Towards Abuse Detection and Prevention in IaaS Cloud Computing

Jens Lindemann

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

August 26, 2015

Introduction

- Cloud computing becoming more popular
- Cloud resources may be abused by customers
- Negative consequences for Cloud Service Provider may ensue
- Infrastructure-as-a-Service (IaaS)
  - e.g. Amazon EC2, Google Compute Engine
  - customer has full control of the OS
  - relatively hard to detect/prevent abuse

Introduction

- Cloud computing becoming more popular
- Cloud resources may be abused by customers
- N… …sue
- In…

eWEEK® Enterprise IT Technology News, Opinion and Reviews

Amazon EC2 Used as Botnet Command and Control

By **Brian Prince** | Posted 2009-12-11

Trend Micro released a report Dec. 9 highlighting what it expects to see as far as security threats in 2010. Among the more interesting predictions -- attacks on cloud infrastructures will increase.

Almost as if on cue, a report surfaced the same day that the Zeus Trojan was observed abusing the Amazon EC2 (Elastic Compute Cloud) for its command and control needs. According to Don DeBolt, CA's director of threat research for its Internet Security Business unit, a server within the Amazon EC2 network was compromised by unknown means and used as the command and control server. Files were placed on the server that bots were programmed to access from across the Internet, he said.

"Zeus Bots would call to the compromised server inside of EC2 to download instructions inside the 'config.bin' file," DeBolt explained. "The Zeus Bot then will post bank account data back to the C&C ... located inside of Amazon ... This shows how aggressive the groups behind malware are today."

This isn't the first time attackers have used an unconventional means of controlling their bots. Earlier in 2009,

Introduction

- Cloud computing becoming more popular
- Cloud resources may be abused by customers
- Negative consequences for Cloud Service Provider may ensue
- Infrastructure-as-a-Service (IaaS)
  - e.g. Amazon EC2, Google Compute Engine
  - customer has full control of the OS
  - relatively hard to detect/prevent abuse

What constitutes abuse in an IaaS context?

Google Cloud Platform Acceptable Use Policy

Use of the Services is subject to this Acceptable Use Policy.

Capitalized terms have the meaning stated in the applicable agreement between Customer and Google.

Customer agrees not to, and not to allow third parties (including End Users) to use the Services:

- to violate, or encourage the violation of, the legal rights of others (for example, this may include allowing End Users to infringe or misappropriate the intellectual property rights of others in violation of the Digital Millennium Copyright Act);
- to engage in, promote or encourage illegal activity;
- for any unlawful, invasive, infringing, defamatory or fraudulent purpose (for example, this may include phishing, creating a pyramid scheme or mirroring a website);
- to intentionally distribute viruses, worms, Trojan horses, corrupted files, hoaxes, or other items of a destructive or deceptive nature;
- to interfere with the use of the Services, or the equipment used to provide the Services, by customers, authorized resellers, or other authorized users;
- to disable, interfere with or circumvent any aspect of the Services;
- to generate, distribute, publish or facilitate unsolicited mass email, promotions, advertisings or other solicitations ("spam"); or
- to use the Services, or any interfaces provided with the Services, to access any other Google product or service in a manner that violates the terms of service of such other Google product or service.

Global Acceptable Use Policy Print version

This Acceptable Use Policy ("AUP") describes activities that are not allowed on the Rackspace® network. The AUP is incorporated by reference in your Agreement with Rackspace.

Capitalized terms used in this AUP and not otherwise defined herein shall have the meaning given in the Agreement.

Inquiries regarding this policy should be directed to abuse@rackspace.com.

