Insider Threats to Cloud Computing: Directions for New Research Challenges

William R Claycomb
CERT® Program
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA, USA
Email: claycomb@cert.org

Alex Nicoll
CERT® Program
Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA, USA
Email: anicoll@cert.org

***Abstract*—Cloud computing related insider threats are often listed as a serious concern by security researchers, but to date this threat has not been thoroughly explored. We believe the fundamental nature of current insider threats will remain relatively unchanged in a cloud environment, but the paradigm does reveal new exploit possibilities. The common notion of a cloud insider as a rogue administrator of a service provider is discussed, but we also present two additional cloud-related insider risks: the insider who exploits a cloud-related vulnerability to steal information from a cloud system, and the insider who uses cloud systems to carry out an attack on an employer's local resources. We also characterize a hierarchy of administrators within cloud service providers, give examples of attacks from real insider threat cases, and show how the nature of cloud systems architectures enables attacks to succeed. Finally, we discuss our position on future cloud research.**

***Keywords*-insider, cloud, security**

I. Introduction

Organizations continue to embrace the advantages of flexibility, scalability, and management provided by cloud computing platforms and services, and often consider security one of their top concerns in cloud environments. One of the most serious challenges, not only to cloud computing, but to data security in general, is the insider threat - a threat well known to security professionals. The CERT Insider Threat Center defines a malicious insider as a "current or former employee, contractor, or other business partner who has or had authorized access to an organizations network, system or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organizations information or information systems." [1] Since 2001, over 700 cases of actual insider crimes have been collected and analyzed by CERT researchers. The crimes collected range across multiple sectors, include small companies to multi-national corporations, and cover several hundred types of exploits used by malicious insiders to harm an organization. Some of these insiders relate to cloud computing, but the subject has not been thoroughly explored. In 2010, the Cloud Security Alliance (CSA) released *Top Threats to Cloud Computing,* describing seven threat areas considered most important to organizations using cloud services, including malicious insiders [2]. The CSA report describes the insider threat in cloud computing as a malicious employee of a cloud provider accessing sensitive customer data. Additional details from the report indicate "76% of respondents believe that the likelihood of Malicious Insiders in the cloud is possible, likely, or frequent." [3]

Yet despite these security concerns, cloud computing use continues to grow. One of many cloud service providers, Amazon.com has been offering commercial cloud computing services for over 5 years, and today, cloud computing is used by millions of people. It has been embraced by governments, academia, and the world's largest corporations. Given the wide-spread adoption and pervasive coverage from personal to business use, one might expect an abundance of cloud-related insider threat incidents. But despite the grim predictions and creative hypothetical attacks presented by researchers, we have little evidence of actual events involving the type of insider described in CSA's document. However, insiders do use the cloud to commit crimes, and the threat should not be dismissed. In this document, we will briefly discuss three types of insider threats related to cloud computing, and share tips for reducing the risk of these types of attacks. We present a hierarchy of service provider administrators, and show how the architecture of cloud computing enables certain types of attacks to succeed. Finally, we share our recommendations for future directions in insider threat research for cloud computing.

II. Three Types of Cloud-Related Insider Threats

We consider the cloud-related insider threat from three different perspectives: the rogue cloud provider administra-

tor, the employee in the victim organization that exploits cloud weaknesses for unauthorized access, and the insider who uses cloud resources to carry out attacks against the company's local IT infrastructure. Though we describe cloud-specific insiders, we believe the people behind these malicious insider attacks will continue to fit the profiles of other insider crimes identified by CERT: theft, sabotage, and fraud [1]. As a result, mitigation strategies may be extrapolated from prior insider threat models, and we will briefly discuss those options as well.

A. Rogue Administrator

Let us first consider the type of insider described by CSA [2] - the rogue administrator employed by a cloud provider. This cloud-related insider is the most commonly addressed by researchers. An attack often posited by this insider is theft of sensitive information, resulting in loss of data confidentiality and/or integrity. The insider described by this threat may be motivated financially, a common motivator for theft of intellectual property or fraud. But another attack possibility to consider is IT sabotage, where an employee seeks to harm an employer's IT infrastructure. Some may dismiss this type of crime in cloud environments, where administrators work for the provider, not the customer organizations. However, this should not be entirely discounted. Even if it is unlikely an insider has a grudge against the victim organization, an insider's a grudge against the cloud provider could result in harm to a victim organization with the intention of damaging the cloud provider's reputation.

