



Interested in learning more about security?

SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Following Incidents into the Cloud

The increased level of complexity which cloud computing has introduced to incident handling is not well understood nor fully engaged. Further complicating this are economic pressures that move companies to commit themselves on a course of cloud adoption without taking into account the inherent risks. The complete array of risks and effects of cloud integration on incident handling are often not fully assessed and mitigated prior to contract signing and implementation. These decisions can have serious consequences de...

Copyright SANS Institute
Author Retains Full Rights

AD

Build your business'
breach action plan.

START NOW

 **LifeLock**
BUSINESS SOLUTIONS

No one can prevent all identity theft. © 2016 LifeLock, Inc. All rights reserved. LifeLock and the LockMan logo are registered trademarks of LifeLock, Inc.

Following Incidents into the Cloud

GIAC (GCIH) Gold Certification

Author: Jeff Reed, jeff360@gmail.com

Advisor: Adam Kliarsky

Accepted: September 20, 2010

Abstract

The increased level of complexity which cloud computing has introduced to incident handling is not well understood nor fully engaged. Further complicating this are economic pressures that move companies to commit themselves on a course of cloud adoption without taking into account the inherent risks. The complete array of risks and effects of cloud integration on incident handling are often not fully assessed and mitigated prior to contract signing and implementation. These decisions can have serious consequences depending on the regulatory requirements of a company or organization and its ability to absorb realized risk. This paper attempts to identify problem areas for incident management in an enterprise considering cloud integration, and provide clarifying questions that should be answered satisfactorily before committing to cloud integration and signing contracts with cloud providers.

1. Introduction

The availability and use of cloud computing continues to grow. Discussions of and references to its benefits and issues grow at a similar pace. As it continues to move from a sort of ‘SOA of the Wild West’ into the mainstream, more companies will face the myriad questions arising from its use. When, why, where and how should integration with the cloud occur? How can one be certain that a cloud provider will survive through an organization’s technology integration lifecycle?

In March of 2010, the UK’s Centre for the Protection of National Infrastructure in their *Information Security Briefing 01/2010 on Cloud Computing* said, “There is, to date, no universally agreed industry definition of cloud computing and it is usual to find conflicting descriptions in any nascent industry” (2010). Nevertheless, sufficient agreement is found upon which this discussion may be based (refer to Mell, Grance, 2009, Catteddu, Hogben, 2009, and Youseff, Butrico, Da Silva, 2008). This paper uses the NIST definition of cloud computing (included in appendix Section 7.1) as its working definition. Part of the NIST definition states it this way:

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models. (Mell, Grance, 2009)

1.1. Cloud Security Climate

Is cloud computing secure and what does that mean? The concerns regarding security are neither insignificant nor are they coming from a “fringe element” of industry professionals. Cisco CEO John Chambers, in a keynote speech to the 2009 RSA conference stated that “[cloud computing] is a security nightmare, and it can’t be handled in traditional ways” (Chambers, 2009). Art Coviello, President of RSA, said that “...you won’t want any part of [cloud adoption] unless service providers can demonstrate their

Jeff Reed, jeff360@gmail.com

ability to effectively enforce policy, prove compliance and manage multi-tenancy” (quoted by Greene, 2010).

Nevertheless, in some cases the cloud will offer a better security posture than an organization could otherwise provide. The 2009 ENISA (European Network and Information Security Agency) report, *Cloud Computing – Benefits, risks and recommendations for information security* (Catteddu, Hogben, 2009), highlights the benefits that some small and medium size companies can realize with cloud computing. A smaller, cost-constrained organization may find that a cloud deployment allows them to take advantage of large-scale infrastructure security measures that they could not otherwise afford. Some of the possible advantages include DDOS (distributed denial of service) protection, forensic image support, logging infrastructure, timely patch and update support, scaling resilience, and perimeter protection (firewalls, IDPS (intrusion detection and prevention services)) (Catteddu, Hogben, 2009).

One type of cloud integration cannot realize all of the benefits offered by nor all of the risks represented in cloud computing. Numerous cloud integration variations are possible, especially when taken in context with an organization’s internal architecture. Each business will need to perform its own analysis to ensure that the organization’s requirements are met for any given cloud integration.

1.2. Background

Cloud computing has built on industry developments dating from the 1980s by leveraging outsourced infrastructure services, hosted applications and software as a service (Owens, 2010). For the most part, the technology used is nothing new. Yet, in aggregate, it is something very different. The differences provide both benefits and problems for the organization integrating with the cloud. The addition of elasticity and pay-as-you-go to this collection of technologies makes cloud computing compelling to CIOs in companies of all sizes.

Cloud integration presents unique challenges to incident handlers as well as to those responsible for preparing and negotiating the contract for cloud services. The challenges are further complicated when there is a prevailing perception that the cloud integration is “inside the security perimeter,” or, “The company has written a contract

Jeff Reed, jeff360@gmail.com

requiring the vendor to be secure, that should be enough.” This sort of thinking may be naïve but, unfortunately, it is not rare. The cloud provider may have a great deal of built in security or they may not. Whether they do or not, incident handling (IH) teams will eventually face incidents related to the integration, necessitating planning for handling incidents in this new environment.

The impacts of cloud integration warrant a careful analysis by an organization before implementation. An introduction of a disruptive technology such as cloud computing can make both definition and documentation of services, policies, and procedures unclear in a given environment. The IH team may find that it is helpful to go through the same process that the team initially followed when establishing their IH capability.

There has been much work done already on the topics of implementing and operating IH capabilities. Two well-known and readily available examples of these are *The Handbook for Computer Security Incident Response Teams, Second Edition*, (West-Brown, *et al*, 2004) and the *NIST Computer Security Incident Handling Guide* (Scarfone, Grance, Masone, 2008). The goal of this paper is to leverage these well-known volumes as representing a base of accepted best practice for incident handling rather than lay, again, the foundation provided by these works. To accomplish this goal an organization must start by analyzing these earlier frameworks against cloud computing implications to give context to the questions arising from the analysis. Section 2 (Perspective: Organizing the IH Response Capability) and Section 3 (Perspective: Incident Handling Lifecycle) of this paper closely parallel much of the outline and topics of the NIST paper (Scarfone, *et al*, 2008). It may prove helpful to the IH team, though it is not required, to have a copy of that paper for additional context while reviewing this one.

The cycle of reviewing the organization’s IH processes against the NIST (or other) framework can bring to light structure, process, procedure, and other unforeseen but required changes. In turn, these will raise questions that will need to be resolved prior to or through the contract process. In some cases, steps may be required to mitigate issues discovered by the review.

Jeff Reed, jeff360@gmail.com

1.3. Questioning: The Approach

This paper asks many questions with the goal of helping an organization to arrive at appropriate answers. With each new cloud integration, the IH team must address many variables unique to an organization's application or environment. Examples of these variables include the cloud service, the cloud vendor providing the service, the legal requirements implied by the physical locations affected, and different scaling approaches. Providing specific answers to the myriad issues arising from cloud integration would require many assumptions regarding these variables. This makes attempting to provide an answer regarding complex integration less useful than having provocative questions. Sometimes the best practical answer surfaces when asking a question that prompts the reader to undertake the analysis necessary to make the best integration decisions for their enterprise at that time.

The issues that incident handlers must resolve when integrating cloud components into their enterprise architecture are too numerous to address comprehensively in this paper (Owens, 2010). Necessarily, there will be significant variation from one organization to another. Each integration will have its own concerns in the full context in which it is made. The IH team, in cooperation with its various departments and partner organizations, will need to collaborate on each cloud integration scenario. It is the goal of this paper to prime and stretch a team's thinking as it works through these integrations and prepares to provide input into the cloud vendor contracting process.

