July 2014 UILU-ENG-14-2202

AN ACTOR-CENTRIC, ASSET-BASED MONITOR DEPLOYMENT MODEL FOR CLOUD COMPUTING

Uttam Thakore, Gabriel A. Weaver, and William H. Sanders

Coordinated Science Laboratory
1308 West Main Street, Urbana, IL 61801
University of Illinois at Urbana-Champaign

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	July 2014	

4. TITLE AND SUBTITLE
An Actor-Centric, Asset-Based Monitor Deployment Model for Cloud Computing

5. FUNDING NUMBERS
FA8750-11-2-0084

6. AUTHOR(S)
Uttam Thakore, Gabriel A. Weaver, and William H. Sanders

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, 1308 W. Main St., Urbana, IL, 61801

8. PERFORMING ORGANIZATION REPORT NUMBER
UILU-ENG-14-2202

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory & Air Force Office of Scientific Research:
USAF, AFMC, Air Force Research Laboratory, 26 Electronic Parkway, Rome, NY 13441-4514

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

Effective monitoring is essential for the security of cloud systems. Although many monitoring tools exist in the cloud domain, there is little guidance on how to deploy monitors to make the most of collected monitor data and increase the likelihood of detecting breaches of security. We introduce an actor-centric, asset-based monitor deployment model for the cloud that enables practitioners to reason about monitor deployment in terms of the security of the cloud assets that they own. We define an actor model that consolidates several roles in the literature to three roles that are motivated by security. We then develop an architectural model that identifies the assets that can be owned by each of those actors, and use it to drive an asset-based cloud threat model. Using our threat model, we claim that a cloud practitioner can reason about monitor deployment to more efficiently deploy monitors and increase its chances of detecting intrusions. We demonstrate the utility of our model with a cloud scenario based on Netflix's use of Amazon Web Services.

14. SUBJECT TERMS
Cloud computing; Security; Monitor deployment; Monitoring; Threat model; Actor model; Modeling

15. NUMBER OF PAGES
10

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

An Actor-Centric, Asset-Based Monitor Deployment Model for Cloud Computing

Uttam Thakore, Gabriel A. Weaver, and William H. Sanders
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign
Urbana, Illinois 61801
{thakore1, gweaver, whs}@illinois.edu

***Abstract*—Effective monitoring is essential for the security of cloud systems. Although many monitoring tools exist in the cloud domain, there is little guidance on how to deploy monitors to make the most of collected monitor data and increase the likelihood of detecting breaches of security. We introduce an actor-centric, asset-based monitor deployment model for the cloud that enables practitioners to reason about monitor deployment in terms of the security of the cloud assets that they own. We define an actor model that consolidates several roles in the literature to three roles that are motivated by security. We then develop an architectural model that identifies the assets that can be owned by each of those actors, and use it to drive an asset-based cloud threat model. Using our threat model, we claim that a cloud practitioner can reason about monitor deployment to more efficiently deploy monitors and increase its chances of detecting intrusions. We demonstrate the utility of our model with a cloud scenario based on Netflix's use of Amazon Web Services.**

I. Introduction

Cloud computing provides consumers with storage and computation services that can scale and are charged according to demand. However, cloud computing also introduces unique security risks, because the cloud is a shared computational resource that is owned and maintained by a variety of actors. According to a recent survey of cloud practitioners conducted by Intel, despite continuing progress in cloud computing security, 57% of respondents marked the inability to measure security services as one of their top three security concerns with cloud computing [1]. Effective monitor deployment is a must to effectively detect security breaches in a cloud system and allow cloud consumers to be confident in the security of the cloud system on which their data and services reside.

Therefore, we explain the motivation behind, formalize, and illustrate the application of a monitor deployment model for the cloud. Our intent is for cloud practitioners to use our model to reason about monitor deployment over the lifetime of a cloud service. One requirement is to enable practitioners to detect breaches of security for cloud assets that span organizational boundaries.

Our monitor deployment model is actor-centric and asset-based because it incorporates both a set of actors and a generic architecture model. In addition, our threat model is specific to the cloud and allows us to reason about the impact of threats at multiple architectural layers.

We explain the motivation for our research in the context of the limitations of existing approaches to cloud monitoring deployment and threat modeling. Existing approaches to cloud monitoring deployment are either general but too high-level to give practical advice on how to deploy a monitoring system, or provide practical advice on how to deploy a monitoring system, but only in very specific scenarios. Furthermore, existing approaches to threat modeling in the cloud may provide a high-level overview of threats, but do not provide practitioners with practical insights on which threats are relevant to them or their assets.

In contrast, our approach uses a simple actor model to express threats relative to the cloud assets owned by each actor. Our intended research contribution is to address gaps in cloud threat modeling to improve cloud monitoring. Thus, practitioners will be able to determine where they need to place monitors (or even potentially cooperate with one another) in order to mitigate a security threat.

The design and formalization of our framework were inspired by the layer-based approach to cloud monitoring taken by Spring in [2] and [3] and by the understanding that cloud monitors must monitor either the behavior of assets in the cloud or the communication interface between assets.

Our intent is to provide a general framework that practitioners can use to perform an analysis of their set of cloud assets and deploy monitors to detect security breaches. We design an approach that is both actor-centric and asset-based and is driven by the threats to the assets. We identify a set of roles in cloud scenarios so that practitioners in those roles may easily evaluate security threats relative to their own cloud assets. We also identify the set of assets upon which monitors can be deployed and the threats relevant to each asset. The intent of our monitor deployment model is to identify threats in a manner that accommodates a variety of different technologies used to implement cloud services.

The remainder of this paper is organized as follows. Section II explains the motivation for a threat model that is specific to cloud computing, as well as the need to reason about the impact and relevance of the included threats. Section III defines the three components of our approach—the actor model, architecture model, and threat model—and describes how we developed an approach to monitor deployment in the cloud using these components. In Section IV, we describe how we apply our model to analyze security-driven monitor deployment in an artificial but realistic cloud architecture based upon Netflix's use of Amazon Web Services. Section V discusses future work. We conclude in Section VI.