The following excerpt is adapted from an actual case in CERT's database of insider threats, and describes an attacker motivated by financial gain:

The insider was employed as a system administrator at a data-mining firm contracted by a victim organization to process customer information. Though unnecessary for the job function, the insider had access to servers and data owned by the victim organization. An unprotected file containing encrypted password information was found on one of these servers. The insider brute-force attacked over 300 passwords, accessed data belonging to dozens of the victim organization's customers, and downloaded millions of personal records. Fortunately, the information was never sold or released by the insider prior to arrest.

This example highlights a case involving a trusted business partner, in a relationship very similar to cloud computing. But despite CERT's ongoing process for gathering current insider threat incidents, this is one of only a few cases we have describing a cloud-related scenario. However, we do have many cases involving contractors, temporary employees, and employees of trusted business partners exploiting authorized access to an organization's systems and information to steal data. That isn't quite cloud-related, but does pose some of the same challenges; specifically, who does an organization choose to trust with access to sensitive information, and how does the organization control the access it grants?

Another often-overlooked threat posed by the rogue administrator is impact to customer data *availability.* Another incident captured by CERT describes a system administrator at an organization very similar to the one described above, one that managed data and operations for other companies. However, rather than compromise customer information, this insider simply removed critical software from the hosting systems, preventing the provider from responding to customer requests. While customer data remained intact and confidential, customers were adversely affected by the lack of access to important information.

Different Types of Rogue Administrators: It is important to note that the threat of rogue administrators is layered differently for a cloud architecture than a standard enterprise environment. There are at least four levels of administrators to consider in the cloud:

- Hosting Company Administrators
- Virtual Image Administrators
- System Administrators
- Application Administrators

We propose this is a hierarchical relationship, that higher level administrators have the capabilities of lower level administrators. Figure 1 shows this relationship, and gives a summary of example attacks and/or vulnerabilities each level of administrator could exploit.

B. Exploit Weaknesses Introduced by Use of the Cloud

A second type of cloud-related insider threat, often overlooked by security researchers, is the insider within the organization who exploits vulnerabilities exposed by the use of cloud services to gain unauthorized access to organization systems and/or data. This may be malicious or accidental, and is sometimes enabled by differences in security policies or access control models between cloud-based and local systems. This threat may also be successful because direct administrative control of systems and data can difficult for an organization to effect quickly. This type of insider is most likely looking to gain access to sensitive information to sell (fraud) or use for future employment opportunities (theft of intellectual property), and the cloud may provide the easiest way to compromise security measures with the least chance of detection. But once again, sabotage attacks should not be discounted. It is unlikely a local insider would try to sabotage the cloud infrastructure itself, considering the resilience and stability of cloud based systems, in addition to the remote location of cloud systems. A local system may be a better target for sabotage, unless the insider seeks to harm the company by leaking sensitive or embarrassing company information.