1.4. Regarding Private Clouds

It is important, in a paper about incident handling in the cloud, to explain the lack of attention to one of the four cloud deployment types (for cloud deployment types see Appendix 7.1). Though the concepts and questions throughout the paper will aid the IH team in raising and addressing private cloud integrations, the focus remains on cloud integrations external to the company. An organization often engages their IH team late in the acquisition process, if at all. In the public, community, or hybrid models of cloud deployment, the IH team must rapidly perform complex analysis and research in order to meet deployment deadlines and ensure corporate security. Developing a private cloud capability in a business typically takes considerably longer. Additionally, the technical details of the architecture are open to the internal IH team, as is access to the cloud

Jeff Reed, jeff360@gmail.com

systems. Integrations with an internal cloud do not face many of the challenges presented by the other cloud deployment models. Motivations, priorities, jurisdictional concerns and other factors related to integration with another corporate infrastructure do not affect the enterprise with a private cloud.

However, when developing a private cloud, an enterprise will benefit from the review recommended by this paper. For private cloud, as in the other cloud models, the IH team must address new technologies and concepts such as elasticity, resource isolation, and fine-grained access control (see Section 4 | Additional Considerations). Many of the other challenges discussed here, even in the private cloud case, require changes to IH processes and skill sets. Therefore, although not directly addressing private cloud, the paper will still provide significant benefit to the IH team.

2. Perspective: Organizing the IH Response Capability

The review process should begin by reviewing and analyzing the factors that went into the initial establishment of an organization's incident response capability. This section poses cloud integration considerations and questions arising from a review of these factors.

2.1. Events, Incidents, & Response

2.1.1. Terminology & Semantics

A very early step in the process of establishing an IH capability is defining the terminology and reaching agreement on semantics. All parties involved in the IH process must agree on the definitions of terms such as "incident," "breach," "event," and "threat" and adequately document them. Make sure that you include those who will be receiving reports, notifications, and other communications from the IH process.

Bringing all parties into agreement on these terms and collaborating to formalize them in policies and practices can be a difficult process within a single organization or company. When integrating with a cloud provider, the difficulties increase. For one organization, the terms 'incident' and 'breach' may be differentiated by whether legal notification is required as a result (Blanton, Schiller, 2010). Their cloud provider may have arrived at different definitions. The results of negotiations may be to 'agree to disagree' in writing. However, the process the IH team follows is typically tied to these terms creating the very real possibility of misdirection or delay of an incident investigation at a critical moment. The problem can be ameliorated, to some extent, by clearly documenting the semantic differences. However, the more cloud providers with differing definitions are in the mix the more the interpretation by those with vested interests can create conflicts. If there is an IH team member rotation, how will all members keep the various semantic differences they will encounter straight when operating with multiple cloud components?

2.2. Policies, Plan, & Procedures

"The plan, policies, and procedures should reflect the team's interactions with other teams within the organization as well as with outside parties, such as law

Jeff Reed, jeff360@gmail.com

enforcement, the media, and other incident response organizations” (Scarfone, *et al*, 2008). The very nature of cloud computing creates ambiguity with regard to a team’s relationship to the technology and to external teams with which the organization has to collaborate. To make matters worse, as an organization expands into more types and models of cloud integration, it must address these ambiguities and the differences between the providers and their offerings.

NIST further states, “A *computer security incident* is a violation or imminent threat of violation of computer security policies, acceptable use policies, or standard security practices” (Scarfone, *et al*, 2008). Building on these policies and practices for cloud computing implies the need for collaboration between two organizations which most likely have differing security postures. It also suggests the need for revisiting policies and practices to ensure that they sufficiently address cloud integration.

It is important to develop cloud IH use cases to help the IH team think through the possible scenarios that may arise. Clearly defined use cases will also serve as guides for working through each aspect in an organization’s planning.

2.2.1. Policy Elements

Policy governing incident response is framed for each organization in unique ways (Scarfone, *et al*, 2008). Whatever commonalities there may be between the ways that organizations frame their incident response policies, it is important for the IH team to review the differences between theirs and the cloud provider’s. Consider the areas of the policy affected by the proposed cloud integration under review. It is also helpful for the policy to encompass more than this one integration if possible. In reviewing the organization’s policies, one should consider what cloud architecture and cloud relationship scenarios are possible given current business and IT architecture directions. What modifications to the policies are necessary to ensure that other cloud possibilities are appropriately included without unduly constraining the current integration or the organization?

When looking at the cloud provider’s policies, the IH team needs to ensure that they are consistent with their own. Even if the cloud vendor uses the same terms as the customer does, the business needs to be certain these terms are used the same way and in

Jeff Reed, jeff360@gmail.com

the same contexts. If not, resolve these differences and, if necessary, clarify them with the vendor in writing. It is best to give the whole area of cloud provider policy review the broadest possible thinking, as some questions may not be obvious. Ensure that the cloud provider's policies articulate prioritization that is consistent with the security mission and posture of the organization. Though it may not be immediately obvious, the IH team should be able to derive motivations, position and intent by carefully considering statements of management commitment, purpose, objectives, levels of team authority, incident prioritization, and performance measures. If there is an issue with the cloud provider's prioritization, is there a way to reconcile this to the customer's satisfaction? If not, what additional steps should the customer take?

Regarding policy elements, NIST says that they "should include the authority of the incident response team to confiscate or disconnect equipment and to monitor suspicious activity, and the requirements for reporting certain types of incidents" (Scarfone, *et al*, 2008). Such actions may be impossible depending on the cloud scenario. What are the provider's policies and procedures for handling these scenarios? The cloud vendor may have the dilemma of confiscating or disconnecting the organization's systems, or at least the virtual systems and data in their cloud. The cloud integration will be in a shared (virtualized) environment. Typically, taking physical measures in the cloud will affect more than one organization or company. An IH team should ensure that the cloud vendor's processes include the customer in decision making related to taking physical measures. The vendor's processes should at least provide the customer with sufficient notification for operations teams to affect an orderly shutdown and inform users of the outage. Depending on criticality, the business may need to provide for the execution of business continuity processes. Does the contract cover the matter of taking physical measures to the legal department's satisfaction?

There are additional questions a business will need to ask regarding what incidents the cloud provider is reporting, to what level of detail, when, and to whom. Confirm whether legal and executive management are in agreement with the answers to these and, if not, if there is a way to mitigate this through contract language that the provider will accept. If there is no way to mitigate the gaps, what steps can the customer take?

Jeff Reed, jeff360@gmail.com

An organization may have performance measures for IH capabilities and processes to consider. These performance metrics may require data from the provider. Furthermore, there may be process components to enable the creation and proper interpretation of these performance measures. Are there other reporting or analytic data needs such as SEIM (Security Event and Incident Management) or other correlating event data that your IH team requires? Work with the purchasing team in negotiating these processes and data with the vendor.

2.2.2. Plan Elements

The mission statement of an organization or cloud provider can also impact incident handling in the cloud. A mission statement is broad, attempting to encompass the entirety of the IH field of responsibility in a few words. Because of this, it is unlikely that there will be substantial differences between the customer's mission statement and that of the cloud provider. Nevertheless, consider it carefully to identify the possibility of conflicts. Does the vendor's mission statement indicate, as was mentioned earlier, what their priorities or motivations will be during an incident engagement? Their motivations drive strategy or tactical measures they employ in a crisis. The vendor's priorities influence their reporting, communications and notifications, which, in turn, affect the customer. The mission statement may be too broad to determine these specifics. However, as stated in the previous section, this will be one of the indicators of a divergent set of priorities.

