CERIAS

The Center for Education and Research in Information Assurance and Security

PURDUE UNIVERSITY

Forensic Evidence in Apache's CloudStack (a work in progress)

Will Ellis, Dr. Sam Liles

Abstract

Apache's CloudStack allows service providers to create Infrastructure as a Service (IaaS) solutions. The open source nature of the software is a draw for organizations wanting to provide low cost solutions either in the form of private clouds or reselling a cloud infrastructure to their clients. As services such as CloudStack grow and compete with offerings from Amazon, Microsoft, or Google there will be a growing need to gather forensic evidence from cloud environments. Law enforcement is able to request information from the major vendors, but an open source platform brought online to solve the needs of a small organization poses a significant problem for local law enforcement. The study will attempt to provide a compendium of forensic evidence locations in Apache's CloudStack environment. This will enable law enforcement agencies to know where and what to look for in the environment, should they encounter an instance during the course of their investigations.

Methods

Utilizing Oracle's VM VirtualBox and Apache's DevCloud instance of CloudStack we will implement a small private cloud. The private cloud will allow for an end user, or "customer", to perform actions such as bringing up a virtual machine on the platform. The customer will perform some "illegal" activity such as attempting a denial of service attack from the virtual platform to another instance in the cloud environment. It is expected there will be logs indicating the instantiation of the virtual instance as well as logs of the traffic from the virtual box at the cloud administration level. The key is finding these logs to provide a method of extracting the evidence for attribution.

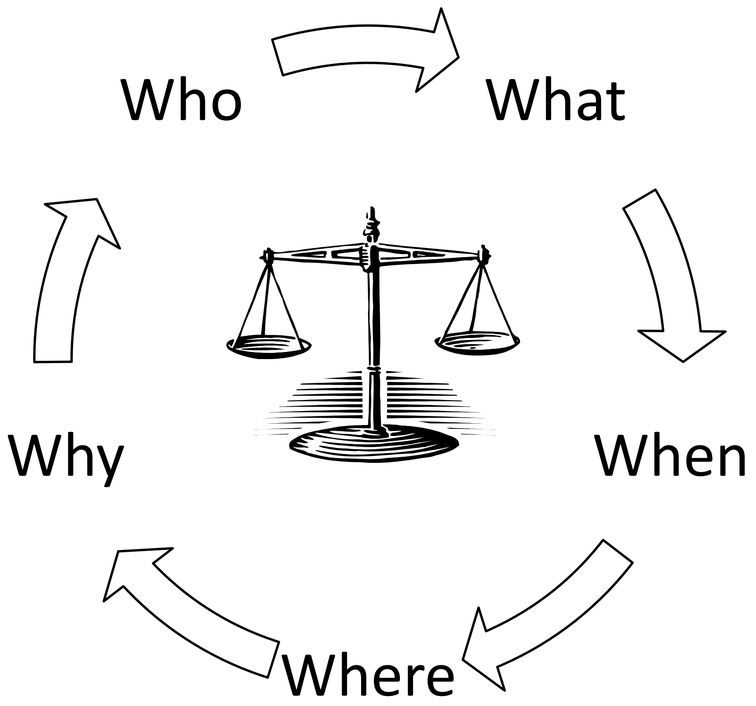

Significance

According to Cisco Global Cloud Index, annual global cloud IP traffic will reach 4.3 zettabytes by the end of 2016. This represents 355 exabytes per month (up from 57 exabytes per month in 2011) [1]. This amount of traffic is beyond the scope of searching by an individual without tools to aide in log analysis and capture. Before such tools can be designed research must be conducted to identify whether logs exist and identify the data gaps in the logs.

References

[1] Cisco. (2012). *Cisco Global Cloud Index: Forecast and Methodology, 2011–2016* [White paper]. Retrieved from http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns1175/Cloud_Index_White_Paper.pdf

cloudstack

CERIAS

Discovery Park e-Enterprise Center