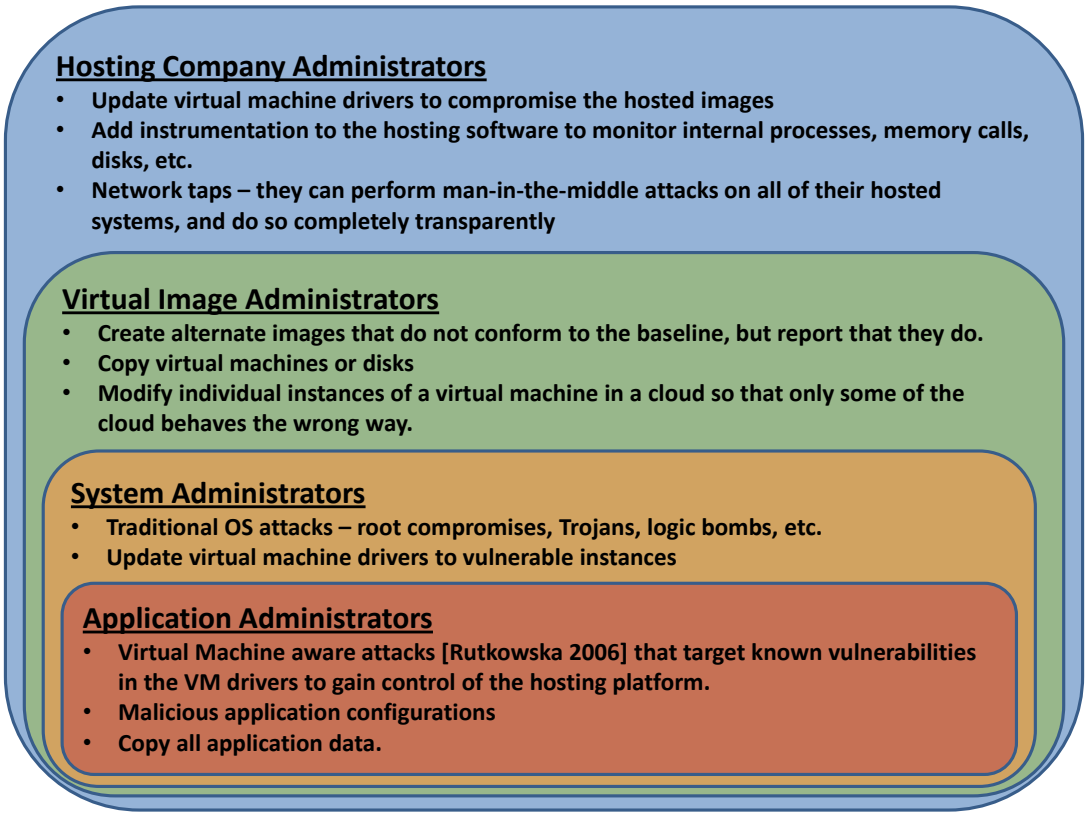

Figure 1. Hierarchy of Cloud-Based System Administrators and Potential Vulnerabilities and/or Exploits

This type of incident is described in the following excerpt from an actual case in the CERT insider threat database:

An employee in the victim organization was tricked by a malicious outsider into opening a document infected with malware. Using that exploit, the attacker was eventually able to gain access to the organization's email service, hosted by a cloud computing provider. Though aware of the attack in progress, the victim organization was unable to terminate email service quickly enough to prevent sensitive data loss. This delay was exacerbated by the inability of the organization to validate its identity with cloud provider support personnel.

This attack actually describes an unintentional insider attack, which means the employee did not intend to cause harm to the organization. However, the outside attacker was able to obtain credentials of an inside administrator, and use those credentials to attack as an inside administrator might. The exploited weakness that allowed this attack to succeed was the lack of direct control of email services for the victim organization.

1) Replication Lag Exploit: To further illustrate this type of insider, but from the perspective of fraud, consider an attack exploiting the increased latency, or *replication lag*, between servers in a cloud architecture. With constraints such as high server load, multiple network segments and layers between servers, and geographic separation, replication of changes from one server to others can take significantly longer for cloud systems than those hosted on-site, dedicated to the organization, and using the same network infrastructure. An insider who understands the hosted application environment can take advantage of that knowledge to devise an attack. First, he must be aware of an upcoming change (or be able to initiate one), which is introduced at the top of the hierarchy and replicated to nodes further down. He would then introduce a malicious change at a point further down the replication hierarchy, knowing his change will only exist for a very short period of time. The insider would take advantage of that short window to carry out the attack on the target node. This is very similar to the Byzantine Generals Problem [4], which deals with malicious nodes during message replication. However, the situation described here does not assume malicious nodes, it simply inserts what appears to be an authorized message immediately prior to replication and takes advantage of the temporary inconsistency caused.

As a specific example, consider a sample organization

with authoritative price server A, which replicates prices to servers B_1 and B_2, which have 1 and 2 seconds of latency, respectively. Server B_1 replicates prices to servers C_1 and C_2, which have 2 seconds of latency each. Server B_2 replicates prices to server C_3 with 4 seconds of latency. Assume an insider wants to buy a large number of a $20 item from his company, but he only wants to pay $10 each. If he knows about an upcoming price change for the item, say from $20 to $18, he could stage a false replication notice incorrectly listing the new sales price as $10, and send that notice to server C_3 so that it arrives four seconds after the initial price change is initiated. Then he carefully times his purchase, from C_3, before the correct replicated message is received two seconds later, overwriting the incorrect price and potentially removing evidence of the attack.

C. Using the Cloud to Conduct Nefarious Activity

A third type of cloud-related insider is one who uses cloud services to carry out an attack on his own employer. This is similar to the previous type of insider, who targets systems or data in the cloud. In contrast, the third type of insider uses the cloud as the tool to carry out the attack on systems or data targeted that are not necessarily associated with cloud-based systems. Though more uncommon than the previous two examples, this type of attack could present itself in the following scenarios:

- A financially troubled insider exploits the processing power of cloud services to crack password files, allowing unrestricted access to company bank accounts.
- A disgruntled insider uses several relatively cheap, easily configured cloud systems to launch a distributed denial of service attack on his organization, hindering incident investigation and limiting forensic analysis.
- A insider planning to leave the company leverages cloud storage to consolidate and exfiltrate sensitive information to take to a new job with a competitor.