II. Motivation

As organizations increasingly migrate their services to the cloud to leverage economies of scale, these organizations also increase their services' attack surface. The intent of our ongoing research is to use monitor deployment, driven by a set of threats to a particular actor's cloud resources, as a tool to improve the security of the cloud. We now explain how our research responds to the limitations of monitor deployment and threat modeling in the context of the cloud.

A. Cloud Monitor Deployment

Our research addresses a gap that we see in the current cloud monitoring literature. We claim that the current literature provides either (1) a general, high-level discussion of cloud monitoring that does not help practitioners *implement* monitor deployment in specific instances, or (2) practical implementation advice that is limited to specific cloud scenarios and is not generally applicable.

High-Level Cloud Monitor Deployment: Existing research does provide high-level analysis of where to deploy monitors within the cloud. For example, Spring [2] uses the 7 layers of cloud computing defined by the Cloud Security Alliance to organize monitor placement. Although the paper provides a good list of considerations and best practices for deploying monitors and access control mechanisms to secure cloud offerings, several practical monitoring issues fall outside its scope. For example, the paper neglects actors other than cloud providers that also use a cloud. As a result, the paper cannot be applied to multi-level scenarios, such as those in which a Software-as-a-Service (SaaS) provider does not actually own the cloud infrastructure.

In addition, several surveys provide overviews of the state-of-the-art in cloud monitoring. For example, Aceto et al. explain cloud monitoring and provide an extensive survey of the current state-of-the-art of comprehensive cloud monitoring systems [4]. They identify a set of desired properties for monitoring systems that focus primarily on performance and quality of service and evaluate different commercially available cloud monitoring software services and platforms using these properties. The properties Aceto et al. use, however, do not translate easily to desired security properties, which makes their evaluation less useful for monitor deployment for intrusion detection. Furthermore, their analysis is top-down and looks only at complete monitoring systems; no consideration is given to constructing good monitoring systems using individual monitors or combinations of monitors and monitoring systems. Our proposed approach provides a bottom-up method for deploying monitors within and around an actor's cloud assets to increase the utility of the monitors in detecting security breaches.

Monitoring Approaches for Specific Cloud Scenarios: In contrast, a number of approaches have been proposed for cloud monitoring in very specific scenarios. Clayman et al. propose a new monitoring framework for RESERVOIR federated service clouds [5]. Khandelwal et al. use a set of indirect information to attempt to monitor network health information in Amazon EC2 infrastructure clouds [6]. González et al. propose a specific monitoring architecture for dynamic security monitoring in virtualized environments [7]. While those approaches may have some merit in the specific scenarios they address, they are not generally applicable.

B. Threat Modeling in the Cloud

Practitioners need to be able to model threats in the cloud. Our ongoing research addresses a gap in threat modeling for the cloud: the lack of a practical way for cloud practitioners to evaluate the threats that are relevant to their assets. This is important because understanding where to deploy monitors in the cloud to maximize the probability of detecting intrusions requires a detailed understanding of how general security issues and threats relate to the specific assets one owns and needs to protect. As with cloud monitoring, many of the cloud-based threat models are useful, but somewhat high-level. In addition, to the best of our knowledge, there is no cloud-based threat model that clearly associates security issues with the parties they affect.

A variety of organizations and governmental agencies have come out with their own high-level evaluations of threats to cloud computing and recommendations to mitigate these threats (see reports by ENISA [8], Gartner [9], and NIST [10]). Much of the work in the area has been done by the Cloud Security Alliance (CSA) in their "Security Guidance" [11] and "Top Threats to Cloud Computing" [12] publications. Although the CSA and others classify threats in terms of different cloud security models, they do not give practitioners an easy way to understand which threats affect the specific assets that they own or depend upon.

Finally, although many researchers have surveyed security threats within cloud computing, they have not yet expressed these threats in terms of how individual practitioners, or actors, are affected. Subashini and Kavitha survey security issues and relate them to each of the three prominent service models (IaaS, PaaS, and SaaS) [13]. Modi et al. identify security issues and mechanisms to mitigate them in terms of 7 layers in cloud computing that they developed [14]. However, neither of these papers provides a clear way to map the cloud assets owned by different actors to the service models or layers, which makes it difficult to associate security issues with the parties they affect.

III. Cloud Monitor Deployment Model

We approach monitor deployment in the cloud from an actor-centric and asset-based perspective.

A. Cloud Actor Model

We distill the roles defined in the literature into the following three primary roles [4], [15].

- *Cloud provider (CP)*: Provider and owner of the physical infrastructure. The service provided by the CP may be infrastructure (as in the case of EC2), platform (as in the case of Azure or AppEngine), or even software (as in the case of SalesForce), so long as the CP does not actually own the data or computation that is run through its cloud offering.
- *Cloud service provider (CSP)*: User of a cloud provider, but also a provider of a service offering to consumers. The hallmark of a CSP is that it is not the sole end-user of the cloud software chain, and passes on some of the risk to a

cloud service consumer. The CSP does not have control over the physical cloud infrastructure it uses; for that, it relies on a service offered by the CP.

- *Cloud service consumer (CSC)*: End-user of a cloud service. This party only consumes the cloud service and has no control over the security of the infrastructure or software behind the service.

Initial Evaluation: The roles we define are independent of the service model. A single actor may take on multiple roles. For example, a software development organization may use Amazon EC2 to host its SaaS offering and thus act as a CSP, but it may also use EC2 virtual machines for testing purposes, thus taking the role of a CSC. Furthermore, it is possible that in a given scenario, there is no CSP (as in the case of a software development organization that uses EC2 instances as test machines) or there are multiple CSPs, potentially themselves acting as CSCs of other CSPs. The roles we define are sufficiently flexible to classify actors in a variety of cloud scenarios.

Our roles are similar to the "users" and "providers" described in [16] and hold true to the model of the cloud presented in that paper. Armbrust et al. [16] define a cloud provider as an actor that makes cloud infrastructure available as a utility computing service. Amazon Web Services is commonly classified as an IaaS provider and Microsoft Azure and Google AppEngine are commonly classified as PaaS providers; Armbrust et al. simply classify all three as cloud providers. By their definition, the fundamental characteristic of a cloud provider is the ownership and provision of the cloud infrastructure in a pay-as-you-go manner, which is the same definition we use for the CP role. In their model, cloud infrastructures are used by SaaS providers, which in turn provide services to SaaS users. We expand the roles defined in their model in light of newer use cases for clouds to allow CSPs to provide a service from any service model and to allow CSCs to be direct users of CPs or CSPs, irrespective of the service model. However, we stay true to the understanding behind the roles, in that a CSP does not own the cloud infrastructure and a CSC is an end-user of cloud services.