The statements of goals and strategies may be another matter. It is typically here that any organization's priorities and intents begin to emerge. They will target specific protections desired to meet their corporate objectives. These will be very legitimate areas of corporate concern for the cloud vendor. However, the cloud provider's priorities do not always align with the customer's objectives. The differences in priorities often surface in decision-making during an incident. In some types of incidents, they might have a tendency to err on the side of self-preservation. One area of concern is that of reporting. For instance, if the cloud provider sees that an attacker or malware have made it to one of the customer's platforms, but do not find evidence of stolen or modified data, will they decide against notifying the customer? Another area of concern is incident

Jeff Reed, jeff360@gmail.com

impact isolation. What happens if the cloud provider sees an incident originating from the customer's integration and decides to take the customer off-line to preserve their reputation or reduce production impacts? This is a very legitimate concern on their part, but one that can have significant impact on the customer, especially if the customer has no failover capability for this aspect of their production processing. If the organization sees differences in the goals and strategies, it will be important for the team to clarify them with the provider before proceeding with the integration.

In addition, while reviewing the above, allow for necessary plan accommodations or restructuring to address cloud integration. It is best to think this through carefully. It is particularly important to update the plan if there are not already provisions for hosted third party scenarios. Furthermore, if the organization has not taken into consideration elasticity, multi-location, multi-tenancy, virtual machine (VM) migration and other cloud characteristics, now is the time. The IH plan will look very different depending on the cloud model with which the organization is integrating (refer to Section 4.1 Structural Perspectives). How far into the provider's infrastructure will the IH team need to go while investigating incidents? An organization's planning and analysis may call for investigating further into the cloud infrastructure than the cloud provider will allow. A customer will need to negotiate where the lines are. The cloud vendor may provide a mature incident response capability as a service to their customers. Nevertheless, they will necessarily have limits as to what and how they work with the customer in IH. The organization will need to know if and where this is an issue and modify the plan accordingly.

When planning how the incident response team will communicate with the rest of the business, an organization will need to consider the composition of both the local IH team and the provider's IH team. What the cloud vendor is doing for IH affects their customer and vice versa in very intimate ways. Who will communicate what with whom, when, and to what extent? Check that the cloud provider's approach to communication is at least acceptably consistent with the organization's mandates. If it is not reasonably consistent, determine what mitigating steps the team requires to make the approach acceptable. Again, carefully consider the vendor's motivations. If the vendor already has

Jeff Reed, jeff360@gmail.com

language in contracts, SLAs or other statements of service, ensure that the language is sufficiently clear and definitive. If the language appears to meet the organization's requirements, have the semantics of the language been defined clearly enough that the organization can be certain that it is sufficient?

If the organization's mission requirements are such that there is a need to acquire incident metrics from the cloud provider, there are other questions to ask. Can the team accurately interpret the IH metrics from the vendor given contextual differences between the cloud environment and organization's environment? Clear definitions and an unambiguous process context will help in determining if the number of incidents handled is a meaningful statistic. However, other contextual characteristics will be necessary for a full understanding of the vendor's IH metrics. If the provider sends the organization numbers regarding incidents in the environment, given elasticity characteristics, will it be possible to tell what these numbers means statistically? For instance, interpretation of the numbers that the vendor provides possibly depends on the time or time range represented, number of hosts, instances, processes or threads running during those times, and more. Is it possible to get the level of specificity the business requires to understand the vendor's IH metrics? The organization must take steps to gain an appropriate understanding of the metrics they are requesting. All of these factors will be important in making fully informed and meaningful statistical analyses and evaluations.

The issues discussed in this and the previous section affect the team structure discussed in Section [2.4, Team Models, Structure, Staff, & Dependencies](#). A business will need to train its IH team members in the cloud provider's processes and capabilities. The customer will also need to adjust their processes to integrate the required cloud interactions. IH team members may need experience in virtual machines, virtual network devices, Software as a Service (SaaS) administration and auditing or numerous other skills that may be unique to this integration. If the vendor provides a self-service portal, the IH team will need to know the level of role definition granularity provided and how to administrate those roles (see Section [4.1.1, Elasticity Concerns](#)). The IH team will also need to understand the audit logging and monitoring controls provided by the cloud tools. What does audit logging look like for the elasticity mechanisms provided by the cloud

Jeff Reed, jeff360@gmail.com

vendor? Does it provide sufficient information to understand the scaling events or history it is representing? This is just a sample of the types of questions the organization will need to ask regarding team preparation and training. Every instance of cloud integration will be unique. The organization will need to understand the administrative components well enough to structure processes and training for them.

2.2.3. Procedure Elements

As the team is reviewing or developing procedures to handle the cloud integration scenarios, they will need to keep asking the question: Are there cloud contract and process implications here? The organization may need tools to be effective in the given deployment such as reporting, monitoring, tracking, auditing, forensic, and real-time analysis. Information regarding whether the cloud provider allows or supports these tool requirements and the impacts to the organization will aid customers in implementing effective IH procedures for the cloud integration.

2.3. External Interactions

How an IH team approaches external interactions is a highly sensitive area of the IH process. When, how and to whom the IH team communicates should be carefully coordinated with the organization's legal and public relations teams (Scarfone, et al, 2008). This is an area where governance issues arise due to cloud integration. Instead of an internal team with only internal policies governing their responses to an incident, a company now has an external team with a different set of policies and leadership directives governing their responses. Is the onus on the customer to think through every possible circumstance and get agreement from and the cloud vendor for each one? In the case of regulations driving organization policy, the answer to this is most likely, yes. In the cases where the organization's reputation is at stake, the answer is also most likely yes. Will the cloud vendor work with the customer to settle on agreeable answers to all of these questions? In general, the vendor will probably help the organization to resolve these questions, since the customer's success is important to the vendor's success and the vendor's reputation is at stake. Will the cloud vendor provide the cloud customer with measurable assurances regarding communication response agreements? An organization may consider the cloud provider an 'internal entity' for policy interpretation purposes. If

Jeff Reed, jeff360@gmail.com

so, what does their being an internal entity mean in practice and will they or can they comply with internal policies? The provider's agreement to your policies seems unlikely unless internal policy requirements are coincident with theirs.

What will a company do if the provider is also looking for the organization to comply with vendor policies? Most, if not all, cloud vendor's terms of use agreements require the customer to comply with the vendor's policies. What conflicts does this introduce for the organization? Can the conflicts be resolved satisfactorily? A business should gain a complete understanding of the responsibilities imposed by their vendor and resolve them prior to contract acceptance.

What has any of this go to do with IH? Violations of policy are incidents. Having either policy alignment or strict policy separation between a business and their cloud vendor helps simplify the task of positive incident identification.

2.3.1. Media

Public communication of incidents can quickly become an issue in cloud integration if not appropriately negotiated with the vendor. Is it all right for the cloud provider to go public with the organization's incident without notifying or negotiating with the customer first? The vendor may not have a choice in the matter. NIST points out that:

An organization may want to—or be required to—communicate incident details with an outside organization for numerous reasons. The incident response team should discuss this at length with the organization's public affairs office, legal department, and management before an incident occurs to establish policies and procedures regarding information sharing. Otherwise, sensitive information regarding incidents may be provided to unauthorized parties; this action could lead to greater disruption and financial loss than the incident itself... The incident handling team should establish media communications procedures that are in compliance with the organization's policies on appropriate interaction with the media and information disclosure. (Scarfone, *et al*, 2008)

Each country, province, state, county, prefecture and city has different regulations and ordinances whose requirements the customer and vendor must meet when operating

Jeff Reed, jeff360@gmail.com

in that jurisdiction. The cloud provider may be required to communicate in ways that could jeopardize the customer. The customer must understand what the provider's policies dictate for these external communication situations including under what circumstances, and using which processes. Are the cloud provider's criteria compatible with the customer's policies and requirements? Care should be taken to work through each possible scenario and account for it based on acceptable risk calculations (Anderson, 2005, Tipton, Henry, 2007, Tipton, Krause, 2007). As NIST suggests, the organization may want members of the public affairs and legal department to participate in all incident discussions with the media due to the sensitive nature and possible legal consequences (Scarfone, *et al*, 2008). While a business may not want to undermine the positive aspects of appropriate disclosure by how they address specific needs in cloud provider communications, the customer will need to negotiate disclosure processes that could affect them. NIST points out the "importance of not revealing sensitive information, such as technical details of countermeasures (e.g., which protocols the firewall permits), which could assist other attackers" (Scarfone, *et al*, 2008). Again, semantics and organizational context could lead to potentially harmful miscommunications if not resolved in the planning phase of integration.