There are very few empirical cases of the first two examples. However, CERT has cataloged many cases of the third - insiders using cloud-based services to steal information. These are usually instances of theft of intellectual property. Often the attacks use web-based email (i.e. Gmail, Hotmail, etc.) or file-sharing services (i.e. DropBox), which may circumvent controls in place to filter and/or monitor corporate email attachments. More information on this type of crime is presented by Moore et al. [5].

III. Securing Against Cloud-Related Insiders

Security of cloud computing is a popular research topic, and insider threats in the cloud is no exception. Unfortunately, as cloud computing is primarily a collection of previously existing technologies used in a new way, many solutions to cloud security concerns are merely repackaged solutions to other problems. Though responsibilities may differ, there are few fundamental differences between a rogue administrator at the cloud provider and a rogue administrator within the customer organization; both insiders have root access to systems and data, and both may employ similar types of attacks to steal information. However, architecture differences and trust issues between organizations and cloud providers does present the need for specialized approaches to insider security in the cloud.

A. Protecting Against Rogue Administrators

The remediation listed in CSA's document is quite applicable to the rogue administrator [2]:

- Enforce strict supply chain management and conduct a comprehensive supplier assessment.
- Specify human resource requirements as part of legal contracts.
- Require transparency into overall information security and management practices, as well as compliance reporting.
- Determine security breach notification processes.

Many of these items can be achieved through careful management and enforcement of service level agreements (SLAs) with cloud providers. Though enforcement of SLAs is difficult, due to transparency issues with cloud providers, Maurer et al. [6] and Emeakaroha et al. [7] present methods for SLA enactment and monitoring that organizations may find useful to consider.

Several researchers suggest encryption as a method of protecting data in the cloud. Two novel solutions are described by di Vimercati, et al. [8] and Itani et al. [9], but these are clearly not the only options proposed for secure data storage in the cloud. One vulnerability of data encryption with respect to rogue administrators is that encryption keys stored or used on cloud systems are subject to eavesdropping. A motivated attacker could potentially recover decryption keys using memory analysis on the host system. Using cloud services to simply store and/or transfer encrypted information, without introducing the associated keys to the cloud system, is a potential way to protect that data from a rogue cloud administrator. For some organizations, this could be a way to protect the data from rogue local administrators as well. That is, a cloud-based rogue administrator would have access to the encrypted data, but not the associated keys, and a local rogue administrator would have access to the locally-stored keys, but not the encrypted data. It is easy to point out a weaknesses with this suggestion - how does the organization prevent the local rogue administrator from stealing credentials to access the cloud-based data? Proper enforcement of local separation-of-duties policies [10] could be a viable approach to that problem.

Other issues that illustrate the risk posed rogue administrators are the lack of involvement an organization has in the hiring process, access control procedures, and monitoring of system administrators at the cloud provider. The concerns are generally as follows: How does an organization know the

cloud provider is enforcing strict controls for administrator access? How does the customer know the cloud provider enforces strict hiring guidelines? How can the organization be assured the cloud provider is adequately monitoring for insider attacks? The insinuation could be made that cloud providers hire any system administrator they can find, regardless of qualifications or security concerns. But current events do not seem to support the notion that cloud providers have no security vetting process and hire unknown and untrusted administrators. Otherwise, cases of nefarious insiders stealing sensitive information from within cloud providers would abound. Rather, it seems cloud providers have a vested interest in hiring carefully screened administrators that meet the security requirements of their customers. This would seem to be particularly true for very large and visible cloud providers seeking to attract business from multinational corporations, governments, etc., such as the U.S. Department of the Interior, which recently announced a 7-year, $35 million contract for cloud email and collaboration services [11]. Because data protection is critical to business success, cloud providers simply cannot afford a rogue administrator incident - and they have enormous resources available to ensure system administrators are carefully vetted prior to hiring, given very limited access to systems with customer data, and are carefully monitored for indications of malicious activity.