B. Cloud Architecture

In order to understand where to deploy monitors, it is first necessary to identify the assets in the cloud that can be monitored. To ensure that we identify the assets that are commonly found in public clouds, we examine the cloud layer by layer. Based on the work of the Cloud Security Alliance, Spring identifies 7 layers of cloud computing within which cloud assets can exist [2]. We have extended these layers by separating the management and security infrastructure from the other layers. The layers are listed below in bold, and their assets are in italics. Figure 1 illustrates our cloud architectural model.

- **Facility**: The physical infrastructure of the datacenters that support the computing and networking hardware. This layer includes the *buildings*, *power supplies and backup power infrastructure*, *control hardware and software*, and *physical security hardware and software*, such as CCTVs and biometric authentication devices.

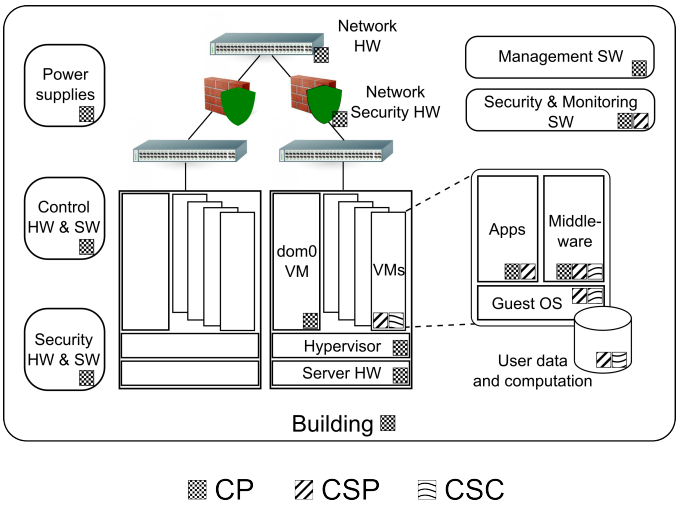

Fig. 1: Cloud architecture as reflected in our monitor deployment model. Shaded boxes represent ownership.

- **Network**: Within the datacenters, the *network hardware* (e.g., switches, routers, and ToR switches) that routes traffic within the datacenter. *Network security hardware*, such as dedicated firewalls, IPS boxes, IDS boxes, and network proxies, also falls in this layer. While the Internet networking hardware responsible for routing traffic between availability zones and between the datacenter and consumers also falls within this layer, it is generally not under the direct control of any of the cloud actors, so we have excluded it from the list of assets.
- **Hardware**: Also within the datacenters, the *servers* that contain the physical computing and storage hardware, which include memory, CPU, HDD, and other peripherals.
- **Operating System**: The operating system and virtualization software (virtual machine monitor), which include the *hypervisors* and *administrave domain (dom0)* virtual machines running on the servers themselves and the operating systems running on the *client virtual machines*.
- **Middleware**: The *middleware* located on the virtualization software and operating systems running on the client VMs. Middleware is software that interacts with cloud applications and the virtual machine operating systems. An example of middleware might be an enterprise service bus that allows multiple cloud apps to communicate with one another.
- **Application**: *Software and services* that run on the cloud infrastructure.
- **User**: The end-user of cloud services. As the CSC and its assets are outside the cloud, we exclude them from our list of cloud assets. However, user *data* and *computation* promised to a user by a Service Level Agreement (SLA) both reside on the cloud and are assets that must be protected.
- **Management & Security**: The software that is not visible to the CSPs and CSCs but is nonetheless executed by the CP to deploy, manage, meter, and protect the cloud infrastructure. Examples of software in this layer include *cloud management and orchestration software* and *auditing, monitoring, and security software*. An attacker who is

TABLE I: Table of cloud asset ownership.

Asset	Owned by	Layer
Buildings	CP	Facility
Power supplies and backup power infrastructure	CP	Facility
Control hardware and software	CP	Facility
Physical security hardware and software	CP	Facility
Network hardware	CP	Network
Network security hardware	CP	Network
Servers	CP	Hardware
Hypervisor	CP	Operating System
Administrative domain (dom0)	CP	Operating System
Client VMs	CSP or CSC	Operating System
Middleware	CP, CSP, or CSC	Middleware
Software and services	CP or CSP	Application
Data	CSP or CSC	User
Computation	CSP or CSC	User
Cloud management and orchestration software	CP	Management & Security
Auditing software	CP or CSP	Management & Security
Monitoring software	CP or CSP	Management & Security
Security software	CP or CSP	Management & Security

aware of vulnerabilities in the management and security software can cause loss of availability, confidentiality, or integrity of the CP's infrastructure and its consumers' data and services without needing to compromise assets in other layers.

Further, in order to reason about monitor placement within assets, we must identify for each actor the set of assets within which the actor can independently place monitors (the assets it "owns"). The relationship is important because actors can directly deploy monitors only within the assets they control through ownership. While it is sometimes possible for an actor to monitor assets outside of its control through agreements with other actors, we focus primarily on direct monitor deployment. We briefly examine monitoring of unowned assets in Section III-E.

In Table I, we identify the owners of cloud assets and the layers in which they reside.

C. Cloud Threat Model

We motivate monitor deployment on cloud assets as a mechanism to detect security breaches on those assets. To aid in detection of security breaches, we first identify the threats to the assets that each actor owns. Understanding which threats are relevant to an actor's assets can help the actor decide where deploying monitors will most increase its chance of detecting intrusions.

It is important to note that we do not attempt to identify the specific vulnerabilities in the cloud assets, as these vary among different hardware and software platforms and change with the release of software patches and new versions. Practitioners can use an understanding of the generic threats to their assets to identify the specific vulnerabilities in their own instances of the cloud, and can use this knowledge to drive the placement and tuning of monitors in the cloud.