Further questions regarding outside communication arise when considering the repercussions of appropriate disclosure. Is the provider willing to collaborate on a decision making process for incident communications external to them and the customer and put the agreement in writing? A company will need to know how this agreement will be affected by or differ from agreements and processes they have with other providers for similar scenarios. They will also need to assess the impacts of any differences between the various agreements. How can the organization effectively manage the differences as the number of cloud vendor relationships grow (see also Section 4.4 | Scaling)? Find out if the provider is willing to negotiate a joint communication approach. Such an approach may not be possible for them given the volume of cloud customers they may now be, or anticipate serving in the future. Are there regulations that make the company liable for how the provider communicates these issues (as an 'agent' or business partner)? To prevent undesired legal and credibility consequences from an appropriate disclosure,

Jeff Reed, jeff360@gmail.com

collaborate with internal legal, public relations and purchasing departments to get answers to these questions.

Ask the cloud provider how they address the customer's notification requirements. Does the provider notify the customer in ways and in timeframes consistent with the customer's policies and procedures? If the vendor requires the customer to notify them of an incident, the customer needs to ascertain whether they can meet the vendor's timeframes. What if the vendor will not meet the customer's requirements or vice versa? The answer to this question can create serious difficulties for a business if not addressed until late in the procurement cycle. A business will need to define the possible scenarios well enough to provide for each communication need arising from them.

2.3.2. Law Enforcement

Depending on the cloud integration type, the organization may have critical data in the cloud. An IH team will need to know where there data is and be able to retrieve it if they are to meet the requirements of law enforcement. Indeed, the IH team will need to know the location of the data simply to decide which jurisdictions to contact. If an incident response team must appear in court, can they obtain the information and evidence required when the organization needs it? This might seem like a straightforward case of remotely hosted data. However, with cloud integration, it is not straightforward to discover where the organization's data is or where it was when an incident occurred. The cloud customer must determine if the vendor provides sufficient forensic information to meet the law enforcement demands for each affected jurisdiction (see Section 4.3 Legal).

To meet the demands of law enforcement an IH team must make incident documentation available. In describing the appropriate incident documentation level for potential law enforcement use, West-Brown, *et al*/ state:

Essentially the task is to identify the minimum level at which the CSIRT [Computer Security Incident Response Teams] events (especially the incidents) should be documented, and also to identify the right way of doing this. The "minimum" is meant as that which is required by law, and that which may be

Jeff Reed, jeff360@gmail.com

required (or come in very handy) in obvious court cases. The “right way of doing it” means that the evidence (the documents, logs, archives, etc.) should be gathered so that it will receive high marks for completeness (within the set purpose), logic, and reliability when the material is legally requested or is investigated in a court case. This is less trivial than it sounds. (West-Brown, *et al*, 2004)

Complete incident documentation may mean that an IH team provides their own with supporting documentation from their vendor’s IH team. The cloud provider and customer must reach agreement on the alignment of legal requirements. Research the cloud provider’s evidence handling practices to understand if they are equivalent to, stricter than or more lax than the organization’s. The IH team will need to know if the vendor can provide the evidence or information necessary to meet the team’s legal requirements. It is likely that the customer must also meet the vendor’s legal requirements. To ensure that the IH team will have what it needs when called for might only require a review by the team’s legal department of the cloud provider’s written position and procedures. If either their documentation is unsatisfactory or they do not allow others access to this information, the organization may have to bring the respective legal teams together to negotiate a workable solution. Whether the customer must meet the vendor’s requirements or vice versa, the customer must know what the gaps are and if there is a way to reconcile or mitigate the differences.

The prospective cloud customer must also decide what sensitive data to allow the provider’s IH team to see. If the customer’s sensitive data traverses the vendor’s processing paths unencrypted, then the data is visible to vendor personnel. In some cases, it may not be possible to encrypt the sensitive data and still utilize the cloud service proposed. In other cases, it may not be feasible, financially or otherwise, to encrypt the data. The IH team needs to know this so that they can prepare to handle incidents in the resulting environment. If a customer’s sensitive data is traversing the cloud unencrypted, the IH team will need to know that they have sufficient audit log data from the vendor to show who has accessed the data. The vendor may choose not to provide the customer with this sensitive internal information. What are the customer’s

Jeff Reed, jeff360@gmail.com

options at that point? A company can go on with the cloud integration and accept the risk or look for other technology options.

Search and seizure issues are an important consideration for potential legal situations. What does the organization do if law enforcement wants to seize their equipment or data? Does it include, by implication, any of the cloud vendor's equipment? Remember that an organization must consider the laws in all potentially affected jurisdictions. Does the customer have the authority to surrender the vendor's equipment? If not, is the organization liable? What happens in the case where the cloud provider is required to surrender the customer's data or equipment (hard disks, backup media, or other processing equipment) containing the customer's data? An organization's legal department or management will have concerns about this. However, the customer may have no say in the matter. Contracts can only get the organization just so far when it comes to meeting various legal demands related to search and seizure.

If the IH team requires a single point of contact (POC) as liaison with law enforcement during an incident, the team will need to determine if the cloud provider's team structure, policies and procedures can accommodate it. NIST states, "you typically should not contact multiple agencies because doing so might result in jurisdictional conflicts" (Scarfone, *et al*, 2008). What are the impacts of having the cloud vendor contact a law enforcement agency near their operational headquarters and the customer's IH team doing the same? It may be that the laws in the affected jurisdictions require both companies to do exactly that. If not so required by law, then the customer will need to get agreement with the vendor regarding the process for handling law enforcement notification. It may be that the customer cannot get agreement from the cloud provider. The organization must decide if they are willing to assume the risk. Even if the cloud provider does agree to your coordination requirements, there may be other cloud depth concerns to address (see Section

4.1.2	Clouds of Clouds
-------	------------------

).

Assuming that it is legal for a business not to report a specific incident to law enforcement for the affected jurisdictions, the customer and cloud provider will still need to reach agreement regarding delayed reporting or not reporting. The customer must define incident types, circumstances or threshold levels at which they do not want to

Jeff Reed, jeff360@gmail.com

report or file charges. Then they will need to negotiate and contract the terms, process, and other particulars with the cloud provider. What steps must the customer take if the vendor is unwilling or unable to agree to meet their requirements? Understand the costs, credibility impacts, and any other risks to the corporation of inappropriate incident reporting when preparing the cloud services contract.

2.3.3. Internet Service Providers (ISPs)

Communication with ISPs and knowledge regarding the location and availability of the organization's data become critical in IH investigations. How will the organization communicate with the ISPs who are supporting the integration? A cloud customer may have no way to ascertain which ISP to contact for a given investigation given the distributed nature of the cloud. Knowing which ISP to contact may be critical during a live investigation. There used to be a public service television commercial in the US that asked parents the question, "It's 11 o'clock, do you know where your child is?" In the same way, does the IH team know where the organization's cloud data is at all times or at any one specific point in time? The IH team will need to determine whether their cloud vendor can or will make available affected ISP contact information and procedures. Cloud approaches to geographically distributed services and elasticity make it difficult to address network path investigation. Ask the vendor if they have tools for determining the network paths used by cloud transactions. This will be necessary for ascertaining which ISP to contact during an incident. Without visibility into network path information in both real-time and historical ways, it becomes more difficult to pursue an incident investigation. If the cloud provider cannot provide this visibility or the customer cannot get this in the contract, the customer will have to decide what their course of action will have to be.

2.3.4. Other Incident Response Teams

Either an organization or their cloud provider may find it necessary to contact another incident response team in the course of an investigation. It is important to make decisions early regarding what is acceptable so that the business can best negotiate third party IH team contact with the vendor. A business will also need to know whether they or the cloud vendor are constrained from doing so in some way, legally or otherwise.