Abuse

You may not use Rackspace's network or Services to engage in, foster, or promote illegal, abusive, or irresponsible behavior, including:

- Any activity or conduct that is likely to be in breach of any applicable laws, codes or regulations, including data protection and privacy laws and laws relating to unsolicited commercial electronic messages;
- Use of an internet account or computer without the owner's authorization;
- Unauthorized access to or use of data, systems or networks, including any attempt to probe, scan or test the vulnerability of a system or network or to breach security or authentication measures without express authorization of the owner of the system or network;
- Monitoring data or traffic on any network or system without the express authorization of the owner of the system or network;
- Introducing intentionally, knowingly or recklessly, any virus or other contaminating code into the Services;
- Collecting or using information, including email addresses, screen names or other identifiers, by deceit, (such as, phishing, Internet scamming, password robbery, spidering, and harvesting);
- Use of any false, misleading, or deceptive TCP-IP packet header information in an email or a newsgroup posting;
- Distributing software that covertly gathers or transmits information about a user;
- Distributing advertisement delivery software unless: (i) the user affirmatively consents to the download and installation of such software based on a clear and conspicuous notice of the nature of the software, and (ii) the software is easily removable by use of standard tools for such purpose included on major operating systems (such as Microsoft's "add/remove" tool);
- Any conduct that is likely to result in retaliation against the Rackspace network or website, or Rackspace's employees, officers or other agents, including engaging in behavior that results in any server being the target of a denial of service attack (DoS);
- Any activity intended to withhold or cloak identity or contact information, including the omission, deletion, forgery or misreporting of any transmission or identification information, such as return mailing and IP addresses;
- Interference with service to any user of the Rackspace or other network including, without limitation, mail bombing, flooding, deliberate attempts to overload a system and broadcast attacks;
- Any action which directly or indirectly results in any of our IP space being listed on any abuse database (i.e. Spamhaus);
- Conducting any gambling activity in violation of any required licenses, codes of practice, or necessary technical standards required under the laws or regulations of any jurisdiction in which your site is hosted or accessed; or
- Any action that is otherwise illegal or solicits conduct that is illegal under laws applicable to you or to Rackspace.

Offensive Content

You may not publish, transmit or store on or via the Services any content or links to any content that Rackspace reasonably believes:

- Constitutes, depicts, fosters, promotes or relates in any manner to child pornography, bestiality, non-consensual sex acts, or otherwise unlawfully exploits persons under 18 years of age;
- Publish, transmit or store any content or links to any content that is excessively violent, incites violence, threatens violence, contains harassing content or hate speech, creates a risk to a person's safety or health, or public safety or health, compromises national security or interferes with an investigation by law enforcement;
- Is unfair or deceptive under the consumer protection laws of any jurisdiction, including chain letters and pyramid schemes;
- Is defamatory or violates a person's privacy; or

What constitutes abuse in an IaaS context?

- Portscanning
- Distribution of malware
- (Distributed) Denial of Service
- Sending unsolicited e-mail
- Forged headers
- Exploiting security vulnerabilities
- Botnets

What constitutes abuse in an IaaS context?

- Survey of Acceptable Use Policies (see paper)

Prohibited activity	Amazon	Google	Microsoft	Rackspace	Softlayer	HP	ProfitBricks
illegal activities (generically)	✓	✓	✓	✓	✓	✓	✓
sending unsolicited e-mail	✓	✓	✓	✓	✓	✓	✓
distributing malware of any kind			✓		✓		
distributing harmful content	✓					✓	✓
distributing viruses	✓	✓	implied	✓	✓	✓	✓
unauthorised access to other systems	✓		✓	✓	✓	✓	✓
fake mail headers	✓		✓	✓	implied	✓	✓
fake IP addresses in network headers			implied	✓		implied	✓
intentional interference with a system	✓	✓		✓	✓	✓	✓
activities harmful to the CSP's operations	✓	✓	✓			✓	
scanning for vulnerabilities of a system	✓			✓		✓	✓
IRC, anonymisation, streaming, download, P2P services and linking to these							✓
content that compromises national security				✓			
receiving unsolicited e-mail (?!?)					✓		

How to detect abuse?