B. Protecting Against Cloud Exploits

Protecting against the insider who uses weaknesses exposed through use of cloud services is also challenging, but can be addressed via diligence and planning in implementing, transitioning to, and maintaining cloud services. Enforcing fundamental security controls such as separation of duties, least privilege, consistent auditing, data loss prevention, etc., on cloud-hosted systems is important. Organizations should not assume that because the system is hosted by a cloud provider that security is also handled externally. Current research on this topic includes solutions by Shin et al. showing methods for authorization [12] and access control [13].

Additionally, organizations should have agreements and policies in place with cloud providers to handle cloud-based security incidents. A plan for incident response, including offline credential verification, is essential for a timely and efficient reaction to an attack in progress. System administrators within the organization should be familiar with configuration tools for their cloud-based systems, including procedures for quickly changing access controls or even disabling cloud-based services if necessary.

C. Protecting Against Those Using the Cloud Against You

Detecting insiders who use cloud-based services to carry out attacks on local resources can be challenging, particularly if an organization permits internal access to these services, such as web-based email accounts. Data loss prevention tools and techniques can be effective in detecting sensitive data being sent via email or uploaded to cloud-based storage. Limiting employee access to external resources via network or host-based controls (i.e. firewalls, proxies, etc.) is another option for some organizations.

IV. Future Research

Cloud computing security is ripe with new opportunities for future research, including cloud-related insider threats. As mentioned previously, we do not believe the nature of the insider will change due to cloud computing's impact, but the opportunities for attacks will broaden. Researchers should take note of these new opportunities and respond accordingly to prevent, detect, and respond to new cloud-related insider attacks. Some important future research topics are:

- Socio-technical approach to insider threats
- Predictive models
- Identifying cloud-based indicators
- Virtualization and hypervisors
- Awareness and reporting
- Normal user behavior analysis
- Policy integration

A. Socio-Technical Approaches and Predictive Models

CERT has long advocated that insider threat prevention requires a combination of non-technical ("socio-") and technical input, as shown in Figure 2. Examples of non-technical input include information on workplace behavior (tardiness, conflicts with others, etc.), personal behavior (drug or alcohol abuse, overwhelming debt, etc.), or human-resources data (demographics, salary, position). Technical input is a more common data source used by researchers, and includes system logs, intrusion detection or data loss prevention systems, etc. One of the first authors to identify the need to combine these sources need is Schultz, who notes, "many different potential indicators of internal attacks exist," and suggests several indicators including technical and non-technical components. Others have championed the same cause [14]–[16], but few have implemented successful real-world socio-technical monitoring systems, and this remains an open research challenge.

Corresponding to that challenge is the ability to rank or weight the importance of specific indicators. The UK Centre for the Protection of National Infrastructure advocates an ongoing insider threat risk management program that includes screening, access controls, promoting effective security culture, social engineering, protective monitoring and intrusion detection, and investigations. [17]. Greitzer et al., advocate a predictive model that identifies several weighted indicators of insider risk. Furthermore, Greitzer et al, discuss development of a reasoning system to integrate multiple data sources and help analysts identify high-risk

