

## Mitigating Data Exfiltration in SaaS Clouds

Existing processes and methods for incident handling are geared towards infrastructures and operational models that will be increasingly outdated by cloud computing. Research has shown that to adapt incident handling to cloud computing environments, cloud customers must establish clarity about their requirements on Cloud Service Providers (CSPs) for successful handling of incidents and contract CSPs accordingly. Secondly, CSPs must strive to support these requirements and mirror them in their Service Level Agreements. Intrusion Detection Systems (IDS) have been used widely to detect malicious behaviors in network communication and hosts. Facing new application scenarios in Cloud Computing, the IDS approaches yield several problems since the operator of the IDS should be the user, not the administrator of the Cloud infrastructure. Cloud providers need to enable possibilities to deploy and configure IDS for the user - which poses its own challenges. Current research and commercial solutions primarily focus on protecting against Denial of Service attacks and attacks against the Cloud's virtual infrastructure. To counter these challenges, we propose a capability that aims to both detect and prevent the potential of data exfiltration by using a novel deception-based methodology. We also introduce a method of increasing the data protection level based on various threat conditions.

*Index Terms*—Cyber Deception, Intrusion Detection, Cloud Storage Provider.



**Dr. Duane Wilson** has been a practitioner in the field of cyber security for almost 20 years. After earning his bachelor's degree in Computer Science, Dr. Wilson went on

to earn a Master of Engineering in Computer Science from Cornell University, a second Master of Science in Security Informatics, and a Ph.D. in computer science from Johns Hopkins University. The balance of both scholastic and operational experience allows him to have insight both into the theory and practice of cybersecurity. As a subject matter expert in cyber technology research, Duane's areas of expertise range from cyber security, business development, computer science engineering, and network protection. Duane provides subject matter expertise toward various Cyber Training & Education initiatives through his consulting firm Wilson Innovative Solutions LLC. The company's focus is to provide innovative cyber research and development, advanced cyber security consulting, and cyber security training for all levels of experience in the commercial and federal sectors.

### I. INTRODUCTION

According to [1], security incident handling, an integral part of security management, treats detection and analysis of security incidents as well as the subsequent response activities (i.e., containment, eradication, and recovery) as outlined below:

- Detection: Discover indicators of possible security incidents
- Analysis: Ascertain that indeed a security incident is at hand and understand what exactly has happened/is happening

- Containment: Contain the incident before it spreads and overwhelms resources or increases damage
- Eradication and Recovery Eliminate system changes caused by the incident and recover normal operations
- Preparation and Continuous Improvement: Set up/adapt incident handling activities according to changing requirements and system/threat landscape.

Existing processes and methods for incident handling are geared towards infrastructures and operational models that will be increasingly outdated by cloud computing. Research has shown that to adapt incident handling to cloud computing environments, cloud customers must establish clarity about their requirements on Cloud Service Providers (CSPs) for successful handling of incidents and contract CSPs accordingly. Secondly, CSPs must strive to support these requirements and mirror them in their Service Level Agreements. Typically, security incidents will cross the boundaries of both customer and CSP in each of the common deployment models (i.e., Software-as-a-service, Platform-as-a-service and Infrastructure-as-a-service or SaaS, PaaS, and IaaS respectively) denoting both joint responsibility and access. This must be taken into account when setting up incident handling or prevention for a cloud infrastructure. Our research focuses on attacks confined to the SaaS deployment model due to the limited level of control a user has over the security of their data once it is stored within the cloud.

Intrusion Detection Systems (IDS) have been used widely to detect malicious behaviors in network communication and hosts. IDS management is an important capability for distributed IDS solutions, which makes it possible to integrate and handle different types of sensors or collect and synthesize alerts generated from multiple hosts located in the distributed environment. Facing new application scenarios in Cloud Computing, the IDS approaches yield several problems since the operator of the IDS should be the user, not the administrator of the Cloud infrastructure. Extensibility, efficient management, and compatibility to virtualization-based context need to be introduced into many existing IDS implementations. Additionally, the Cloud providers need to enable possibilities to deploy and configure IDS for the user. To date, a number of theoretical frameworks have been proposed to address these concerns. [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17]. None, however, have been implemented in practice, hence, the concepts that are presented cannot be effectively validated. Other research focuses on protecting the Virtual Machine (VM) within the cloud infrastructure from attacks - primarily focused on monitoring VM resources to maximize utilization or prevent Denial of Service attacks. [18], [19], [20], [21], [22], [23], [24], [25], [26], [27].