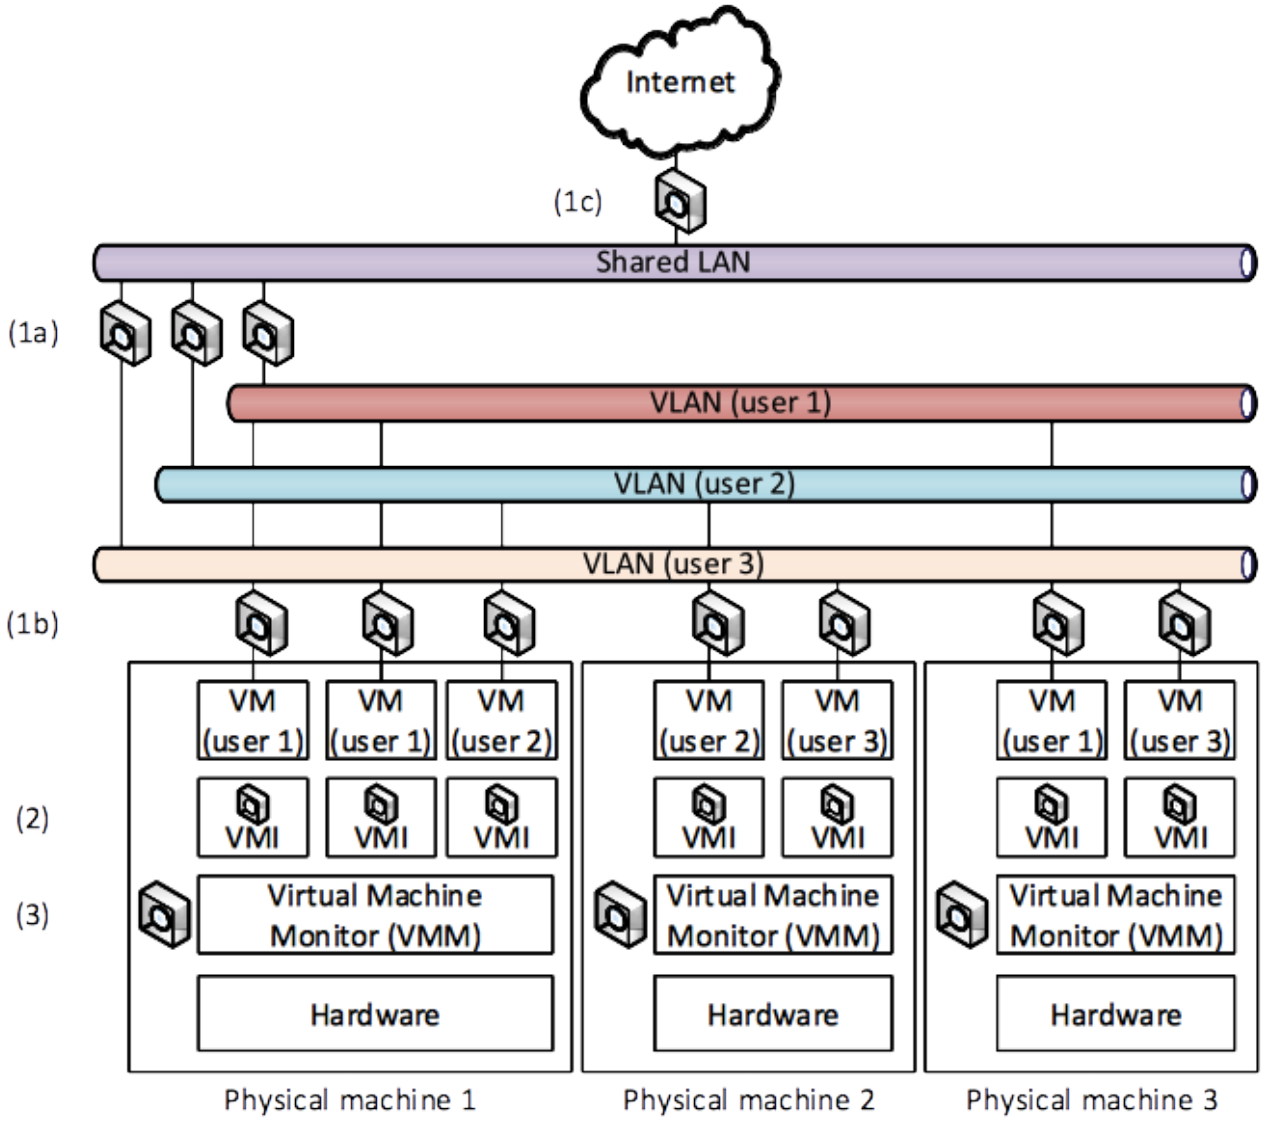

~~How~~ Where to detect abuse?