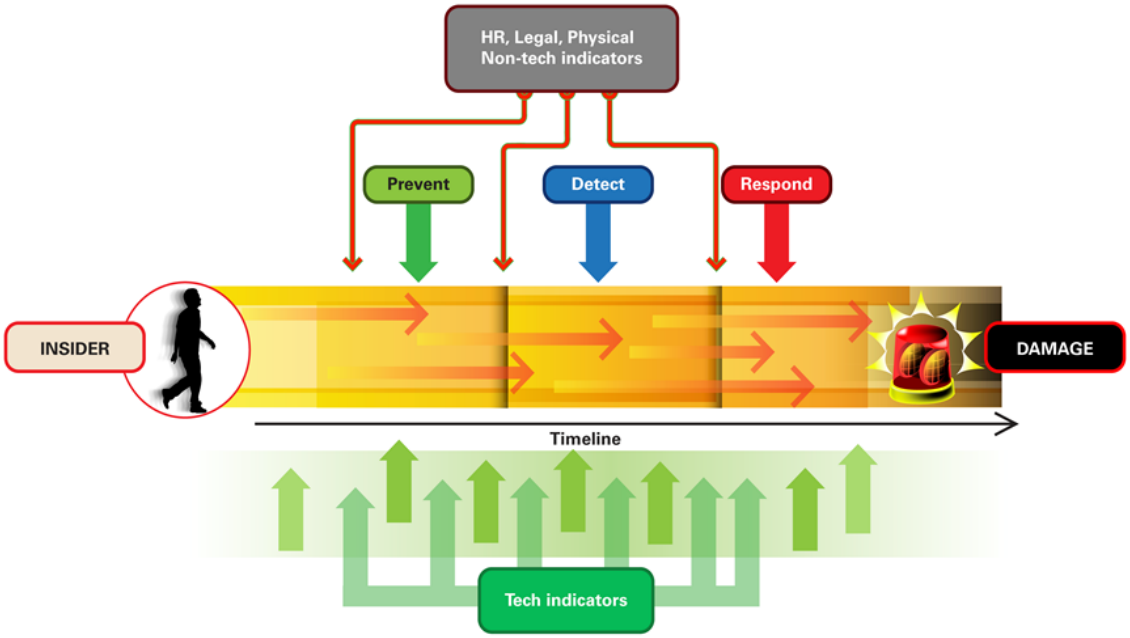

Figure 2. Opportunities for prevention, detection, and response for an insider attack

events [18]. Successfully ranking a combination of technical and non-technical indicators is a very challenging topic for future research.

A socio-technical approach to insider threats related to cloud computing is not directly applicable from the perspective of an organization concerned with the rogue administrator at the cloud provider, but it is useful when looking for employees who exploit cloud weaknesses or use the cloud against the employer. On the other hand, organizations may have some insight into certain important non-technical aspects of the cloud provider, such as hiring processes including pre-employment screening. Understanding how pre-employment screening [19] can identify potential threats is critical to reducing overall insider threats, both for cloud service consumers and providers. New directions in this area could include careful analysis of which pre-employment screening practices are most effective at identifying potential insider threat issues.

B. Identifying Cloud-Based Indicators of Insider Threats

Identifying *indicators* of insider threats is another subject of ongoing work [20]. However, many indicators suggested for cloud-based insider threats are simply reworded versions of malicious behavior indicators for non-cloud systems (i.e. access outside normal work hours, abnormal search patterns, obtaining back-door access to company data.) While these should not be discounted, identifying indicators unique to cloud environments could significantly improve the likelihood of detecting cloud-based insider attacks. For instance, some technical indicators of rogue administrators at the cloud provider could be the following: violation of SLAs, improper virtual machine management, using suspicious software, or performing similar activities across different platforms and customer systems. Non-technical indicators might include those that indicate a lack of concern for company policy or the protection of others' data (i.e. carelessness, indifference towards customer concerns, etc.)

Researchers should be careful to identify a wide-range of potential indicators. Ilgun et al. note four types of intrusion detection methods, which also apply to insider threat detection: threshold, anomaly, rule-based, and model-based [21]. Each method has advantages and limitations, as noted by Greitzer and Hohimer [14]. For instance, threshold based methods can be foiled by remaining within set limits; anomaly-based methods can be manipulated by clever insiders; rule-based methods are limited to a strictly defined set of criteria and eliminate the detection of novel attacks; model-based methods are expressive enough to encompass different behaviors, but often focus on audit records alone. Additionally, we find model-based methods difficult to implement as specific detection methods without becoming too rule-based. The combination of different methods would form a more holistic picture of insider behavior that could reduce false-positives and increase the chances of finding clever and/or novel insider threat attacks in the cloud.

C. Virtualization and Hypervisors

Examples of virtualization and hypervisor exploits [22], [23] highlight the need for work on enforcing virtual machine isolation. These attacks are technically sophisticated, and practically necessitate some level of insider access to the system being attacked. Indeed, it is difficult to imagine

accidental data loss due to hypervisor vulnerabilities. Potential new research could include new technologies that could more completely implement virtual machine segmentation, perhaps using hardware enforced mandatory access controls and process separation.

D. Awareness and Reporting

In May 2012, the FBI released a news story titled *Economic Espionage: How to Spot a Possible Insider Threat,* including a list of insider threat warning signs and potential contact information [24]. In fact, many of the insiders described in CERT's database were detected via reporting by others (co-workers, customers, management, etc.) Improving insider threat awareness and reporting programs is critical to improving the ability of others to identify signs of potential insider activity and increasing employee confidence in raising concerns to the appropriate authorities. Exploring which types of awareness campaigns are most effective for specific audiences, as well as developing measurably-improved reporting mechanisms will give organizations a better chance at detecting attacks as soon as possible.