Jeff Reed, jeff360@gmail.com

Likewise, the organization may have mandatory reporting requirements to other incident response teams that place the business in conflict with cloud vendor requirements or the requirements of the jurisdiction in which the cloud processing or storage takes place.

2.3.5. Affected Parties

Incidents sometimes affect parties that are external to the organization. There are many regulations, domestic and foreign, regarding notification of third parties when an incident occurs involving their data.¹ Being sure to take all possible jurisdictions in mind, here are some examples of questions a business should be asking regarding notification:

- When should notification go to the affected party?
- Does the company have a say in whether, when, or what is communicated to an affected party?
- Which jurisdiction dictates when notification should go to an affected party?
- In which third party notification cases is each jurisdiction authoritative?
- Are there multiple jurisdictions involved and how will the organization know?
- If multiple jurisdictions are involved, are there any conflicts between the respective legal requirements?
- Does the customer know where their data is well enough to know which jurisdiction is affected?
- Does the cloud provider have the ability to provide information regarding the location of data at the time of an incident and affected jurisdiction to the customer?

¹ Websites providing legislation details related to privacy & notification: State Laws Related To Internet Privacy: <http://www.ncsl.org/default.aspx?tabid=13463>, Telecommunications and Information Technology – Issues & Research: <http://www.ncsl.org/default.aspx?TabID=756&tabs=951,71,531#951>, European Union – Information Technology laws: http://europa.eu/legislation_summaries/information_society/, International Privacy Laws and Resources: <http://arielsilverstone.com/resources/international-privacy/>, US Privacy laws, rules, regulations & resources: <http://arielsilverstone.com/resources/us-privacy/>

- Will the cloud provider notify the affected party or will notification fall to the customer?

Find out if any or all of the above has contract language or is in other materials provided to the organization prior to purchase. A business will need to know what the impacts are if the organization does not have a choice regarding notification. Are the impacts enough to prevent the organization from going forward with this cloud integration?

Privacy and notification laws will continue to be in flux in many jurisdictions. Additional regulations are in the works with new ones published frequently. Does the cloud vendor provide a service to inform the organization of legal changes in jurisdictions that affect their customers? If not, the customer must periodically apply to the cloud vendor for new cloud locations and jurisdictions to ensure compliance with the latest laws affecting their data, processing, and security. The processes the customer modifies will need to take into account cyclical research into new regulations when reviewing Sections [2.2.2](#) & [2.2.3](#).

2.4. Team Models, Structure, Staff, & Dependencies

In the cloud case, as in the case of other hosted services, the IH team extends outside of the organization. What are the implications of this for the IH team? This section looks at cloud issues related to models and structures for IH teams, staffing considerations and procurement when moving into cloud computing.

2.4.1. Models & Structure

NIST speaks of three possible models for the incident response team; Central Incident Response Team, Distributed Incident Response Team and Coordinating Team (refer to Scarfone, *et al*, 2008 for more complete information regarding these team models). The first model requires an IH team to be centralized and is not possible in the cloud scenario since the very nature of adding a hosted service extends the concept of the team to include the cloud vendor's IH team. What remains, for cloud integration, will typically be a hybrid of the latter two models, distributed and coordinating. In a distributed IH team model, there are separate incident teams for logical or physical

Jeff Reed, jeff360@gmail.com

segments of the business with centralized IH oversight. Whereas the coordinating IH team model performs as an IH team advising other IH teams over which the former has no authority. The nature of the IH processes the organization defines for their cloud use cases will drive how the team structure looks at any given point in time. The structure will likely tend to morph somewhat over time depending on the integration and the types of situations the organization anticipates. For instance, the decision matrix used in determining when to make first contact with the cloud IH team for Infrastructure as a Service (IaaS) will probably differ from the SaaS case. The team roles responsible for owning the incident will likewise differ.

If the organization outsources some or all of its IH capability, this will affect the model used and, additionally, it may create contractual difficulties. Will the cloud vendor allow collaboration with a customer's third party IH organization? Will they place restrictions on what data access to allow the third party? Much of the information an organization will need to know if there will be issues with involving a third party contractor is required to be in the organization's privacy statements. Much of whether or not the cloud provider can allow a third party IH vendor access to the data may come down to an interpretation of the laws in the affected jurisdictions. Again, the organization will need to stay informed regarding the introduction of new jurisdictions by the cloud integration, and what process the vendor will provide for notifying the organization regarding the introduction of a new jurisdiction by cloud expansion (see Section 4.1.2 Clouds of Clouds). What will need to be in the contract (e.g. processes, procedures, contact information, internal notification, reporting, and access (self-service portal, data, services)) to adequately support the organization's third party IH capability? If the contract does not explicitly provide for a third party IH team, add language to cover all of their access, process, contact and other information needs. Does the organization have mechanisms to ensure that they and their cloud provider are dealing with a valid member of the third party IH team (whether the third party contacts by phone, email, in person, or via any other contact medium)? Will the cloud provider's processes support these mechanisms? Bringing in a third party IH team creates opportunity for social engineering attacks against the company and the cloud vendor. Be certain to close any possible loopholes in communications related to your third party contractor.

Jeff Reed, jeff360@gmail.com

Whether or not a company is dealing with cloud integration, the organization's IH team members will need to be some of the most highly trained personnel in the IT department. However, the team will have the need of additional skills to support cloud IH. Speaking to the need to have a high skill level in IH team members, NIST says: "Technical inaccuracy in functions such as issuing advisories can undermine the team's credibility, and poor technical judgment can cause incidents to worsen" (Scarfone, *et al*, 2008). The IH team will need to be proficient in the specifics of both the organization's and the cloud provider's cloud architecture. The new technologies and the vendor's approach to delivering them will require team training in new portals, cloud administration, APIs, audit capabilities, troubleshooting tools and architectural concepts for integration. There will be education needs regarding business continuity and disaster recovery. What capabilities does the vendor have to enable or support these training activities? Are the vendor's cloud capabilities and tools well documented? If not, ask the vendor if they will agree to contractual provisions for team-to-team knowledge transfer and consultation sufficient for the organization's education needs. There may not be any documentation or training for some of the vendor's cloud technologies and processes. In this case, what options are available to the IH team? The team will need time to learn what they need to make appropriate decisions, consult with cloud vendor technical staff (or other cloud customer technical staff) to understand the issues and to plan, or collaborate in developing a combined plan. The organization will want to negotiate for this consulting and allow time for required preparation if the integration is to be successful.

2.4.2. Staff Considerations

The experience of a company's staff is critical to accurately understanding the indications of an unfolding event. NIST asserts that "technical staff members within the organization usually have much better knowledge of the organization's environment than an outsourcer would, which can be beneficial in identifying false positives associated with organization-specific behavior and the criticality of targets" (Scarfone, *et al*, 2008). Looking at both sides of the relationship, that of the cloud provider and that of their customer, there are different concerns. Can the organization tap into the cloud provider's

Jeff Reed, jeff360@gmail.com

skills directly at the level of infrastructure under investigation? When an incident occurs that crosses into the cloud infrastructure, for all practical IH intents and purposes the cloud provider's personnel are a customer's primary technical resources. The customer requires the cloud IT operations personnel for situational awareness and incident decision-making expertise. Nevertheless, does the contract, cloud vendor's professional services structure and vendor relationship allow for this level of intimacy in IH investigations? Some level of staff collaboration will likely prove to be critical to cloud IH.