We now discuss security objectives for each of the actors in our actor model in relation to the cloud assets that they own.

1) Cloud Provider: As the owner of the physical infrastrucure, the CP assumes all responsibility for its protection and availability. The CP must ensure that the physical facilities are available as per its SLAs, that access to the facilities remains restricted to authorized parties, and that the infrastructure within the facilities is protected from malicious and unintentionally harmful activity. The CP must ensure the confidentiality, integrity, and availability of its network hardware, network security hardware, and servers, ensure that network traffic within the datacenter is not a source of data leakage or attack, and ensure that access to computing resources is fairly distributed among co-tenants.

The CP is also responsible for the protection of the virtual assets it controls. It must ensure that the hypervisor and administrative domain VM responsible for the virtualization of cloud hardware are protected from attack and that they do not behave maliciously as a result of compromise. If the CP is providing its consumers with access to middleware applications, it must also be concerned with the security of the middleware applications. Furthermore, as the CP is responsible for the physical storage of consumer data, it is partially responsible for the protection of the data against leakage and usurpation by malicious parties. It is also responsible for the integrity and availability of the data.

In addition to its responsibilities to consumers, the CP must ensure that it complies with federal regulations and laws. That means it must audit the use of its cloud infrastructure, monitor the cloud for illicit activity, and perform measures to ensure data confidentiality and integrity as required by data privacy laws. As the customers of the CP may be CSPs with their own obligations to abide by regulations and laws, the CP must also be able to provide audit information to its customers on demand.

Table II lists the threats to each asset owned by the CP that affect the CP's security goals.

2) Cloud Service Provider: Since it cannot fully trust the cloud infrastructure, the CSP must secure its resources from attack from both outside and inside the cloud. It must ensure that any operating system it is using is secured from attack, both from outside attacks and from compromised virtualization infrastructure. The CSP must also ensure that its middleware, software, and services are available to its consumers, are protected from attack or loss of integrity, and do not behave maliciously. In addition, the CSP must ensure that its customers get fair use of the software (e.g., it must prevent denial of service or resource hogging).

In addition to its interest in the security of its own resources on the CP's infrastructure, the CSP must also be concerned with the security of its clients' data and computation performed on the cloud. The CSP must ensure the security of its consumers' data from malicious parties both inside and outside the cloud. That requires that the CSP take necessary precautions to prevent threats due to vulnerabilities in the CP's data storage mechanisms.

As a business entity, the CSP, like the CP, must ensure that it complies with laws and regulations, particularly in its handling of customer data. Additionally, the CSP is interested in ensuring that its SLA with the CP is being met. Thus, the CSP must monitor its usage of cloud services, using both

TABLE II: Threats to the cloud provider's assets.

Asset	List of threats
Buildings	Natural disasters, man-made disasters (such as terrorism), power outages, unauthorized parties' gaining of access to the datacenter facilities.
Power supplies and backup power infrastructure	Power outages, bugs in failover code leading to improper response to power failures.
Physical security and control hardware and software	Malicious insider tampering with the hardware, failure of or damage to hardware, exploitation of vulnerabilities in the software or hardware.
Network and network security hardware	External (D)DoS by way of network traffic blackholing or flooding, exploitation of vulnerabilities in networking hardware code, unauthorized access to or malicious insider tampering with the hardware, improper physical care of hardware (e.g., poor cooling), failure of hardware.
Servers	Exploitation of bugs in cloud management software, unauthorized access to or malicious insider tampering with the hardware, malicious hardware controller software, improper physical care of hardware (e.g., poor cooling), failure of hardware.
Hypervisor	Exploitation of vulnerabilities in hypervisor code that can be exploited by tenant VMs, exploitation of covert channels.
Administrative domain	Exploitation of vulnerabilities in dom0 OS, DoS of administrative functionality.
Middleware	Exploitation of vulnerabilities in middleware code, improper security mechanisms used by developers of middleware, malicious or deceptive middleware.
Client data	Unauthorized access to data by malicious insiders, unauthorized access to data by outsiders who have taken control of data storage hardware, data leakage due to improper destruction of old hardware, PATRIOT Act, confiscation of storage hardware by the government, compromise of data availability by DoS, corruption of data due to faulty hardware, data leakage or corruption during transit through the cloud infrastructure.
Client computation	Theft of encryption keys stored in memory, leakage of client activity profiles through side-channels, unauthorized access to client computation information through control of compute hardware (e.g., Blue Pill hypervisor rootkit) [17].
Cloud management, orchestration, and security software	Exploitation of predictability in scheduling policies [18], externally visible covert channels, bugs in cloud management and security software that could result in data leakage, unavailability of management software, downtime of cloud security or auditing software (which could result in violations of laws/regulations or high risk to cloud infrastructure).

information acquired from the CP and its own observations, to determine whether the CP is violating its SLA, and if it is, whether to maintain its relationship with the CP.

Table III lists the threats to each asset owned by the CSP that affect the CSP's security goals.

3) Cloud Service Consumer: If a CSC acts as a direct consumer of a CP's services, it inherits some of the threats associated with the CSP. If it has control over the operating system being run on a client VM, it is responsible for the security of the operating system against attack, both from outside attacks and from compromised virtualization infrastructure. If it controls or uses middleware, it must be concerned with the secure execution of said applications.

As a consumer of a cloud service, the CSC's primary roles in security are to audit its provider to confirm adherence to the SLA, and to verify the security of its own data and computation as they are sent to and processed in the cloud. It can take steps to protect the data in transfer and to ensure that the CP or CSP is using encryption and access control to prevent unauthorized use or tampering with the data on disk.

Table IV lists the threats to each asset owned by the CSC that affect the CSC's security goals.

D. Security-driven Monitor Deployment

We now motivate monitor deployment using our actor-centric and asset-based cloud threat model.

Monitor deployment may have economic and computational resource costs, so it must be done efficiently. It may cost money to purchase or support the monitoring software or hardware, especially in the case of dedicated monitoring hardware (such as an online network traffic monitor). Monitors use computational resources that have a nontrivial power cost, potential performance hit to production infrastructure. Furthermore, the amount of data collected by monitors is substantial, and there is a cost to store and analyze the data. In addition, monitor deployment can increase the attack surface of a system. Therefore, practitioners must selectively deploy monitors to maximize their utility and minimize the risk introduced by monitor deployment. We discuss considerations that should be made in the following steps.