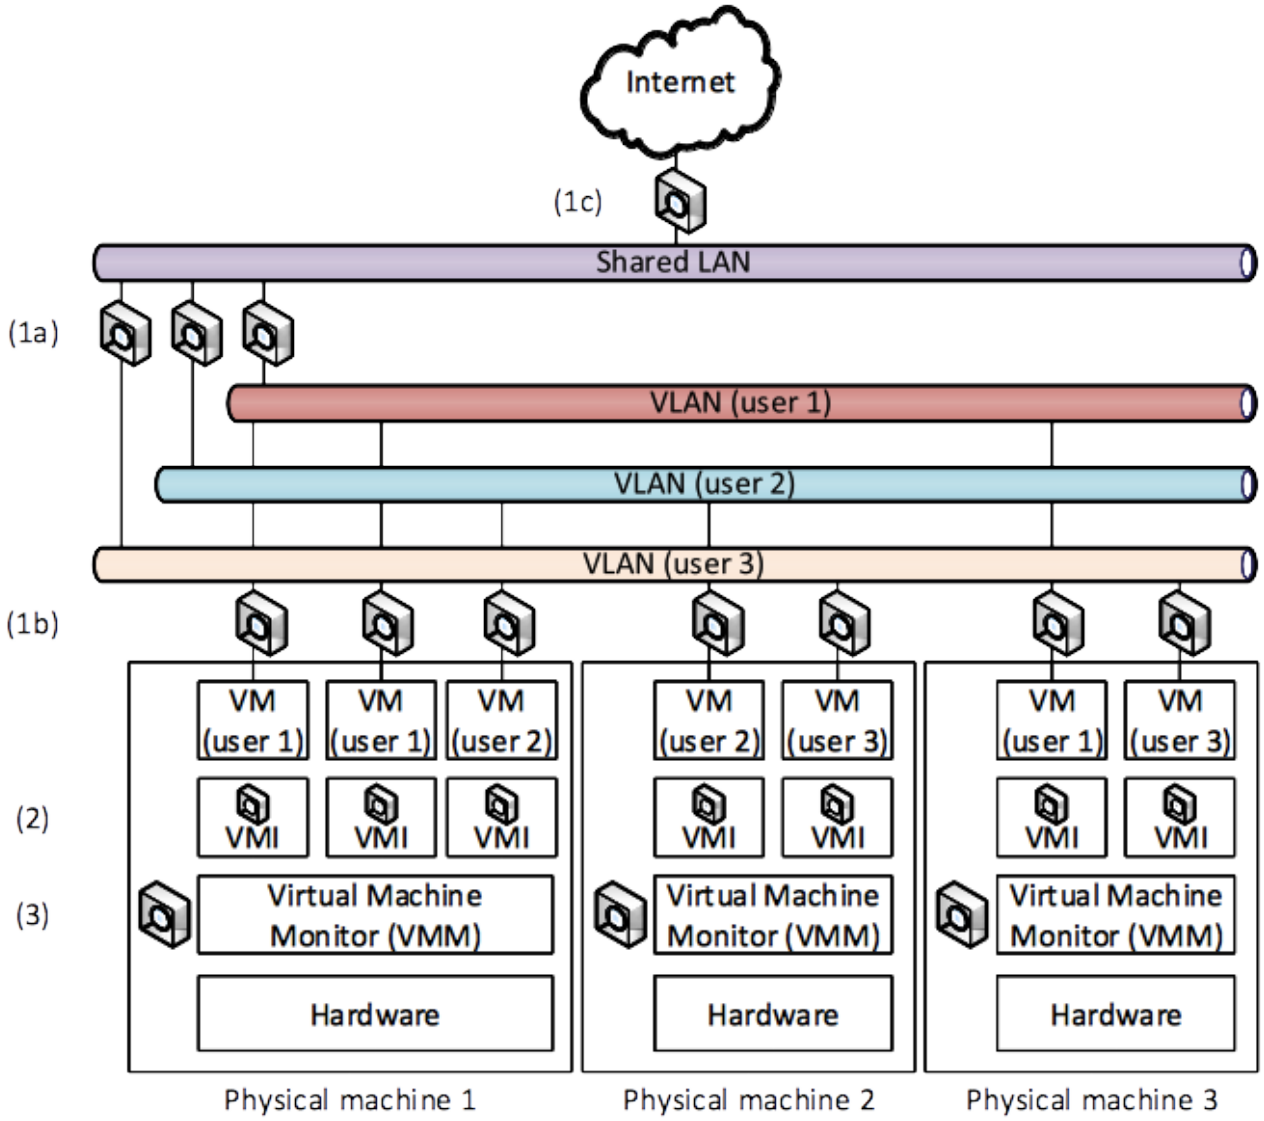

Option 1: Network-based sensors

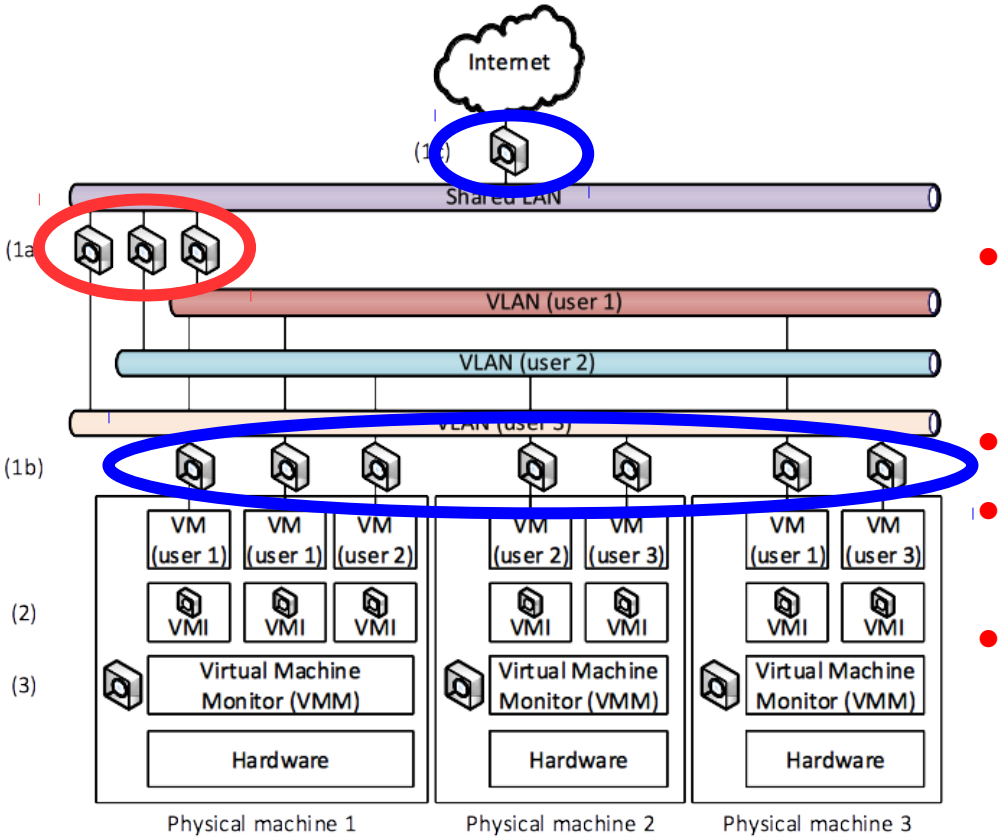

- Best compromise between visibility and ease of deployment: customer VLAN ↔ shared LAN