One major gap identified in our assessment of the current state of affairs in applying IDS concepts in a Public cloud setting is data exfiltration detection and prevention. Data exfiltration is defined as the unauthorized transfer of data from a computer. To address this particular gap, we propose a novel method for detecting and preventing data exfiltration using cyber deception [28]. According to [28], cyber deception is a *deliberate act to conceal activity on a network, create uncertainty and confusion against an adversary's efforts to establish situational awareness and to influence and misdirect adversary perceptions and decision processes*. To accomplish this, we focus on the Detection and Containment aspects of the incident handling process. In the case of our research, containment refers to the efforts aimed at preventing the exfiltration from occurring - not preventing the spread of a cyber infection (e.g., worm or other malware). Our approach is implemented such that cloud users do not have to rely on the protection mechanisms offered by the

Cloud Storage Provider (CSP) which are often limited to common Data Loss Prevention techniques - which operate based on pre-defined policies. The limitations of these approaches include: 1) can be bypassed by the use of encryption 2) are also under the control of the CSP (client-side or server-side) and 3) are not designed to prevent against unknown attacks (i.e., that are not specified in the deployed policy). To address these and other limitations we propose the following contributions:

- A novel data anonymization technique using format-preserving encryption for producing deception file objects based on input data
- The design of a novel capability to detect data exfiltration attempts in Software-as-a-Service (SaaS) cloud storage providers without having to rely on the provider or a middleware vendor using cyber deception.
- The design of a novel data protection mechanism that adjusts protection level based on threat conditions.
- A proposed prototype implementation (future work section) of our design across several popular cloud storage vendors to show its viability in a number of cloud settings The presented approach is not designed to replace existing solutions, but provide a mechanism that increases the Defense in Depth of data within a SaaS cloud environment.

The rest of this paper is organized as follows:

- Section 2 provides a background of the threat space we address with this work and presents related research,
- Section 3 describes the design of a proposed Cloud Storage Environment that is protected via Cyber Deception concepts,
- Section 4 concludes the paper, and
- Section 5 provides areas for future work - primarily focusing on a prototype implementation across multiple Cloud Storage Providers.

## I. BACKGROUND

As mentioned above, existing research largely has focused on the protection on Virtual Machines within Cloud infrastructures as well as on proposing frameworks that would aid in the detection/protection of data within the cloud. None

of the aforementioned methods resulted in implementations that could be used by the public to protect data in SaaS storage infrastructures. The following research outlines a variety of methods attempting to remedy this. One common weakness with those discussed below is that the user is not in control of the protection mechanisms applied to their data. Hence, they have to trust that 1) the protections offered are adequate, 2) the audit information is accurate, and 3) the protections will not fail in the event of adversarial compromise.

#### A. Threat Model

When user data is stored in a Software as a Service (SaaS) Public Cloud setting, it is subject to the security features provisioned by the CSP. As mentioned above, this leads to a number of threats that user's data can be subject to without their knowledge or control. We focus our efforts on the methods by which sensitive data can be exfiltrated once it has been transferred to the cloud. Our research does not examine uploads for policy conformance. We focus specifically on data that is resident on the cloud - as we assume that users have followed any relevant policy that governs the storage of data in the cloud. Within the cloud, there are 3 primary ways data can be accessed and hence subsequently exfiltrated: 1) Via Sharing (individual or group), 2) Via download (web interface or client application), and 3) Via the cloud interface (within the cloud environment). Our research focuses on the detecting and prevention of data exfiltration via sharing and download requests. As previously stated, SaaS environments are not under the control of the user - hence we cannot modify the cloud infrastructure to monitor accesses within the environment itself. An additional feature we provide is the ability to be aware of varying threat conditions - providing a higher level of protection when the adversarial threat is at its peak.

Within the cloud setting, data exfiltration can be prevented via the following methods: 1) Detect an attempt and prevent it before it happens - would require the monitoring all potential paths of theft in cloud and intercepting malicious looking attempts, 2) Allow/Disallow information from being downloaded based on individual or organizational policies (after data is stored), or 3) Allow access to information in cloud only when given permission or authorization. In the SaaS deployment model, method 1 is not possible without the modification of

the Cloud infrastructure. Method 2 focuses on protecting information from being exfiltrated via the Cloud

- not after the information is stored. This has also been addressed by CloudFilter (a paper we discuss below). We aim to control data exfiltration via the 3rd method to prevent attempts by malicious adversaries who have compromised a user account [30]. We assume that users are generally trustworthy (i.e., not malicious insiders) and are able to make wise decisions about who needs to have access to their data

- prior to it being uploaded. Our approach can be used as an alternative or augment to client-side encryption cloud storage vendors that offer a method of protecting user data from the cloud provider themselves.