Step 1: The first step a practitioner should take in using our monitor deployment model is to enumerate all assets under its control and identify the specific versions of each asset. To the extent possible, the practitioner should also identify the versions of assets being used by its providers and consumers, as awareness of threats to other actors can motivate deployment of monitors. The practitioner should then figure out its system architecture (i.e., the connections and relationships between its assets and those of its providers and consumers), as a compromised asset can be used as a platform for further attacks.

Step 2: Once a practitioner has created a detailed picture of its cloud system architecture, it should examine the list of threats in the threat model in Section III-C and determine which threats are applicable to its assets. For example, if a CP utilizes network hardware that has known vulnerabilities, then a threat that exploits one of those vulnerabilities may be important and warrant targeted monitoring. Further, if the practitioner offers a service that is known to be subject to a particular type of attack, then more monitors should be deployed to the assets that are directly affected by the attack and can detect the attack early. For example, a CSP offering software-as-a-service that is susceptible to (distributed) denial of service ((D)DoS) attacks should deploy network traffic monitors in front of its load-balancing service and analyze the monitor data for the first signs of a DoS attack.

Step 3: Practitioners should also consider the importance of an asset to the correct performance of the system or the criticality of their mission. For example, if a government agency uses the cloud to host its public website and also to store and perform mission-related big data analytics, the agency would likely consider placement of monitors on the mission-related big data processing resources to be of much

TABLE III: Threats to the cloud service provider's assets.

Asset	List of threats
Client VM	Exploitation of a vulnerability in OS code by an outside party, an application running on the VM, or a malicious hypervisor or dom0; improper allocation of physical resources by malicious or compromised VMM.
Middleware	Exploitation of a vulnerability in software code by an outside party, another application running on the VM, or a malicious hypervisor or dom0; improper security mechanisms used by developers of software; malicious or deceptive middleware or third-party software.
Software and services	Code injection, denial of service, exploitation of a vulnerability in application code, misbehavior or inadvertent damage by insiders or administrators, loss of availability due to government confiscation of hardware associated with co-tenants' illicit activity, loss of security due to malicious or compromised VMM or OS, loss of availability due to CP's loss of availability.
Client data	Improper management of storage hardware by CP, leakage of data during transit between cloud and users, leakage of data during transit within cloud, PATRIOT Act, confiscation of storage hardware by the government, leakage of data through covert channels or mismanaged co-tenancy (shared memory, unauthorized access to disks), malicious or inadvertently harmful CP or CSP insider.
Client computation	Theft of encryption keys stored in memory, leakage of client activity profiles through side-channels, unauthorized access to client's or own computation information through usurpation of provider's or own resources.
Monitoring and security software	Exploitation of predictability in usage of CP infrastructure, unavailability of management software, unavailability of security or auditing software, misinformation provided by malicious or compromised CP management and security software.

TABLE IV: Threats to the cloud service consumer's assets.

Asset	List of threats
Client VM	Exploitation of a vulnerability in OS code by an outside party, an application running on the VM, or a malicious hypervisor or dom0; improper allocation of physical resources by malicious or compromised VMM.
Middleware	Exploitation of a vulnerability in software code by an outside party, another application running on the VM, or a malicious hypervisor or dom0; improper use of security mechanisms; malicious or deceptive middleware or third-party software.
Data	Improper management of data security by CP or CSP, leakage of data from client machine, leakage of data during transit between cloud and CSC, leakage of data through covert channels.
Computation	Theft of encryption keys stored in memory, leakage of activity profiles through side-channels, unauthorized access to computation information through usurpation of provider resources.

greater importance than the security of the website. Mission-criticality analysis of assets would likely be done as part of a risk analysis on the system.

Step 4: Additionally, practitioners should consider how they will use the collected monitor data. Specifically, a practitioner should deploy a monitor only if the value of the information collected through the monitor outweighs the cost of purchasing, deploying, and maintaining the monitor and analyzing the data it collects. For example, network packet captures are an information-rich monitoring mechanism. However, they are extremely verbose and can become prohibitively expensive to store. If a CP is simply attempting to detect inter-VM traffic relationships, it would likely be better served with a ToR switch- or hypervisor-level monitor that specifically tracks the source and destination IP addresses of network flows.

Step 5: A practitioner must also consider how it can validate the data produced by a monitor. It is important to realize that the monitors themselves are subject to some of the same threats as the remainder of the cloud infrastructure, and may therefore be compromised and return bad data. For example, consider a CP that places a system call monitor within the hypervisor in its servers. If a zero-day exploit is used on the hypervisor to install a rootkit that intercepts system calls and returns fallacious results to tenant VMs, the system call monitor itself will report fallacious data. The amount of trust placed in a monitor becomes more important the higher in the cloud stack the monitor is placed, as compromise of any of the infrastructure at lower layers in the cloud can cause compromise of the monitor.

There are a variety of approaches a practitioner could employ to validate monitor data. One such approach is to deploy redundant monitors. Deploying multiple monitors that collect similar information would allow a practitioner to corroborate information collected by other monitors. Monitors should also be deployed with some level of coverage in mind. In the previous example of a compromised hypervisor, if the CP were to place redundant monitors in assets in the layers surrounding the hypervisor, such as an activity profile monitor in the server hardware, it might be able to detect the discrepancies between the reported system calls and the actual use of the hardware, and detect compromises despite the failure of the hypervisor monitor. In general, practitioners should consider placing monitors in and around as many assets in their system as possible. One primitive measure of coverage could be to ensure placement of a monitor in each layer of the cloud.

E. Monitoring Unowned Assets

Often, it is also necessary to collect monitor information from layers or assets not directly under one's ownership. There are multiple methods for performing such monitoring.

The most straightforward method for collecting monitoring information from other actors' layers is through a service provided by the other actors. For example, users of Amazon Web Services cannot directly monitor hardware-layer metrics, such as CPU utilization, latency, and request counts, as they do not have ownership of the hardware-layer assets. However, through Amazon's CloudWatch service, they can gain access to those and other CP-controlled metrics [19]. Similar services are also available for other providers [20].