- Need relatively few sensors
- Can observe traffic between customers
- Cannot observe internal traffic of customers

Option 2: Virtual Machine Introspection

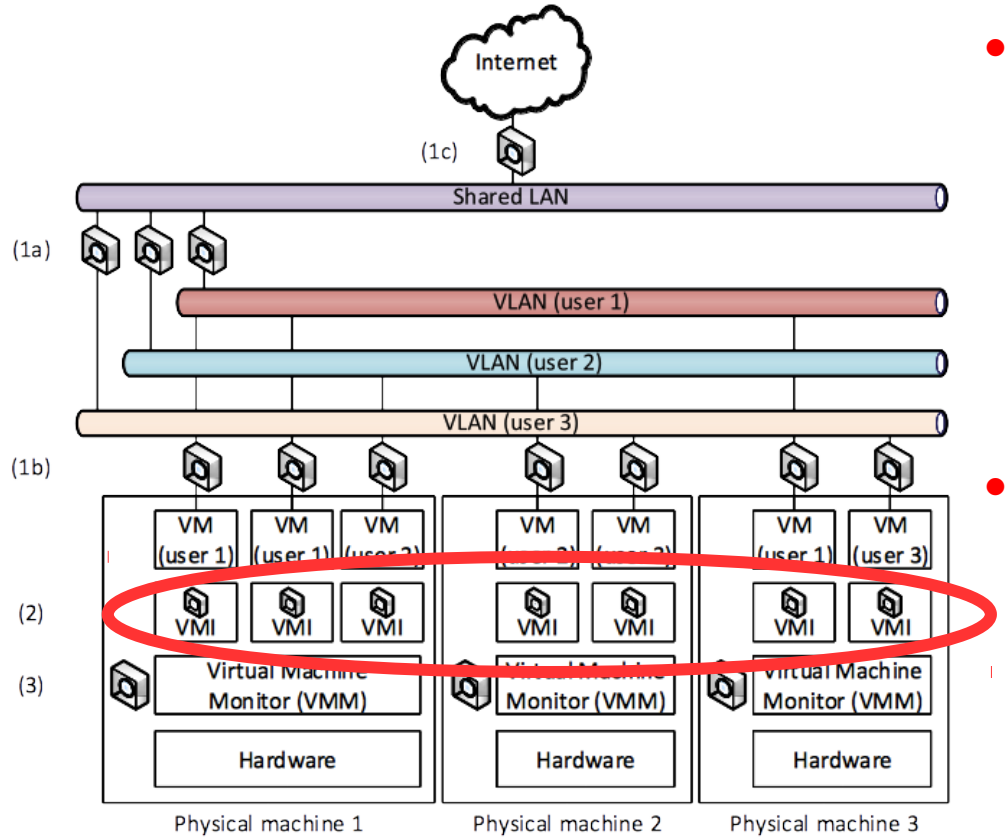

- Traditional host-based sensors problematic
  - Customer would have full control over OS (and sensor)
  - If VM is compromised, attacker could tamper with sensor
- Possible solution: Virtual Machine Introspection