#### B. Related Work

##### 1) Cyber Deception:

The art of deception has been used as a technique to lure attackers away from important data and learn about their techniques and target [35], [43]. In the space of cloud data protection, two methods have been proposed: Fog Computing and a Honeypot implementation for the cloud. In [31], the authors propose a novel deception-based concept that focuses on the protection of insider data theft attacks by monitoring data access in the cloud and detecting abnormal access patterns. They accomplish this by profiling the normal behavior of users and upon the detection of an anomaly, they produce and return decoy information (e.g. docs, honey files, honeypots) to the attacker. Being that there is no implementation in an actual cloud environment, it is difficult to ascertain the viability of their approach due to the fact that quantifying normal user behaviors is a difficult problem [44] - in general - and much more in a cloud computing environment that is largely multi-tenant. Other approaches [32], [33] attempt to apply anomaly-based detection mechanisms to the cloud but suffer from the same issues - profiling users in such a complex environment and not being implemented in an actual cloud environment.

Similarly, [34] aims to employ deception concepts in a private cloud using a honeypot. The authors deploy the Snort Network Intrusion Detection System (NIDS) in the Eucalyptus private cloud as a proof of concept implementation. They incorporate

a honeypot into the NIDS to further entice attackers and provide a mechanism to learn more about similar future attacks. The primary weaknesses of their approach is that 1) it would require a major modification to existing cloud infrastructures to implement their methodologies, 2) it uses a signature-based approach for detection, and 3) their deployment would require awareness of Cloud infrastructure to enable full OS coverage (e.g. Windows, Linux, OS X). Neither the Fog Computing, Honeypot, or Anomaly-Based approaches as described would be feasible in practice due to these limitations. Our proposed solution extends the concepts present within Fog Computing by providing an implementation in an actual cloud storage environment, providing a novel mechanism for users to generate decoy information that is derived from their uploaded files, and focusing specifically on data exfiltration problem. Additionally we provide a mechanism that can adjust to current threat conditions as we subsequently discuss.

**Commercial Cloud Providers:** There are a number of Commercial Cloud Providers that provide features to protect the security of cloud user data. The primary security features they all provide are: the protection of data in transit and data at rest and two-factor authentication. Other than Box, they all offer limited oversight of the activities associated with files/folders. Within Dropbox, the event log tracks when users create, delete, and restore Dropbox folders. Once an action is taken, the event log records the name and ID of the Dropbox folder, the action, the user who made the change, and the date it was performed [36]. For Google Drive, the types of file activities that are recorded are the: 1) moving and removing, 2) renaming, 3) uploading, 4) sharing and unsharing, and 5) editing and commenting. There are no specific additional security features provided in the documentation [37]. SugarSync<sup>1</sup> monitors any edits to a file, any new files created, files deleted, and changes to a folder or sub-folders [38]. Box [39] offers the most advanced native security features of all the existing CSPs (to include the ones described above). Box provides comprehensive reporting, logging, and audit trails in order to track account activity, file access, settings changes and nearly everything else that occurs in Box.

These features allow users to monitor for access violations only, but are not combined with a

mechanism to prevent data loss. Box, however, does provide a separate mechanisms to detect and prevent data loss via extended security policies that allows users to keep content confidential, mitigate data loss, flag risky sharing requests, and white list certain domains that should be trusted. Some native features are present, but others are offered via collaborative security partners. CipherCloud (one such partner) offers of a secure middleware service that provides users with insight into cloud data [38]. They provide a range of features to include: 1) User activity monitoring, 2) Anomaly Detection, 3) Comprehensive Audit Logs, and 4) Ongoing Data Loss Prevention Monitoring. This solution has a number of promising features, however, these all require trust of an entity like CipherCloud with their data (risk transference) and it also comes at a premium cost and is geared towards organizations - not individual users.

**CSP IDS Implementations:** In [40], Greg Roth and Don Bailey present a detailed overview of Amazon Web Services (AWS) Intrusion Detection capabilities. These capabilities are a combination of several features: AWS's Identity and Access Management, Multi-Factor Authentication, Amazon S3 Bucket Logging, Security Audit Role, Write Once Storage for Data Provenance (Versioning), and Auditing Logs. Their approach is to use this compendium of features to detect unauthorized access to user data based on user roles. The major weakness of this role-based approach is that it is a script based IDS, not an actual system that the general user can use. Another weakness is dynamically adjusting to changing roles may be difficult to capture and categorizing employee's using clear cut roles may be difficult to define. Due to the number of components involved, it would require a security expert to setup, configure, and interpret the resulting alerts. It is limited in scope - focusing primarily on access permission violations not specific attack types (e.g., data exfiltration, integrity breaches) and is specific to Amazon EC2, hence, not interoperable with other CSP infrastructures.