Another method for such monitoring is through an agreement made with the other actors. As part of its SLAs with its providers or consumers, a practitioner could require that its providers or consumers place monitors in certain locations within their own cloud assets. While that method is contingent on agreement by the external actor, it provides further control over the exact information collected if a provider or consumer does not provide its own monitoring service.

1) Limitations of our monitor deployment model: The intent of our ongoing research is to help practitioners to deploy monitors in an efficient manner backed by practical concerns, which are discussed here, and theoretical principles, which are part of ongoing research. Our monitor deployment model allows practitioners to inventory the vulnerabilities within their cloud assets, which can motivate efficient deployment of monitors so as to increase the chances of detecting an attack.

Although our model may help practitioners deploy monitors relative to potential points of attack, our approach has several limitations. We do not claim that our model will keep a cloud system running in the presence of attacks. We claim that the effectiveness of a deployed monitor depends on the configuration of the monitor and on its ability to detect exploitation of the vulnerabilities specific to an actor's cloud assets. Furthermore, although we envision our monitoring system as a means to drive responses that support intrusion tolerance, we do not provide recommendations for these responses.

IV. Application

We illustrate the application of our monitor deployment model with a scenario that is based on Netflix's use of Amazon Web Services to host its streaming video service. Where possible, we use actual infrastructure and application information given by Netflix in [21] and [22], but we take some liberties in constructing the scenario. We also hypothesize about the threats that might be relevant to each of the actors in our scenario.

A. Scenario

In our example scenario, shown in Figures 2 and 3, a streaming video service company, which we call StreamPics, hosts its website, development operations, and streaming video service on a private cloud. In order to reduce operating costs and increase its service's availability, StreamPics is planning to migrate its services to a public cloud provided by CloudSpace, an IaaS cloud provider. To start, StreamPics will deploy its entire service in just one of CloudSpace's cloud datacenters. The actors in this scenario are CloudSpace, which is the CP; StreamPics, which is the CSP; and StreamPics's customers, which act as CSCs.

CloudSpace's datacenters have a three-tiered fat tree architecture, with top of rack (ToR) switches making up the bottommost tier. CloudSpace offers its consumers a compute service, within which consumers can deploy virtual machines, and a storage service, which implements a highly durable and available key-value object datastore. The compute servers run a Xen hypervisor with an Ubuntu LTS-based dom0, and each server can host up to eight of CloudSpace's smallest-sized virtual machines. The storage servers provide a hardened API that allows users to write, read, and delete objects. Compute clusters and storage clusters within each datacenter reside on separate racks, so communication between them passes through higher-level switches. CloudSpace additionally provides its clients with access to a usage monitoring service for their VMs and use of the storage service. Figure 2 illustrates CloudSpace's datacenter architecture.

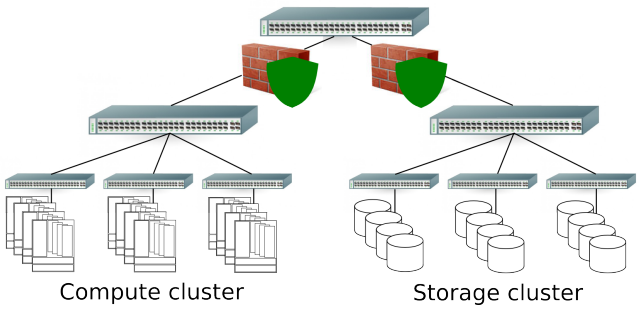

Fig. 2: Architecture of CloudSpace's cloud datacenter.

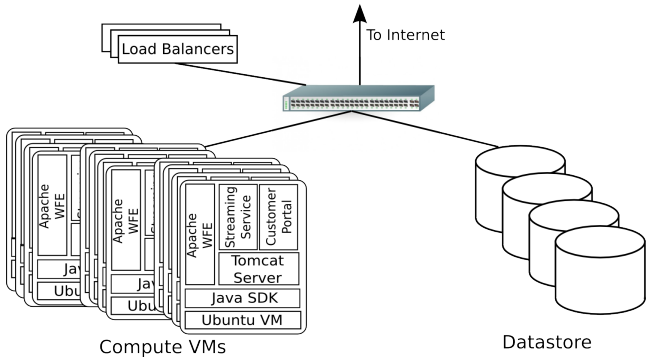

Fig. 3: Architecture of StreamPics's cloud-based streaming video service.

The storage service and compute cluster hypervisors use ACLs to meter access. CloudSpace also uses hypervisor-level firewalls to reduce the possibility of attacks by co-tenant VMs. Network traffic within the datacenter is restricted through the use of firewalls and configurable ACLs, and network traffic with the outside Internet is run through signature-based intrusion prevention systems (IPSes) that can be configured to check for simple known attacks.

StreamPics uses both CloudSpace's storage and compute services to host its frontend website and backend streaming service. All of its virtual machine images (VMIs) are a custom version of Ubuntu that prevents SSH between virtual machines, blocks all access into the virtual machine except through a specific set of Web portals (e.g., customer portal and development portal), and runs each process with a predefined user ID and security group. The VMIs are preloaded with monitoring and diagnostic software, a recent version of the Java JDK, Tomcat Server, and an Apache management frontend. No other software is allowed to run on StreamPics's virtual machines. StreamPics stores its data (video objects, logs, user authentication information, etc.) in an at-rest encrypted format using CloudSpace's storage service with unencrypted keys. It configures CloudSpace's ACLs to restrict access to the storage servers solely to compute VMs. Figure 3 illustrates StreamPics's service architecture.

StreamPics's customers access the streaming service through the customer portals running on StreamPics's virtual machines. Customers are assigned to customer portals through a set of load-balancing servers. Each customer has a user ID

and password that it uses to authenticate to the portals. Once authenticated, customers can make application requests that modify account information or stream a video to their client device. When a customer requests a video stream, the customer portal retrieves video data from the storage service, decrypts it, and sends the video through an encrypted connection to the client.

B. Application of Monitor Deployment Model

We can use our monitor deployment model to reason about monitor deployment from the perspective of all three actors in our scenario.

