# CloudIDEA - Cloud Intrusion Detection, Evidence preservation and Analysis

Benjamin Taubmann * Hans P. Reiser * Thomas Kittel ** Andreas Fischer *
Waseem Mandarawi * Hermann de Meer *

Universität Passau*, Technische Universität München**
{*firstname.lastname*}@uni-passau.de, kittel@sec.in.tum.de

## 1. Problem statement

Virtual machines (VMs) hosted on Infrastructure-as-a-Service (IaaS) clouds are an attractive target for attackers. Cloud providers and cloud customers who want to detect, analyze, and preserve evidence about malware attacks in IaaS clouds are faced with multiple problems: Cloud *customers* cannot use existing intrusion detection tools that require access to physical hardware or make use of virtual machine introspection (VMI). Cloud *providers* lack contextual knowledge about the system executing within the VM and do not know which intrusion detection heuristics fits for it.

## 2. Goals

We want to enhance the security of IaaS clouds by designing an architecture for malware and intrusion detection, analysis and evidence collection. Our proposed architecture achieves the following goals:

- Offer customizable security services to cloud customers, e.g., VMI-as-a-Service.
- Enhance protection of cloud infrastructure and other VMs against attacks originating from a VM.
- Be sufficiently lightweight to be usable in production environments with negligible overhead.
- Provide detailed insight into malware behavior and collect conclusive evidence about attacks.

## 3. Architecture and analysis framework

The CloudIDEA architecture is designed to be modular, scalable and offers semantic aware introspection. It consists of a decentralized analysis framework and a central management component. The analysis framework itself is part of every physical cloud node and contains several introspection and tracing plug-ins. This includes VMI, network traffic, hypercalls to the hypervisor and further performance statistics. The central management component is composed of the *decision engine*, *behavior database* and the *virtual network management*.

Whenever a VM behaves abnormal, the *decision engine* defines actions based on external inputs, such as user configurations or service level agreement (SLA), and internal information, such as available resources, VM interdependencies, and migration cost. Depending on the expected attack and the resource intensity of the analysis modules, it can activate additional analysis modules, replace a malicious VM with a fresh instance, or trigger the virtual network management to migrate and isolate a suspicious VM.

We divide between two types of tracing methods. *Lightweight* plug-ins are used to detect intrusion in a production environment causing only a negligible overhead. *Heavyweight* plug-ins are used for further investigation on VMs that might be infected and are only enabled if required. As the heavyweight plug-ins are more resource intensive, the system under analysis can be migrated to a dedicated investigation host, based on the expected attack and available cloud resources.

The log data that is obtained by the different plug-ins of the analysis framework is stored in a central database, the *behavior database*. This database is used in order to create behavior models for each VM using machine learning algorithms and also stores and updates these behavior models. The decision engine can then use these models to decide on how to configure the current monitoring behavior.

The *virtual network management* module is used by the decision engine to assign interconnected VMs to physical cloud resources. It has to take into account both, the underlying network, and VM interdependencies. It ensures under these constraints that after migration all VMs still meet their SLAs.

## 4. Contributions

CloudIDEA is a novel architecture for malware detection, analysis and evidence collection in IaaS based cloud data centers. It leverages several monitoring techniques in order to learn more about benign behavior and to create behavior models of all VMs at runtime. The overhead of the analysis can be configured at runtime so that it can be minimized on production environments. If anomalies are detected in a VM, the system can be analyzed more intensely in order to detect if the anomaly is caused by malware. Additionally, CloudIDEA provides an interface for cloud customers to be informed about intrusions and offers forensics means for evidence collection and malware analysis. Therefore it is able to offer VMI-as-a-Service.

## Acknowledgments

The research leading to these results was supported by the Bavarian State Ministry of Education, Science and the Arts as part of the FORSEC research association.