E. Normal User Behavior Analysis

Some observable insider activities are clearly harmful to the organization; for instance, an insider deleting critical applications from the organization's servers. However, not all insider activity is so blatantly malicious. A clever insider seeking to avoid detection will attempt to use authorized access to the target information/systems, and do so in a manner unlikely to raise suspicion. In reviewing the literature, we find many novel proposals for detection of specific insider-related activity, but few that compare the proposed insider behavior to similar non-malicious behaviors, or even acknowledge the necessity of doing so. One counter-example of this trend is Greitzer and Hohimer, who note, "There are several reasons why development and deployment of approaches to addressing insider threat, particularly proactive approaches, are so challenging: (a) the lack of sufficient real-world data that has 'ground truth' enabling adequate scientific verification and validation of proposed solutions; (b) the difficulty in distinguishing between malicious insider behavior and what can be described as normal or legitimate behavior." [14]

Few publicly-available data sets exist that characterize normal user behavior in relation to indicators of insider threats, much less indicators related to cloud-based insiders. Researchers addressing the challenge of collecting and analyzing normal user behavior should be careful to include attributes useful for cloud-based research as well. For instance, correlating access requests across multiple disparate systems, exploring how often and how much data users transfer from the organization to cloud-based systems (web-based mail, etc.), or how often cloud-based administrative tools are used. Collecting and sharing such information will greatly enhance the ability of other researchers to propose and validate indicators of malicious cloud-related insider behavior.

F. Policy Integration

A final suggestion for future research topics is exploring how organizations can better manage discrepancies among cloud-based security policies. These may arise due to conflicts between local and cloud-based policies, different policies for each service consumed, or the use of multiple cloud service providers, each with different security policies. Other barriers further exacerbate seamless policy integration, such as differences in operating systems and less control of auditing capabilities in the cloud (i.e. physical). Takabi et al. propose developing a trust management framework for policy integration and an ontology to address semantic heterogeneity among policies [25]. Researchers should carefully consider the danger of combining inadequate cloud policy management with the limited resources many organizations have to implement costly or complicated policy management systems. One solution would be to propose automated, easy to understand, and easily verifiable policy management techniques for cloud-based systems.

V. Conclusion

Insider threats are a persistent and increasing problem. Cloud computing services provide a resource for organizations to improve business efficiency, but also expose new possibilities for insider attacks. Fortunately, it appears that few, if any, rogue administrator attacks have been successful within cloud service providers, but insiders continue to abuse organizational trust in other ways, such as using cloud services to carry out attacks. Organizations should be aware of vulnerabilities exposed by the use of cloud services and mindful of the availability of cloud services to employees within the organization. The good news is that existing data protection techniques can be effective, if diligently and carefully applied.

Future research on cloud-related insider threats should focus on identifying and addressing unique vulnerabilities posed by the use of cloud computing services. We caution against simply casting previous solutions to other problems in the light of a cloud environment; this has little benefit to the community and should be avoided unless a distinct advantage can be obtained and measured. Rather, we suggest an approach grounded on solid information assurance principles and on focused on finding new solutions that address real threats to cloud computing.

Acknowledgment

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

REFERENCES

[1] D. Cappelli, A. Moore, and R. Trzeciak, *The CERT Guide to Insider Threats: How to Prevent, Detect, and Respond to Information Technology Crimes (Theft, Sabotage, Fraud)*, ser. SEI Series in Software Engineering. Addison-Wesley Professional, 2012.

[2] C. S. Alliance, "Top threats to cloud computing, version 1.0," Cloud Security Alliance, Tech. Rep., March 2010. [Online]. Available: http://www.cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf

[3] ——, "Star-205 cloud security alliance - top threats to cloud computing v2.0.pdf," in *RSA Conference Europe*, 2010. [Online]. Available: http://365.rsaconference.com/docs/DOC-2819

[4] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Trans. Program. Lang. Syst.*, vol. 4, no. 3, pp. 382–401, Jul. 1982. [Online]. Available: http://doi.acm.org/10.1145/357172.357176

[5] A. P. Moore, D. M. Capelli, T. C. Caron, E. Shaw, D. Spooner, and R. F. Trzeciak, "A preliminary model of insider theft of intellectual property," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, vol. 2, no. 1, 2011.

[6] M. Maurer, I. Brandic, and R. Sakellariou, "Self-adaptive and resource-efficient SLA enactment for cloud computing infrastructures," in *5th International Conference on Cloud Computing (IEEE Cloud)*, 2012.