Lastly, in [41], Ioannis and Pietzuch describe CloudFilter - a system for the practical control of

---

<sup>1</sup> [www.sugarsync.com](http://www.sugarsync.com)

sensitive data propagation to the cloud. Their scheme monitors HTTP File Transfers to the Cloud independent of the CSP being used. Upon upload of each user file or files, CloudFilter enforces data propagation policies to confirm that the upload request is authorized. Their approach is limited to just monitoring the HTTP protocol which would cause CloudFilter to miss sensitive data being propagating via another protocol. One major strength of the described approach is that no modification is necessary to the CSP infrastructure, however, it does require modification of files upon upload and web browsers for user identification. This solution does not address data exfiltration once it is resident within the cloud, because it only monitors uploads.

## II. PROTOTYPE DESIGN

The goal of our prototype is to provide a transparent and practical solution for protecting sensitive data stored in Software-as-a-Service (SaaS) Cloud Storage Provider, which represent stored data as files. An important requirement is for the system to be applicable across different cloud storage providers with minimal end-user configuration. We accomplish this by leveraging a cloud agnostic API (discussed in next section) to avoid having to adapt to the API of specific cloud storage service. Secondly, the deceptive information should be automatically generated based on the input provided to accompany the files that are uploaded. Lastly, we provide a solution that can be employed in a representative scenario based on the criticality of the cyber situation and/or INFOCON status. We accomplish this without modification of the cloud infrastructure. Our prototype provides the following component engines all integrated with a cloud storage provider:

- **Generation engine** that produces deceptive Objects or Decoy Documents with high similarity to original content
- **Detection engine** that monitors access requests and verifies originating source.
- **Prevention engine** that controls access to legitimate documents via deception-based cloud access layer
- **Threat engine** that makes adjustment of deception-based protection “level” based on varying threat conditions (e.g., INFOCON Level)

Below, we progress through an exfiltration attempt

against our framework. There are several steps that take place between a user (or enterprise) and a cloud provider during an upload operation. Prior to step 1, the initial threat level is set to INFOCON Level 5 - indicating the lowest threat environment. In step 1, a user submits a file via our prototype application and several host system identifiers are embedded within the file to include: Computer MAC Address, Network IP Address, Network Hostname, UserID, and Hash of 4-tuple identifiers (MAC, IP, Hostname, UserID).

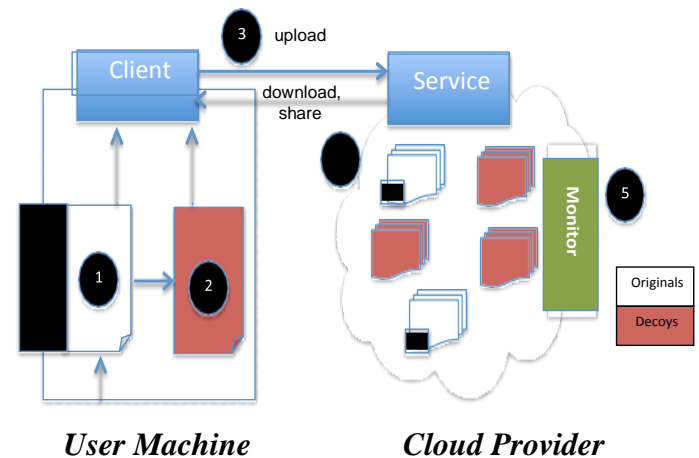


Figure 1: Proposed Prototype Architecture: Prior to entering deceptive file editing, INFOCON Level is set. 1) A file is submitted and host system identifiers are embedded in the file 2) Files are deceptively modified 3) Upload and store original and decoy document to cloud 4) A file within the cloud is access 5) Monitor exfil attempt(s) for a particular file.