# CloudIDEA - Cloud Intrusion Detection, Evidence preservation and Analysis

Benjamin Taubmann*, Hans P. Reiser*, Thomas Kittel**, Andreas Fischer*, Waseem Mandarawi*, Hermann de Meer*
*University of Passau
**Technische Universität München
{firstname.lastname}@uni-passau.de, kittel@sec.in.tum.de

## Problem

Cloud infrastructures are an **attractive malware target,** as many security-sensitive services are provided

Common **IDS solutions are impractical** for cloud environments because the provider lacks detailed contextual knowledge about the system running in the virtual machine

Customers have detailed information about their virtual machine but can not use tools that require **physical hardware access** or make use of **virtual machine introspection**

## CloudIDEA Architecture

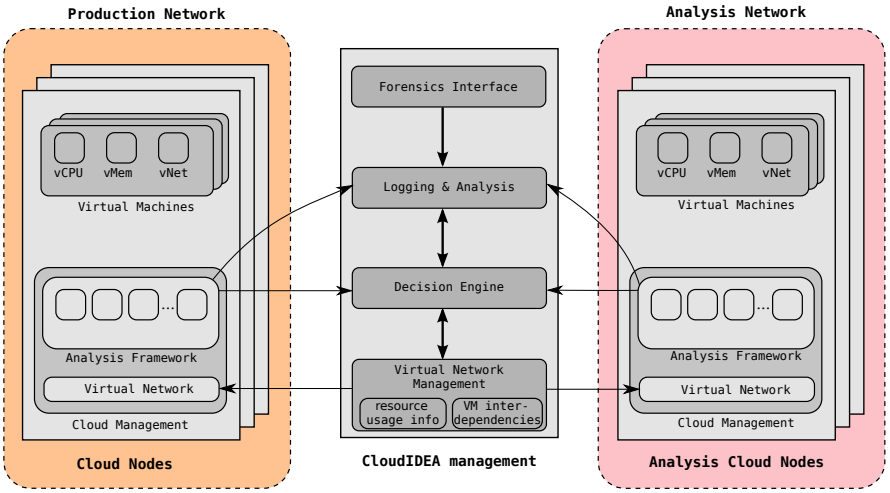


## Analysis Framework

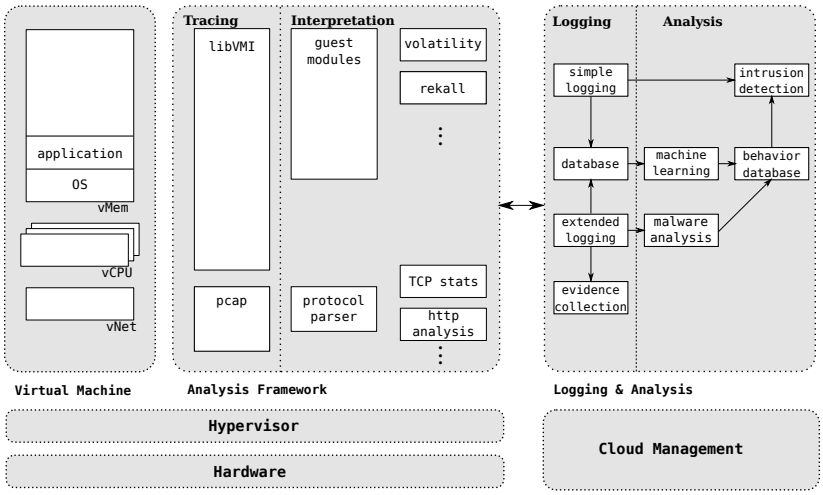


### Plug-in based monitoring
- virtual machine introspection
- network traffic
- virtual machine hypercalls
- performance statistics

### Logging
- central storage
- standardized storage format
- provide data for malware analysis
- preservation of evidence

### Analysis
- combine the output of several plug-ins
- use **machine learning** to create **behavior models** of virtual machines
- anomaly based **intrusion detection**

## Goals

- provide **customizable security services** to cloud customers, e.g., VMI-as-a-Service
- defend **large scale cloud** environments with a combination of monitoring and attack detection techniques
- maintain **scalability** by the application of analysis plug-ins with different resource requirements
- **select tracing plug-in automatically** based on the current threat level and available cloud resources
- support **attack knowledge exchange** with central data collection and analysis
- enhance security by **on demand virtual machine migration** and isolation when intrusions are detected
- enable **live malware forensics** by leveraging a dedicated analysis environment
- provide customers an interface for **preservation of evidence**

## Contact

Benjamin Taubmann
University of Passau

UNIVERSITÄT PASSAU

FORSEC
Sicherheit hochgradig vernetzter IT-Systeme

Technische Universität München

The research leading to these results was supported by the "Bavarian State Ministry of Education, Science and the Arts" as part of the FORSEC research association.