Option 3: VMM-based sensors

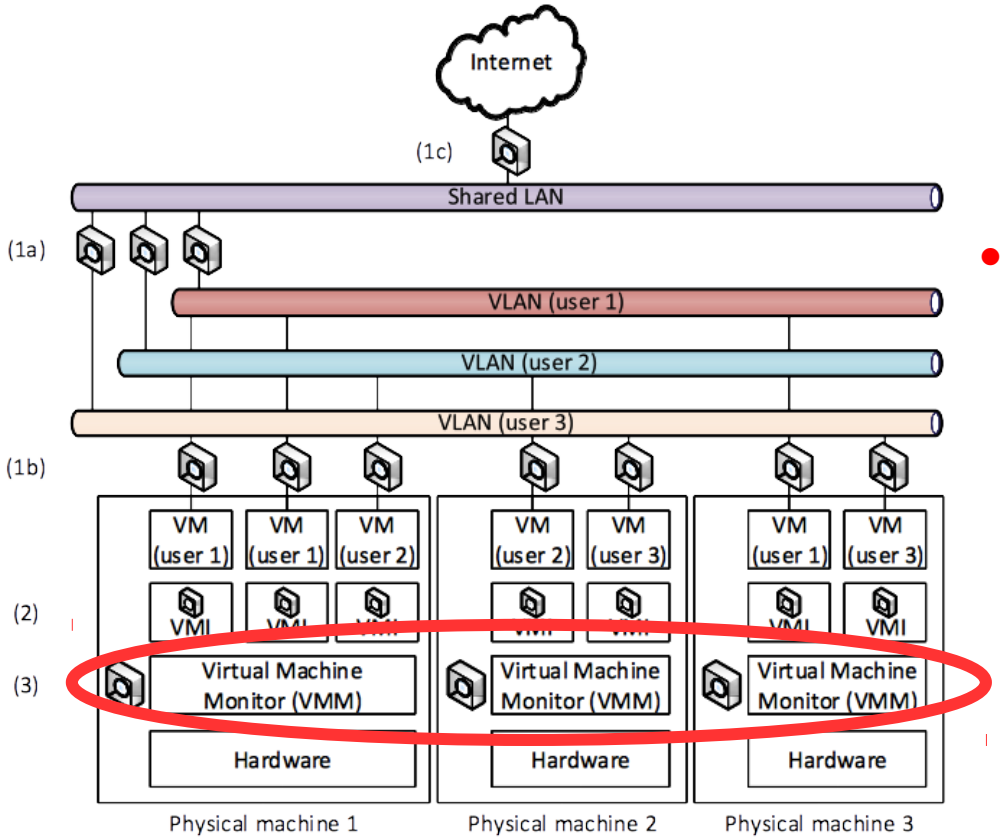

- The Virtual Machine Monitor (VMM) can also provide data that can be used for detecting abuse, such as
  - Creation and deletion time of VMs
  - Resource utilisation

Further Challenges

Prevention

- Apply existing techniques from IPS?
- Account verification
- Financial incentivisation

Reporting

- Not all abuse can be prevented automatically
- Assign reputations?
  - Event very likely to be abuse → higher impact
  - Reputation of individual VM
  - Reputation of user

Privacy Considerations

- Providers may process confidential customer data
- Violates privacy laws if customer did not explicitly agree
- Surveyed providers do not give much information on what types of data they currently process for abuse/intrusion detection
  - HP and Amazon reserve the right to investigate any suspected violation of their AUPs (no more detail given)
- Privacy policies not specific to cloud offerings
- Avenue of research: Is it possible to detect abuse in a privacy-friendly way?

Conclusion

- Survey of AUPs → types of abuse
- Architecture – where could abuse be detected?
  - Network communications
  - Virtual Machine Introspection
  - Virtual Machine Monitor data
- Further challenges
  - Prevention
  - Reporting
  - Privacy

- E-Mail: jens.lindemann@informatik.uni-hamburg.de