In step 2, the file is read and each numerical data element of the file is anonymized - creating the decoy document. Sensitive data within a file is usually represented as numerical data (e.g., SSN, Credit Card), but can also be represented as names such as locations and people. Numerical data can be recognized using regular expressions. Name data can be recognized using Natural Language Processing techniques such as Named Entity Recognition[45]. Once identified, sensitive data can be changed to a random value or to a more strategic value using look up tables and online searches depending on the level of security necessary. Step 3 results in both the decoy document and the original document being uploaded and stored in the cloud. Steps 1-3 represent how a file is uploaded to the cloud. Once the file is uploaded, it persists on the cloud as any other file uploaded. Next, assume some time as passed since the files were uploaded to the

cloud. Step 4 represents a request(s) for a particular file stored within the cloud. Lastly, step 5 commences the monitoring process for exfiltration attempts and threat level updates. Steps 4-5 represent an exfiltration attempt on some file. Figure 1 provides an illustration of our prototype architecture.

A few foundational concepts in cyber deception have pre- ceded our design and subsequent implementation - which we outline here for context. Honeyfiles are bait files intended for hackers to access. The files reside on a file server, and the server sends an alarm when a honeyfile is accessed [46]. In [47], Wang et. al. add honey activity to enhance the realism of honeyfiles. The Decoy Document Distributor (D3) System is a tool for generating and monitoring decoys. An accompanying website (FOG) allows users to download files, such as tax documents and receipts, that appear authentic but actually contain spurious information [48]. Lastly, Fox provides a mechanism for placing decoy documents in places that are more likely to get accessed [49]. As previously stated, there has been much previous work in the area of using deceptive data to defend against malicious adversaries. Yet no prior work has attempted to adapt these concepts in a SaaS cloud storage environment - which is out of the control of a user. Secondly, no work has been proposed that provides the ability to adjust to varying threat situations (e.g. INFOCON levels). Additionally, we address the specific threat of data exfiltration in SaaS storage environments. Lastly, our approach uses pre- existing content to generate deception data whereas previous work focuses on generating data based on sensitive properties. For decoy object generation, we follow a different pro- cess than is used to generate a honeyfile, where users are responsible for selecting the honeyfiles in their environment. In our scheme, we rely on data within each file to generate a replica highly similar to original content. Creating simi- lar documents can influence an attacker’s decision making. Not knowing which document is the original and therefore what content is legitimate can force them to make incorrect decisions and potentially critical errors. Creating a higher ratio of replica to original files makes the this technique probabilistic in that the ability of an attacker to guess the correct file if inversely related to the number of

replica files. For the generation process, each file is tokenized (primary delimiter would be spaces) into individual words and the word structure examined to determine its sensitivity. We make a simple assumption that sensitive content will likely take a numerical format (not including alphanumeric characters). Hence, if a word is determined to be sensitive it is converted using format-preserving encryption. In cryptography, format- preserving encryption (FPE) refers to encrypting in such a way that the output (the ciphertext) is in the same format as the input (the plaintext) [50]. Upon completion of the document analysis process, the decoy document will be created. The naming convention for the decoy document will be similar to the original to ensure plausibility As previously stated, there are several ways that data can be exfiltrated via the cloud: Downloading, Sharing, or Cloud Access. To adequately detect and prevent exfiltration through these channels, we provide a system that monitors all requests to information residing on a cloud storage provider’s infras- tructure from a host system. We are unable to monitor access requests originating from the Cloud Provider itself to the design of SaaS environments. Hence, we focus on Download and Sharing operations. Whenever a download or sharing request is made, the default identifier (e.g. host MAC address) is used to confirm the legitimacy of the request. We use the MAC address because it is unique to each machine. Though it can be spoofed, it is a metric that we have access to that is unique to each machine<sup>2</sup>. If the originating host identifier matches the embedded identifier within the file, the actual file is returned to requesting party. Otherwise the corresponding decoy document is returned and the event is logged. Increasing levels of protection are imposed as the threat level increases as will be discussed below.

Cyber threats commonly change over time within a partic- ular environment. We argue that this is also the case within a cloud storage environment. Hence, a solution presented to protect data within that environment, should be able to adapt to the ebb and flow of the threat landscape. Hence our approach proposes an emulated threat environment

---

<sup>2</sup> Ideally, a value that is unique to every machine and stored in a secure location on the device (e.g. a value created and stored on

the Trusted Platform Module[51]) should be used.

that mirrors the IN-FOCON levels<sup>3</sup> used by the US government to designate the presence of a hostile threat. With each threat level increase, our approach requires additional system identifiers to validate an access request. We describe the U.S. government INFOCON levels below:

- INFOCON 5 describes a situation where there is no apparent hostile activity against computer networks. Operational performance of all information systems is monitored, and password systems are used as a layer of protection.
- INFOCON 4 describes an increased risk of attack. Increased monitoring of all network activities is mandated, and all Department of Defense (DoD) end users must make sure their systems are secure. Internet usage may be restricted to government sites only, and backing up files to removable media is ideal.
- INFOCON 3 describes when a risk has been identified. Security review on important systems is a priority, and the Computer Network Defense system's alertness is increased. All unclassified dial-up connections are disconnected.
- INFOCON 2 describes when an attack has taken place but the Computer Network Defense system is not at its highest alertness. Non-essential networks may be taken offline, and alternate methods of communication may be implemented
- INFOCON 1 describes when attacks are taking place and the Computer Network Defense system is at maximum.