Even if the organization can develop and maintain the technical skills related to the outsourced elements of the architecture, can they maintain visibility into ongoing operations for the purpose of attaining an appropriate level of operational awareness? Much of what equals success in this area amounts to following one's intuition or the experiential analysis of the IH team members applied as the incident investigation is evolving. However, some or all of the IT elements under investigation are a 'black box' to the IH team. Understanding the existence and structure of the elasticity mechanisms of a particular cloud component does not provide the experience of seeing it in action. Watching the systems react to multiple tenants, growing and shrinking resources, migrating between servers (both locally and geographically distributed) under various conditions is quite different from knowing how it works. The difference in this understanding can mean the difference in interpreting the incident symptoms as a DDOS attack, a system VM loading issue, a storage bottleneck, a network load balancing issue or something else entirely. Can the organization mitigate the impacts of the lack of visibility and experience? An organization can to some extent by a close collaboration with the cloud vendor's team. The IH team must evaluate whether this is sufficient to satisfy their needs or concerns. Without comprehensive cloud architecture proficiency and visibility into ongoing cloud infrastructure operations, the IH team will be unprepared to engage cloud incidents successfully.

Regarding the issue of multiple cloud locations, does the provider have sufficient IH personnel for on-site investigations across these sites? Do they have enough for all customers? Though it is unlikely that all customers will have ongoing investigations at the same time, it is highly likely that more than one will be under investigation at any

Jeff Reed, jeff360@gmail.com

point in time. Is there a threshold for the number of concurrent investigations after which their IH services break down? How often can the organization sustain the lack of vendor IH team support if this occurs and under what circumstances? What is the likelihood that the vendor will hit that threshold at any given time? Is there any way to mitigate the effects of a vendor on-site IH support gap for the proposed cloud integration type? The IH team may have little or no control over the whether the vendor has sufficient personnel for on-site IH. However, they need to educate management regarding the exposure this lack of personnel presents to the company.

2.4.3. Procurement

The organization will need to work with the procurement department on their processes to achieve the level of rigor the organization requires to address the concerns encountered by cloud integration. Indeed, all sections of this paper speak to the need to collaborate with procurement to some extent. Procurement is typically pleased to have IT teams engaged in helping them ensure their process provides for the best interests of the company. This is a proactive step and covers many areas of corporate endeavor that deal with third party integration, purchasing or collaboration. If this is the first time the organization is dealing with a cloud integration, or if they have not yet matured these processes for cloud cases, the IH team will want to meet with procurement as early as possible. Many of the questions asked throughout this paper will be instructive to them and will aid in developing appropriate decision points in their process for dealing with requests for cloud contracts.

3. Perspective: Incident Handling Lifecycle

Cloud integration multiplies the impacts of the difficulties related to IH. The matter, often articulated as situational awareness, deals with all aspects of the IH lifecycle. Speaking generally of IH, NIST states:

Organizations typically find it very challenging to maintain situational awareness for the handling of large-scale incidents because of their complexity. Many people within the organization may play a role in the incident response, and the organization may need to communicate rapidly and efficiently with various external groups. Collecting, organizing, and analyzing all the pieces of information, so that the right decisions can be made and executed, are not easy tasks. (Scarfone, *et al*, 2008)

The effect of cloud integration on these tasks is to increase the complexity. Why is cloud different from integration with partner organizations? Because, architecturally speaking, the cloud service (application, platform, system and/or network) is internal to the company. Companies typically partition partner integrations from internal systems using firewalls. With cloud, the vendor connections are directly into systems, networks, and applications just as an internal application or subnet would be. The lack of visibility into vital components contained in the cloud hinders the organization's ability to attain complete situational awareness. This may be somewhat mitigated through a variety of services offered by the cloud provider. However, the data provided by these vendor services may be an illusion if an intruder compromises these mechanisms during an incident.

This section addresses cloud challenges related to the steps that NIST recommends for attaining situational awareness (Scarfone, *et al*, 2008). Numerous IH processes and escalations that need to be set in motion by an incident are critical to the organization's situational awareness. The IH team must maintain contacts for both companies and formalize incident response collaboration processes. Notification processes need to be exercised *before* an incident is in progress to ensure that they will work when called upon. Incident prioritization, a critical process element for reducing the risk and impacts of incidents, presents additional problems (See Section [3.2.5](#) [Prioritization](#)). Preparing for communication between incident technical leads in both

Jeff Reed, jeff360@gmail.com

organizations seems straightforward. However, as mentioned throughout this paper, complicating factors in coordination include such things as physical location, reporting structures, competing corporate priorities, and (possibly) language. What can the vendor incident lead share with the customer's IH team, or the customer's IH team with them in the course of maintaining situational awareness? Will the information shared be sufficient for reasonable situational awareness for either team? If not, then what?

All of the planning and preparation discussed in this section multiplies as the organization increases its number of cloud integrations. Getting agreement with one cloud vendor on a definition, priority or approach may create a conflict for a second, or a third. Can an organization with several cloud integrations even successfully coordinate all of the vendors affected by a possible large-scale incident scenario? If so, can the business afford the exercise? Can the business perform the exercise without affecting ongoing IT production and project deployments (that is why the organization is in business)? A company should plan well ahead of the anticipated needs for scaling to meet increased cloud deployments.

The remainder of this section's structure follows NIST's Incident Response Life Cycle model as represented by the following diagram.



01282

Incident Response Life Cycle (Scarfone, *et al*, 2008)

3.1. Preparation

To understand how an IH team's need to prepare for cloud incidents the organization must fully understand the architecture as deployed, planned, and envisioned for the future. Failing to get ahead of needs at the contract stages, including the potential risks and impacts of not being able to provide for IH needs, can be very costly.

Jeff Reed, jeff360@gmail.com

3.1.1. Tools & Resources

An IH team in the course resolving incidents uses many types of tools. The team must properly account for the cloud equivalents of these tools. Will the cloud vendor provide to or work with the IH team's secure communication mechanisms? Will this communication by itself tip off an attacker, especially given that it must traverse open networks? Does the organization have any alternatives? If required, will the organization have access to or the ability to use packet sniffers or other incident investigation tools in the cloud environment? It is easy to see where this would create privacy issues, security issues, and liability concerns due to the multi-tenant nature of the cloud. What are the providers policies regarding the use of computer forensic software in the environment? In an IaaS environment, the organization may be fine. However, how can the organization handle the other cloud types? What support does the vendor provide for the creation of and access to forensic snapshots and backups? Will they, or can they provide the level of documentation necessary for operating systems (OS), applications, protocols, firewall rules, intrusion detection and anti-virus (AV) signatures that the IH team requires? Revealing some of this information may present a problem for the cloud provider's intellectual property. Can they provide network diagrams and critical assets lists, virtual or otherwise, for the related environments? How is all of this accessible (e.g. via self-service portal)? How are OS, application and other media handled? What support does the organization have for physical work in an IaaS environment if it is required? What is in place for secure patch upload and application? What processes are available for restoration and recovery? Take the time to perform the analysis as a team to identify and account for other tools the IH team requires for the cloud integration.

3.1.2. Incident Prevention

The organization will also need to provide for periodic risk assessments to identify appropriate mitigation areas for incident prevention. Will the cloud provider work with the customer on these assessments? Will they provide the customer with enough information about their architecture and environment to allow the organization to determine the risks posed by the vendor's understanding of threats and vulnerabilities? Will the cloud vendor provide their control strategies to their customers to aid in the IH

Jeff Reed, jeff360@gmail.com

team's assessments (e.g. patch management, host/network security, malicious code prevention)? Will the provider do this cyclically according to the organization's requirements? Perform the cloud vendor research and analysis required for appropriate risk prevention.