CloudSpace: CloudSpace owns the physical infrastructure up to the virtualization software and the software that manages the storage service. It owns the following assets:

- Facility: datacenter facilities, which include the building, generators, power supplies and failover mechanisms, physical security mechanisms, and control software.
- Network: network switches, firewalls, and IPSes.
- Hardware: storage and compute servers.
- Operating system: hypervisors and dom0 VMs running on the compute servers and host OSes running on the storage servers .
- Application: storage service software.
- User: client data, identification information, and computation.
- Management & security: VM management and orchestration software, identity and access management software, and host-level firewalls running on the compute servers.

At the facility layer, we assume that CloudSpace is primarily concerned about power failures and unauthorized access to the physical resources, and considers disasters to be of low risk. We assume that its control and power supply rooms house the majority of the security and management software, so it considers the security of these rooms to be of high importance. To increase the chances of detecting power-failure-related incidents, CloudSpace could place monitors on its generators to perform regular health checks and on its failover mechanisms to detect anomalous behavior. Since such incidents are likely to be rare, to minimize the cost of data collection and analytics, the monitors should generate alerts only upon anomalous behavior and report them directly to administrators. To increase the chances of detecting unauthorized physical access, CloudSpace could place multi-factor physical authentication devices on all datacenter entrances, and additionally place such devices on entrances to the high risk rooms. The monitors should generate alerts upon tampering, access by unauthorized parties, and anomalous access patterns. Anomaly-based monitors could also be placed on the management software to monitor for issuance of anomalous commands (such as a request for an unprecedented number of VMs) by administrators or clients.

At the network layer, we assume that CloudSpace is concerned that its network switches, some of which have a known vulnerability that can cause the switch to crash, may be attacked through the use of malicious packets. To detect and potentially mitigate attacks on its network hardware, CloudSpace could deploy signature-based IPSes in front of its core switches and on the ToR switches coming from its compute servers.

At the operating system and application layers, we assume that CloudSpace is concerned about attacks on its storage service from the outside Internet and malicious client VMs and about attempts by malicious client VMs to compromise or take control of the VMM. We assume that it has protections within the hypervisor to fairly distribute access to resources, so it is not concerned about DoS from within so long as the hypervisor is not compromised. To aid in detection of VM attempts to compromise the virtualization software, CloudSpace could deploy a syscall monitor above the hypervisor layer, and provide redundancy by deploying usage monitors on the physical servers, an access log on the requests made to manage the VMs, and, potentially, a usage monitor on the traffic to the VMs. By correlating the information collected by those, CloudSpace could identify intrusions or compromise to the VMMs by detecting mismatches in the monitor data.

At the user layer, we assume that CloudSpace is concerned about attacks on its storage service that aim to steal client data. To increase its chance of detecting attacks on its storage service, CloudSpace could log all access to the storage service. To provide additional redundancy, CloudSpace could also place IDSes on the switches connecting the storage cluster to the rest of the datacenter.

StreamPics: StreamPics owns assets at the operating system level and higher. Its assets include:

- Operating system: its client VMs and custom VMIs.
- Application and middleware: the platform and application software that comprises its portals and streaming service.
- User: the data it stores in CloudSpace's storage service.

At the operating system layer, we assume that StreamPics is concerned that an attacker might gain the ability to execute malicious software on its compute VMs through application or operating system vulnerabilities. We also assume that it is concerned about its dependency on the hypervisor owned by CloudSpace and how a malicious or compromised hypervisor could impact its compute VMs. Additionally, StreamPics is concerned about DDoS attacks on its resources, but it can rely on its SLA with CloudSpace, which stipulates that CloudSpace will provide StreamPics with over 99.99% availability. To increase its chances of detecting compromise of its compute VMs, StreamPics could utilize CloudSpace's usage monitoring service to monitor for unusual activity and deploy anomaly-based IDSes beneath its portal and streaming service. If StreamPics cannot trust the output of CloudSpace's usage monitoring service, it could also deploy application-level heartbeat monitors and usage monitors for redundancy. Further, StreamPics could periodically monitor the contents its VMs' memory to ensure that the software running in memory is only the set of applications it has expressly permitted.

At the user layer, we assume that StreamPics is concerned about unauthorized access to or corruption of the data it stores in CloudSpace's storage service. In particular, it is concerned about theft of encryption keys and client identification data and unauthorized use of video data from inside or outside the cloud, which could violate its agreements with content providers. To aid detection of such security

breaches, StreamPics could enable access logging on its use of the storage service or deploy anomaly-based monitors on its application servers. While StreamPics cannot deploy monitors within CloudSpace's storage service, it could perform sporadic attempts to access its data from anonymous, unauthorized VMs or outside hosts to verify the validity of CloudSpace's protections and access logs.

StreamPics's Customers: StreamPics's customers also own some of the resources in the user layer of the cloud. Their identification information is stored by StreamPics on CloudSpace's infrastructure, and their agreements with StreamPics provide them with a level of availability to the streaming service. In this scenario, StreamPics clients cannot deploy any monitors directly within the cloud infrastructure. However, using our recommendations, clients could use monitors provided by StreamPics to their advantage. To detect violations of the security and privacy of their identification information without StreamPics's cooperation, customers could examine their reported usage of StreamPics's service to detect anomalous usage.

V. Future Work

Our monitor deployment model provides a framework that practitioners can use to communicate and reason more effectively about monitor placement within the cloud. However, much work can still be done to build on our model.

We are investigating and formalizing a variety of metrics for monitor deployment in our ongoing research. For example, redundancy and coverage metrics could increase the trust in collected monitor data and decrease the likelihood that compromise of a subset of monitors would result in attacks going undetected.

In addition, we want to evaluate monitor placement relative to the utility of the data provided by that monitor. We want to classify the statistical data types provided by monitors and use that classification to formalize the concept of the utility of monitor data streams. We also want to survey the types of monitors that can be deployed in cloud systems, and by coupling those results with the formalization of utility and the metrics mentioned above, ultimately drive automated monitor deployment.

Finally, in addition to security, we see other valuable applications of our approach to monitor deployment. We want to investigate a similar, asset-based approach to compliance-, audit-, or performance-driven cloud monitor deployment. We believe that such domains could benefit from a formalization of the types of monitors that can be deployed and the data the monitors can collect.