## CONCLUSIONS

In this paper, we have proposed several methods to address the problem of data exfiltration detection and prevention in a SaaS Cloud Storage environment. Our proposed methods will overcome the limitations of existing approaches that primarily focus on the protection of the virtual infrastructure - not the data that is stored in these environments. We leverage a number of concepts in cyber deception for data protection and introduce the ability to adapt to varying threat conditions. Data Loss Prevention (DLP) is often implemented within a cloud storage

environment to protect against data exfiltration. However, we found that DLP has a number of shortfalls to include: 1) can be bypassed by the use of encryption 2) are also under the control of the CSP and 3) are not designed to prevent against unknown attacks. As a result, we demonstrate the added value of using decoy documents within a SaaS environment for more dynamic data protection. We provide a prototype implementation that interoperates with a number of popular cloud storage providers and offer supporting analysis results. In the future, we shall provide.

## FUTURE WORK

In future work, we will provide a prototype implementation of the design in the section describing our prototype, Figure 1. The goal of our prototype implementation will be to demonstrate the viability of our solution across a number of different Cloud Storage Providers (CSPs). We will present implementations for Dropbox<sup>4</sup>, Google Drive<sup>5</sup>, and Box<sup>6</sup> CSPs. We will leverage a number of 3rd party APIs to enable the key features for our prototype: kloudless, Extensible Metadata Platform (XMP) toolkit, and libffx. For each CSP, the implementation focus will be on 1) Generation of Decoy Documents, 2) Detection and Prevention of Data Exfiltration Attacks, 3) Adjusting the Protection Levels Based on Threat Conditions.

Additionally, we plan to integrate several threat database sources to provide higher fidelity information prior to making a data protection adjustment (e.g. threatconnect, threatscape, enigmadatabase). Additionally, we will identify other actions that could be taken upon notification of an increased threat level (e.g., increase the ratio of decoy objects in the environment or impose a data access timeout period). As a proof of concept, we focused on the action that was most effective in defending against data exfiltration attempts. Addressing other stages within the cyber kill chain model would be a natural evolution of our

---

<sup>3</sup> Strategic Command Directive (SD) 527-1 (2006-01-27). "Department of Defense (DOD) Information Operations Condition (INFOCON) System Procedures" (PDF). DISA Policy and Guidance. Retrieved 2016-04-03. alertness. Any compromised systems are isolated from the rest of the network.

<sup>4</sup> [www.dropbox.com](http://www.dropbox.com)

<sup>5</sup> [www.google.com/drive/](http://www.google.com/drive/)

<sup>6</sup> [www.box.com](http://www.box.com)

concepts. The goal would be to address different aspects of the user experience in the cloud (from a security perspective) with various cyber deception artifacts. In this work we have implicitly focused on mitigating the ‘action on objectives’ stage of the kill chain process - the last phase of the kill chain. Future

work would seek to address the Reconnaissance, Weaponization, Delivery, Exploitation, Installation, and Command & Control phases.



