP
    - **Instance flooding** immediately following launch of instance by launch many instances simultaneously.
  - Achieves a success rate of 40%

Mitigations

- Co-residence checks
  - Prevent identification of dom0/hypervisor
- Co-residence prevention
  - Not allow co-residence at all
    - Beneficial for cloud users, not efficient for cloud providers
  - Offload the placement choice to users

3. How to **EXTEACT** confidential information via a cross-VM attack.

Cross-VM information leakage

- Strategy: Cache-based side channel attack
  - Load measurement
    - Prime+Trigger-Probe
    - Differential coding: mitigating nosie by “odd sets” & “even sets”
- Attack:
  - Estimate victim's work load
    - Cache
    - Network Traffic
  - Stealing cryptographic keys by keystroke timing attack
  - Denial of Service

Mitigations

- Side-channel attack
  - Avoid co-residence in the same physical host

Amazon’s response in 2009

- Amazon downplays report highlighting vulnerabilities in its cloud service
- "The side channel techniques presented are based on testing results from a carefully controlled lab environment with configurations that do not match the actual Amazon EC2 environment."
- "As the researchers point out, there are a number of factors that would make such an attack significantly more difficult in practice."

*http://www.techworld.com.au/article/324189/amazon_downplays_report_highlighting_vulnerabilities_its_cloud_service, Oct 29th, 2009

Pros

- They do reveal vulnerabilities in EC2, like easily retrieve the internal IP address, and mappings
- Use simple tool to launch practical attacks on EC2, which could be done by anyone.
  - Then dedicated attacker may launch more accounts and use more time to accurately attack target
- Introduce the novel threat model: Data and software are not the only assets worth protecting, activity patterns also need to be protected.

Acknowledgement

- Some slides/thoughts borrowed from Edward Wu & Bo Sun.
- Ristenpart T, Tromer E, Shacham H, et al. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds[C]//Proceedings of the 16th ACM conference on Computer and communications security. ACM, 2009: 199-212.

Thanks!

Novelties in the cloud threat model

- Data and software are not the only assets worth protecting, activity patterns also need to be protected.
- Need to accommodate a longer trust chain. (incentives for companies to specialize)
- Competitive businesses can operate within the same cloud computing ecosystem.
- Mutual auditability, between cloud users and providers
- Potentially inaccurate mental models of cloud computing as an always-available service, leads to false sense of security (EC2 Crash)