3.1.3. Large-Scale Incident Handling

Another IH aspect that the contract should provide for is practicing large-scale incident handling. Will the cloud provider support simulated IH exercises? How often can the vendor support such exercises? Cloud represents a significant technological integration for an organization's architecture and being prepared for incidents that cross over into this environment is crucial. Here is a list of some of the large-scale cloud incident exercise questions that the business will want to answer:

- Will the organization want or need the exercises to be "live" with offensive and defensive teams?
- Do the customer and cloud provider have appropriate environments for these exercises?
- Does the vendor impose constraints on the incident exercises (e.g. only once per year, only in development environments, only for one hour per exercise)?
- Are the vendor's constraints acceptable to the company's risk requirements?
- Do the constraints imposed by the proposed cloud vendor overlap with other providers in a way that prevents the exercise altogether (e.g. date/time resource conflicts)?
- Do the constraints overlap with other customers of the provider in a way that prevents the exercise (e.g. personnel availability and scheduling)? If so, then what must the business do?
- Is it possible for the cloud provider's and customer's IH teams to meet face-to-face? Face-to-face meetings for IH team members whether before during or after an incident engagement, provide valuable opportunities to share experience that goes beyond formal incident documentation. In-person communication brings the dimensions of body language and anecdotal, ad hoc and professional intuition sharing to inter-team collaboration.

Jeff Reed, jeff360@gmail.com

- Is it feasible for the customer's IH team to meet face-to-face with cloud provider's IH team given the number of cloud integrations the organization has or envisions having?
- Can the organization afford not to have face-to-face meetings with the cloud vendor and in what cases?
- An organization will have to negotiate costs for cloud personnel planning, exercise, and post-mortem, both internally and with the cloud vendor. Have these contingencies been considered in ROI numbers for the cloud deployment?

If a business is unable to address the needs of large-scale IH in their cloud integration, the business should question the wisdom of proceeding.

3.2. Detection & Analysis

3.2.1. Incident Categorization

The IH team will need to examine their approach to each incident category for cloud consequences and possible coverage requirements that may need to be in the cloud contract. Because cloud integration is different for each company and architecture and there is currently no complete cloud incident taxonomy, the cloud customer must perform the analysis for understanding what incident types are possible for their specific cloud integration. For example, can the business prevent the transformation of a cloud DOS (Denial of Service) attack into a cloud Cost of Service (COS)(see Section 4.1.1 Elasticity Concerns) attack (aka eDOS, economic Denial of Service or Sustainability) and, if so, what steps will the IH team need to take? When and how quickly will the business need to take these prevention steps?

The most common incident category is unauthorized access. An IH team will need to know, if the provider will inform the team when there is an unauthorized access, physical or electronic, in part of the cloud provider's environment that could potentially affect the customer and under what circumstances. There could be legal consequences of the cloud vendor not informing the organization of an unauthorized access depending on the jurisdiction involved. The organization may not know to look for signs that the attacker penetrated from the cloud environment. Furthermore, the cloud provider may not be able to see into the organization's computing environment to detect this (e.g.

Jeff Reed, jeff360@gmail.com

contract constraints or encryption). Will the organization have the audit tools necessary to detect inappropriate usage in their own environment? One of the challenges of the IH team is visibility into all components of the environment. This visibility is difficult when one considers that the environment is not just the integrated cloud component but also includes the underlying infrastructure supporting that component.

The lack of visibility brings up other IH complications. If the cloud provider does not inform a customer of incidents in the cloud provider's environment, how can the IH team (or the cloud provider, for that matter) know if they are dealing with a multiple component incident? Will the vendor provide the IH team with tools and data to determine if what the staff are seeing in the cloud component is part of a multi-component incident? What must the customer do if the vendor will not provide the tools and data? The reverse may also be considered. The cloud vendor may not know that they have a multiple component incident to engage if not informed of customer incidents.

3.2.2. Signs of an Incident

A cloud customer needs to know what the cloud vendor has for incident detection and analysis throughout their infrastructure, processes and personnel and how the customer's IH team might tap into it. Detection typically involves a variety of means including: host and network IDPS, antivirus software and log analyzers. The cloud provider may not allow some or any of these. The chosen cloud type may make it impossible for the vendor to provide their customer with access to the data provided by these tools. In cases where there are no overt signs of an incident, does the provider have the automation (e.g. logs, SEIM, IDPS, and firewall data) to allow the organization to detect other less obvious signs of intrusion? Will the vendor allow a customer access to this detection data and under what circumstances? If the vendor does not allow the organization access to this detection data, it creates an information or visibility gap. The business will have to consider mitigation of any part of this visibility exposure that remains. Does the cloud vendor provide enough cloud architecture detail for the customer to determine what the gaps are? Without detection automation there is virtually no visibility into attack precursors, with the exception of those cases where an attacker announces their intentions. This seriously hinders or eliminates the organization's ability

Jeff Reed, jeff360@gmail.com

to respond to a probable threat. Is this an acceptable position for the organization? If the cloud provider *will* allow access to their detection information, is there a way to integrate it into the organization's existing infrastructure (e.g. logging, or SEIM)? If the cloud provider has provisions for detection, review the design provided to understand any delays introduced in receiving the data. If they will not or cannot allow the IH team access to their detection data, will the vendor allow the organization to establish mechanisms for collecting this data (this largely applies to the Infrastructure as a Service (IaaS) cloud type, though not exclusively)? Detection capabilities are too important to incident response for a business to leave the area unaddressed. If the vendor will allow access to their detection data, there may be constraints to consider.

In some scenarios, the detection capability the organization puts in place may inappropriately expose other tenant's sensitive information if provided to customers without additional event filtering. It would also allow a reasonably funded attacker to purchase similar cloud access and utilize these means for stealth attacks against the organization and the organization's data. These attacks would be difficult or impossible to detect. Therefore, even if the vendor will allow the organization to set up detection capabilities is this necessarily a good thing. A complete understanding of the architecture the organization is deploying into will be crucial to an analysis of the risks presented by multi-tenant detection scenarios.

A business may chose to outsource aspects of their security such as single sign-on (SSO), logging, or SIEM. If the organization has an external security data monitoring or SECurity as a Service (SECaaS) vendor, will the provider integrate with the contracted third party to provide this monitoring? If the vendor cannot or will not, then the design team will need to reconsider the architecture. If the vendor can integrate with the third party service provider, the customer may need to provide provisions in the cloud contracts as well as modify the third party contracts.

Vulnerability and exploit information about deployed applications is essential incident intelligence for IH operations. Some of the applications and platforms that the

cloud provider serves will be COTS² applications and operating systems (OSs). However, they may be serving proprietary applications, especially in the case of SaaS. With COTS applications, there are established information outlets for the latest vulnerability and exploit information, such as <http://nvd.nist.gov>. Some cloud services are beginning to show up in these exploit outlets. However, the organization will need to know that the cloud provider updates at least one of these outlets and which one. Alternatively, ensure that the contract provides this information in a timely manner to the business for the effective management and prevention of incidents.

Another source of precursor or incident intelligence is from personnel within the organization. When either something that is obviously an incident or something unusual happens, the organization's personnel report this to their help desk. This can be the only sign the organization receives that an incident has occurred. Can the organization get agreement regarding notification from the cloud provider when the provider notices something? If so, in what cases will they notify the organization? After what period of time or internal analysis will this notification occur? Referring to earlier comments (see Section 2.2 Policies, Plan, & Procedures), there are motivations that could tempt the provider to withhold needed information from the organization regarding an incident. If the vendor cannot or will not provide the organization with reported incident information, is that acceptable? Is it possible for the organization to put mitigating controls in place for this? A lack of human incident intelligence from the cloud vendor's organization leaves the customer with a significant incident visibility exposure.

If the customer does not have detection data, then they will need to know if they can acceptably mitigate the gaps that remain. Otherwise, the organization will have to do the risk calculations to understand whether the residual risk is acceptable to the organization.

3.2.3. Analysis

An IH team needs to understand their organizations' operating norms so that they can recognize aberrant behavior when it occurs. NIST recommends that, "given the occurrence of indications, the IH team needs to assume that an incident is in progress

² Commercial Off The Shelf

until they can determine otherwise” (Scarfone, *et al*, 2008). This makes it important that the organization understands what the normal operating performance of networks, systems, and applications is in the cloud. This is especially so when the IH team considers that some incidents reflect types of resource issues other than security incidents. The organization will need to treat them the same way until they can make this determination.