## REFERENCES

- [1] Grobauer, Bernd, and Thomas Schreck. "Towards incident handling in the cloud: challenges and approaches." Proceedings of the 2010 ACM workshop on Cloud computing security workshop. ACM, 2010.
- [2] Yassin, Warusia, et al. "A cloud-based intrusion detection service framework. Cyber Security, Cyber Warfare and Digital Forensic (CyberSec), 2012 International Conference on. IEEE, 2012.
- [3] Modi, Chirag, et al. "A novel framework for intrusion detection in cloud. Proceedings of the Fifth International Conference on Security of Information and Networks. ACM, 2012.
- [4] Kholidy, Hisham, and Fabrizio Baiardi. "CIDS: a framework for intrusion detection in cloud systems. Information Technology: New Generations (ITNG), 2012 Ninth International Conference on. IEEE, 2012.
- [5] He, Sijin, et al. "Cloud resource monitoring for intrusion detection. Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on. Vol. 2. IEEE, 2013.
- [6] Tse, Daniel WK, et al. "DATA DRIVEN DETECTION STRATEGY ENGINE FOR BETTER INTRUSION DETECTION ON CLOUD COMPUTING. (2014).
- [7] Modi, Chirag, and Dhiran Patel. "A Novel hybrid-Network Intrusion Detection System in Cloud Computing. (2013).
- [8] Alharkan, Turki, and Patrick Martin. "IdsaaS: Intrusion detection system as a service in public clouds. Proceedings of the 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (ccgrid 2012). IEEE Computer Society, 2012.
- [9] Ficco, Massimo, Luca Tasquier, and Rocco Aversa. "Intrusion detection in cloud computing. P2P, Parallel, Grid, Cloud and Internet Computing (3PGCIC), 2013 Eighth International Conference on. IEEE, 2013.
- [10] Ficco, Massimo, Salvatore Venticinque, and Beniamino Di Martino. "Mosaic-based intrusion detection framework for cloud computing. On the Move to Meaningful Internet Systems: OTM 2012. Springer Berlin Heidelberg, 2012. 628-644.
- [11] Lo, Chi-Chun, Chun-Chieh Huang, and Joy Ku. "A cooperative intrusion detection system framework for cloud computing networks. Parallel processing workshops (ICPPW), 2010 39th international conference on. IEEE, 2010.
- [12] Zargar, Saman Taghavi, Hassan Takabi, and James BD Joshi. "DCDIDP: A distributed, collaborative, and data-driven intrusion detection and prevention framework for cloud computing environments. Collaborative Computing: Networking, Applications and Worksharing (Collaborate- Com), 2011 7th International Conference on. IEEE, 2011.
- [13] Ko, Ryan KL, et al. "TrustCloud: A framework for accountability and trust in cloud computing. Services (SERVICES), 2011 IEEE World Congress on. IEEE, 2011.
- [14] Gul, Irfan, and M. Hussain. "Distributed cloud intrusion detection model. International Journal of Advanced Science and Technology 34 (2011): 71-82.
- [15] Dhage, Sudhir N., and B. B. Meshram. "Intrusion detection system in cloud computing environment. International Journal of Cloud Computing 1.2-3 (2012): 261-282.
- [16] Alharkan, Turki, and Patrick Martin. "IdsaaS: Intrusion detection system as a service in public clouds." Proceedings of the 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (ccgrid 2012). IEEE Computer Society, 2012.
- [17] Wang, Huaibin, Haiyun Zhou, and Chungong Wang. "Virtual machine-based intrusion detection system framework in cloud computing environment." Journal of Computers 7.10 (2012): 2397-2403.
- [18] Bharadwaja, Saketh, et al. "Collabra: a xen hypervisor based collaborative intrusion detection system. Information technology: New generations (ITNG), 2011 eighth international conference on. IEEE, 2011.
- [19] Wei, Jinpeng, et al. "Managing security of virtual machine images in a cloud environment. Proceedings of the 2009 ACM workshop on Cloud computing security. ACM, 2009.
- [20] Nikolai, Jason, and Yong Wang. "Hypervisor-based cloud intrusion detection system." Computing, Networking and Communications (ICNC), 2014 International Conference on. IEEE, 2014.
- [21] Baraka, Hifaa Bait, and Huaglory Tianfield. "Intrusion Detection System for Cloud Environment." Proceedings of the 7th International Conference on Security of Information and Networks. ACM, 2014.
- [22] Gupta, Sanchika, Padam Kumar, and Ajith Abraham. "A profile based network intrusion detection and prevention system for securing cloud environment." International Journal of Distributed Sensor Networks 2013 (2013).
- [23] Gupta, Swastik, et al. "A fingerprinting system calls approach for intrusion detection in a cloud environment." Computational aspects of social networks (CASoN), 2012 fourth international conference on. IEEE, 2012.