[7] V. C. Emeakaroha, T. C. Ferreto, M. A. S. Netto, I. Brandic, and C. A. F. De Rose, "Casvid: Application level monitoring for SLA violation detection in clouds." in *IEEE Computer Software and Applications Conference (COMPSAC 2012)*, 2012.

[8] S. D. C. di Vimercati, S. Foresti, S. Jajodia, S. Paraboschi, and P. Samarati, "Over-encryption: management of access control evolution on outsourced data," in *Proceedings of the 33rd international conference on Very large data bases*, 2007.

[9] W. Itani, A. Kayssi, and A. Chehab, "Privacy as a service: Privacy-aware data storage and processing in cloud computing architectures," in *Dependable, Autonomic and Secure Computing, 2009. DASC '09. Eighth IEEE International Conference on*, Dec. 2009.

[10] R. Sandhu, "Separation of duties in computerized information systems," in *In Database Security IV: Status and Prospects*. North-Holland, 1990, pp. 179–189.

[11] U.S. Department of the Interior, "Press Release: Interior Selects Google Apps for Government for Cloud Email and Collaboration Services," http://www.doi.gov/news/pressreleases/Interior-Selects-Google-Apps-for-Government-for-Cloud-Email-and-Collaboration-Services.cfm, May 2012.

[12] D. Shin, H. Akkan, W. Claycomb, and K. Kim, "Toward role-based provisioning and access control for infrastructure as a service (IaaS)," *J. Internet Services and Applications*, vol. 2, no. 3, pp. 243–255, 2011.

[13] D. Shin, Y. Wang, and W. Claycomb, "A policy-based decentralized authorization management framework for cloud computing," in *ACM Symposium on Applied Computing (ACM SAC)*, 2012.

[14] F. Greitzer and R. Hohimer, "Modeling human behavior to anticipate insider attacks," *Journal of Strategic Security*, vol. 4, no. 2, 2011.

[15] D. Cappelli, A. Moore, R. Trzeciak, and T. J. Shimeall, "Common sense guide to prevention and detection of insider threats 3rd edition version 3.1," Carnegie Mellon University / SEI, Tech. Rep., January 2009. [Online]. Available: http://www.cert.org/archive/pdf/CSG-V3.pdf

[16] J. Montelibano and A. Moore, "Insider threat security reference architecture," *Hawaii International Conference on System Sciences*, vol. 0, pp. 2412–2421, 2012.

[17] C. for the Protection of National Infrastructure (CPNI), "Risk assessment for personnel security: A guide," Centre for the Protection of National Infrastructure (CPNI), Tech. Rep., 2010.

[18] F. Greitzer, L. Kangas, C. Noonan, A. Dalton, and R. Hohimer, "Identifying at-risk employees: Modeling psychosocial precursors of potential insider threats," in *45th Hawaii International Conference on System Science (HICSS)*, January 2012.

[19] C. for the Protection of National Infrastructure (CPNI), "Pre-employment screening: A good practice guide," Centre for the Protection of National Infrastructure (CPNI), Tech. Rep., September 2011.

[20] M. Hanley, "Deriving candidate technical controls and indicators of insider attack from socio-technical models and data," Carnegie Mellon University / SEI, Tech. Rep., January 2011. [Online]. Available: http://www.cert.org/archive/pdf/11tn003.pdf

[21] K. Ilgun, R. Kemmerer, and P. Porras, "State transition analysis: a rule-based intrusion detection approach," *Software Engineering, IEEE Transactions on*, vol. 21, no. 3, pp. 181 –199, mar 1995.

[22] J. Rotkowska, "Subverting vista kernel for fun and profit," Black Hat USA, Tech. Rep., 2006. [Online]. Available: http://blackhat.com/presentations/bh-usa-06/BH-US-06-Rutkowska.pdf

[23] K. Kortchinsky, "Cloudburst: A vmware guest to host escape story," Black Hat USA, Tech. Rep., 2009. [Online]. Available: http://www.blackhat.com/presentations/bh-usa-09/KORTCHINSKY/BHUSA09-Kortchinsky-Cloudburst-PAPER.pdf

[24] Federal Bureau of Investigations, "Economic espionage: How to spot a possible insider threat," May 2012. [Online]. Available: http://www.fbi.gov/news/stories/2012/may/insider_051112/insider_051112

[25] H. Takabi, J. Joshi, and G. Ahn, "Security and privacy challenges in cloud computing environments," *Security Privacy, IEEE*, vol. 8, no. 6, pp. 24 –31, nov.-dec. 2010.