VI. Conclusion

In this paper, we present an actor-centric and asset-based model for monitor deployment that we claim provides parties in the cloud with an easily actionable model for security-driven deployment of monitors in the cloud. We argue that the current literature does not provide such a model, as it focuses on the cloud service models instead of on the actors in the cloud.

To explain the motivation for monitor deployment, we define a set of three actors that are motivated by the need to secure their cloud assets. Working layer-by-layer, we identify the set of assets that exist in a cloud system and associate them with the actors that control them. We then identify the threats to those assets as seen by each actor, and provide a methodology for deploying monitors in the cloud that is based on the threat model we present. We illustrate the use of our monitor deployment model with an example scenario motivated by a real cloud use case.

In this paper, we present the results of our first steps towards a larger goal: to develop a practical monitor deployment model that is useful to practitioners seeking to satisfy their security goals.

Acknowledgment

The authors would like to thank Jenny Applequist and Carmen Cheh for their comments and revisions.

This material is based on research sponsored by the Air Force Research Laboratory and the Air Force Office of Scientific Research, under agreement number FA8750-11-2-0084. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Research Laboratory and the Air Force Office of Scientific Research, or the U.S. Government.

References

[1] Intel IT Pro Research, "Cloud security survey from Intel," p. 31, 2012. [Online]. Available: www.intel.com/content/dam/www/public/us/en/documents/reports/whats-holding-back-the-cloud-peer-research-report2.pdf

[2] J. Spring, "Monitoring cloud computing by layer, part 1," *IEEE Security & Privacy*, vol. 9, no. 2, pp. 66–68, 2011.

[3] ——, "Monitoring cloud computing by layer, part 2," *IEEE Security & Privacy*, vol. 9, no. 3, pp. 52–55, 2011.

[4] G. Aceto, A. Botta, W. De Donato, and A. Pescapè, "Cloud monitoring: A survey." *Computer Networks*, vol. 57, no. 9, pp. 2093–2115, 2013.

[5] S. Clayman, A. Galis, C. Chapman, G. Toffetti, L. Rodero-Merino, L. M. Vaquero, K. Nagin, and B. Rochwerger, "Monitoring service clouds in the future internet," in *Towards the Future Internet*. IOS Press, 2010, p. 115126. [Online]. Available: http://www.future-internet.eu/fileadmin/documents/valencia_documents/plenary/Monitoring_service_clouds_in_the_Future_Internet.pdf

[6] H. Khandelwal, R. R. Kompella, and R. Ramasubramanian, "Cloud monitoring framework," Tech. Rep., 2010. [Online]. Available: http://www.cs.purdue.edu/homes/bb/cloud/h-report.pdf

[7] J. Gonzalez, A. Munoz, and A. Mana, "Multi-layer monitoring for cloud computing," in *Proc. of 2011 IEEE 13th International Symposium on High-Assurance Systems Engineering (HASE)*, 2011, pp. 291–298.

[8] D. Catteddu and G. Hogben, "Cloud computing: Benefits, risks, and recommendations for information security," European Network and Information Security Agency (ENISA), Tech. Rep., 2009. [Online]. Available: http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport

[9] J. Brodkin, *Gartner: Seven cloud-computing security risks*, 2008. [Online]. Available: http://www.idi.ntnu.no/emner/tdt60/papers/Cloud_Computing_Security_Risk.pdf

[10] NIST Cloud Computing Security Working Group, "NIST cloud computing security reference architecture," vol. 500, p. 204, 2013. [Online]. Available: http://collaborate.nist.gov/twiki-cloud-computing/pub/CloudComputing/CloudSecurity/NIST_Security_Reference_Architecture_2013.05.15_v1.0.pdf

[11] Cloud Security Alliance, "Security guidance for critical areas of focus in cloud computing v3.0," Cloud Security Alliance, Tech. Rep., 2011. [Online]. Available: https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf

[12] D. Hubbard and M. Sutton, "Top threats to cloud computing v1.0," Cloud Security Alliance, Tech. Rep., 2010.

[13] S. Subashini and V. Kavitha, "A survey on security issues in service delivery models of cloud computing," *Journal of Network and Computer Applications*, vol. 34, no. 1, pp. 1–11, 2011.

[14] C. Modi, D. Patel, B. Borisaniya, A. Patel, and M. Rajarajan, "A survey on security issues and solutions at different layers of cloud computing," *The Journal of Supercomputing*, vol. 63, no. 2, pp. 561–592, 2013.

[15] US Department of Commerce and NIST, "Important actors for public clouds," Nov. 2010, important Actors for Public Clouds. [Online]. Available: http://www.nist.gov/itl/cloud/actors.cfm

[16] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "A view of cloud computing," *Communications of the ACM*, vol. 53, no. 4, p. 5058, Apr. 2010. [Online]. Available: http://doi.acm.org/10.1145/1721654.1721672

[17] J. Rutkowska, "Subverting Vista kernel for fun and profit," in *Black Hat Briefings*, Tokyo, Japan, 2006. [Online]. Available: http://www.orkspace.net/secdocs/Conferences/BlackHat/Asia/2006/Subverting%20Vista%20Kernel%20For%20Fun%20And%20Profit.pdf

[18] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "Hey, you, get off of my cloud: Exploring information leakage in third-party compute clouds," in *Proc. of the 16th ACM conference on Computer and Communications Security (CCS '09)*. New York, NY, USA: ACM, 2009, pp. 199–212. [Online]. Available: http://doi.acm.org/10.1145/1653662.1653687

[19] "Amazon CloudWatch," 2014. [Online]. Available: http://aws.amazon.com/cloudwatch/

[20] "Monitoring and autoscaling features for Windows Azure with AzureWatch indepth," 2014. [Online]. Available: http://www.paraleap.com/azurewatch

[21] J. Chan, "Cloud security at Netflix," Apr. 2012. [Online]. Available: http://www.slideshare.net/jason_chan/cloud-security-at-netflix

[22] A. Cockcroft, "Netflix global cloud architecture," Oct. 2012. [Online]. Available: http://www.slideshare.net/adrianco/netflix-global-cloud