- [24] Arajo, Josenilson Dias, et al. "EICIDS-Elastic and Internal Cloud-based Intrusion Detection System." International Journal of Communication Networks and Information Security (IJCNIS) 7.1 (2015).
- [25] Li, Bo, Jianxin Li, and Lu Liu. "CloudMon: a resourceefficient IaaS cloud monitoring system based on networked intrusion detection system virtual appliances." Concurrency and Computation: Practice and Experience (2013).
- [26] Zhang, Yinqian, et al. "Homealone: Co-residency detection in the cloud via side-channel analysis. Security and Privacy (SP), 2011 IEEE Symposium on. IEEE, 2011.
- [27] Arshad, Junaid, Paul Townend, and Jie Xu. "An automatic intrusion diagnosis approach for clouds." International Journal of Automation and Computing 8.3 (2011): 286-296.
- [28] Schneier, Bruce. "Schneier on Security." Blog. N.p., 23 Dec. 2015. Web. 24 Dec. 2015.
- [29] Blogger, R. (2012, January 17). Top 10 Common Uses for the Cloud for 2012. Retrieved July 30, 2015.
- [30] Wilson, Duane C., and Giuseppe Ateniese. "To Share or not to Share in Client-Side Encrypted Clouds." Information Security. Springer International Publishing, 2014. 401-412.
- [31] Stolfo, Salvatore J., Malek Ben Salem, and Angelos D. Keromytis. "Fog computing: Mitigating insider data theft attacks in the cloud. Security and Privacy Workshops (SPW), 2012 IEEE Symposium on. IEEE, 2012.
- [32] Nascimento, Gustavo, and Miguel Correia. "Anomaly-based intrusion detection in software as a service. Dependable Systems and Networks Workshops (DSN-W), 2011 IEEE/IFIP 41st International Conference on. IEEE, 2011.
- [33] Vieira, Kleber, et al. "Intrusion detection for grid and cloud computing. It Professional 4 (2009): 38-43.
- [34] Borisaniya, Bhavesh, et al. "Incorporating honeypot for intrusion detection in cloud infrastructure. Trust Management VI. Springer Berlin Heidelberg, 2012. 84-96.
- [35] Stoll, Cliff. The cuckoo's egg: tracking a spy through the maze of computer espionage. Simon and Schuster, 2005.
- [36] "Dropbox." Dropbox. N.p., n.d. Web. 24 Dec. 2015.
- [37] "Google Drive - Cloud Storage and File Backup for Photos, Docs and More." Google Drive. N.p., n.d. Web. 24 Dec. 2015.
- [38] "Back up and Access Your Files. Using Your Existing Folder Structure." Cloud File Sharing, File Sync and Online Backup From Any Device. N.p., n.d. Web. 24 Dec. 2015.
- [39] "Enterprise-Grade Security, Visibility and Control — Box." Box. N.p., n.d. Web. 24 Dec. 2015.
- [40] "Cloud Monitoring — Database Activity Monitoring — CipherCloud." CipherCloud. N.p., n.d. Web. 24 Dec. 2015.
- [41] Roth, Greg, and Don Bailey. "Intrusion Detection in the Cloud." AWS Re:Invent. November 13, 2013. Accessed April 1, 2015.
- [42] Papagiannis, Ioannis, and Peter Pietzuch. "Cloudfilter: practical control of sensitive data propagation to the cloud. Proceedings of the 2012 ACM Workshop on Cloud computing security workshop. ACM, 2012.
- [43] Almeshkeh, Mohammed H., and Eugene H. Spafford. "Planning and integrating deception into computer security defenses." Proceedings of the 2014 workshop on New Security Paradigms Workshop. ACM, 2014.
- [44] Chandola, Varun, Arindam Banerjee, and Vipin Kumar. "Anomaly detection: A survey." ACM computing surveys (CSUR) 41.3 (2009): 15.
- [45] Finkel, Jenny Rose, Trond Grenager, and Christopher Manning. "Incorporating non-local information into information extraction systems by gibbs sampling." Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics. Association for Computational Linguistics, 2005.
- [46] Yuill, Jim, et al. "Honeyfiles: deceptive files for intrusion detection." Information Assurance Workshop, 2004. Proceedings from the Fifth Annual IEEE SMC. IEEE, 2004.
- [47] Wang, Wei, et al. "Catching the wily hacker: A multilayer deception system." Sarnoff Symposium (SARNOFF), 2012 35th IEEE. IEEE, 2012.
- [48] Bowen, Brian M., et al. Baiting inside attackers using decoy documents. Springer Berlin Heidelberg, 2009.
- [49] Voris, Jonathan, et al. "Fox in the trap: thwarting masqueraders via automated decoy document deployment." Proceedings of the Eighth European Workshop on System Security. ACM, 2015.
- [50] Green, Matthew. "Format Preserving Encryption, Or, How to Encrypt a Credit Card Number with AES." A Few Thoughts on Cryptographic Engineering. N.p., 10 Nov. 2011. Web. 24 Dec. 2015.
- [51] Challener, David, et al. A practical guide to trusted computing. Pearson Education, 2007.