Profiling Systems, Networks and Understanding Normal Behaviors

The job of incident analysis is already a difficult one. It becomes more difficult due to the abundance of false positives and false negatives that an IH team must sort through to find evidence of an incident. The IH team can mitigate this to some extent, and in some cases, by using baselines and the judicious use of heuristics. However, complicating the task of sifting through this potentially misleading evidence are various cloud characteristics, such as multi-tenancy and scaling. It may be virtually impossible in some instances to get reliable baseline or profile data from which to perform analysis. How will the organization handle this? The cloud provider may have some answers about baseline and profile data for the organization. Ask them. It is important to understand the impacts of this on the organization’s incident responses and that the IH team researches the cloud provider’s capabilities in this area.

The need for profile and baseline data becomes more important for analysis due to the operations teams’ need for this same data. The operations teams will need to know when to allocate additional cloud resources due to various usage scenarios. The IH team should engage with them early in the purchasing process and in defining the various teams’ respective incident response processes. When the organization sees rapid increases or decreases in network or system utilization, they will need ways for quickly determining the cause. When symptoms occur, the customer may need immediate assistance from the cloud provider to make these determinations. What sort of response times will the customer need from the cloud provider? If the organization needs a highly qualified cloud IH team member to engage with the organization immediately, the SLAs will need to reflect this. Until it is clear that the event the organization is seeing is either “healthy growth” or a “healthy lull” as opposed to a resource problem or security

Jeff Reed, jeff360@gmail.com

incident, the IH team will need to remain engaged, as may the cloud provider's team. Unless the cloud vendor is providing free incident support to the organization, there will likely be real costs to the organization during incident troubleshooting and therefore this should be taken into consideration.

Central Logging & Log Retention

Network, system and application logs are the foundation of incident analysis. As has been stated earlier the organization's cloud services may leave the organization with limited or no access to this essential incident intelligence. It may not matter in all cases, so the organization will have to examine the particular integration to determine if it does for the organization in this instance. If it matters to the organization, what options are available? A careful investigation of options is important. Gartner made the following observation regarding logging concerns:

Cloud services are especially difficult to investigate, because logging and data for multiple customers may be co-located and may also be spread across an ever-changing set of hosts and data centers. If the organization cannot get a contractual commitment to support specific forms of investigation, along with evidence that the vendor has already successfully supported such activities, then your only safe assumption is that investigation and discovery requests will be impossible.

(quoted by Brodtkin, 2008)

It may not matter to the organization that the support does not provide forensic data that is usable in court. It may be sufficient that the vendor can provide the support necessary to find and isolate or eliminate the cause. In either case, the company will have to decide what level of residual risk they can tolerate.

The cloud customer may have legal or policy requirements around log retention. Can the cloud provider meet those requirements or work with the organization to meet them? If the cloud provider makes these logs available to the organization then the onus for retention will be on the organization. If not, then the organization will need to negotiate this with the vendor.

Event Correlation

Whether the organization does event correlation manually and intuitively, uses an SEIM, or contracts with a third party, they will need good log and event data. If the answers to the questions raised earlier leave the organization without this data, can the organization count on an accurate understanding of the cloud-based integration landscape? Depending on the size of the cloud integration, this could be a sizable visibility gap. Because each new cloud integration scenario is different, the IH team will need to perform this analysis for each case.

Filtering Analysis Data

If the organization is able to obtain the cloud environment's log and event data there is another consideration. If the volume of data is too large or affects other critical business processes, it may be necessary to filter the data at the cloud. Does the vendor provide a mechanism for filtering log and event data before forwarding to the customer's site? It may mean that the point of integration for this data is at the cloud provider and that the organization provides a mechanism for this. If so, this will need to go into the integration architecture and into the organization's ROI calculations.

Clock Synchronization

A recurring concern in the cloud forensic and IH debate is time synchronization. Keeping system and network device time in sync with these devices spread across the globe is problematic. This is especially true when the organization must meet the rigor of legal evidence. Irrespective of the approach used by the cloud provider, will it meet the legal forensic requirements of the organization? There may be new training that is required for the legal and IH teams to ensure that the IH processes do not invalidate the use of the data in court. Does the cloud provider support this training activity in ways that meet the organization's requirements? Time synchronization will need to be considered carefully as the cloud customer integrates the new cloud service.

Incident Analysis Knowledge Base

Research what the cloud provider has by way of a knowledge base that the IH team can tap into for understanding such things as deployment into the vendor's

Jeff Reed, jeff360@gmail.com

environment, recognizing ‘normal’ events, and understanding issues with deployment tools. Some of this will come in the form of Frequently Asked Questions (FAQs). However, the IH team will require specific IT operations and IH related information. Does the cloud vendor provide other IH supporting information, such as a diagnosis matrix to aid in focusing an ongoing incident investigation? This type of information can save the IH and operations teams time and save the organization money. An understanding of the tools, information and data provided by the vendor can help the team make quick determinations regarding event causes.

Encryption Keys

There are many details to resolve surrounding key management in the cloud. The ENISA report provides a good overview in the description of some of the issues in their vulnerability v11 (Catteddu, Hogben, 2009). What primarily matters to the IH team is, where are these keys stored and how will the team get access to them, when required, for an investigation? The organization’s design team will be one source for this information. The IH team will have to understand what happens to these keys during cloud related events, such as: scaling, elasticity, virtual host migration, VM snapshots, backups, etc. When there are new resources allocated for upward scaling, keys created, and then the resources released, are the affected keys retained at all? Once again, the IH team will need to ask, at the time of the incident under investigation, where was the cloud resource that was being accessed for key exchange. The IH team will need to know how to discover this if they are to find the keys when required. The cloud provider or the organization’s design team may have answers. In either case, the IH team will need to know what the answers are in order to investigate an intrusion effectively.

3.2.4. Documentation

Critical to both court cases involving incidents and process improvement is the need for documenting each incident case appropriately. Do the cloud vendor’s processes for documenting incidents meet the customer’s requirements? They may operate with a framework that fully meets the jurisdictions in which they operate; however, it may not meet the demands of the jurisdictions in which the customer operates. Additionally, if the customer is required to produce their incident documentation for the vendor, does it

Jeff Reed, jeff360@gmail.com

meet the cloud vendor's requirements? An organization needs to know whether they have to meet the vendor's requirements and what those requirements are. The companies incident documentation processes need to reflect the differences imposed by cloud integration.

3.2.5. Prioritization

NIST (Scarfone, *et al*, 2008) and CMU (West-Brown, *et al*, 2004) suggest that organizations should create written guidelines for prioritizing incidents. SLAs should include language addressing this to the organization's satisfaction. This should take into consideration prioritization from both the cloud vendor's and the organization's points of view. Because either organization may detect an incident first, a written agreement between the organizations on correct prioritization can make a big difference in the delay between the start of an incident and the start of real engagement. Agreeing on incident prioritization may be problematic, however. As mentioned at the outset of this paper the customer will have one perspective of mission, regulatory demands and other requirements while the cloud vendor's view may be very different (again, refer to Section 2.2 Policies, Plan, & Procedures). If the vendor is working on an incident that may potentially affect a customer, what is motivating them regarding when or whether to notify the organization? What must a customer do if they disagree with the cloud vendor's decision rationale? What will happen if the customer never knows what the vendor's prioritization rationale is? A business will need to decide if the possible gaps are significant enough to eliminate the cloud offering from the organization's enterprise solution. NIST highlights this by saying, "Incident handlers may be under great stress during incidents, so it is important to make the prioritization process clear" (Scarfone, *et al*, 2008). There is a compounding of these effects due to teams unfamiliar with one another, with differing (and potentially competing or unstated) priorities, and with different or incompatible processes. Who contacts whom? At what point in time? At what level in the organization structure? The confusion resulting from a lack of clarity regarding prioritization can result in significant financial and credibility costs.

Jeff Reed, jeff360@gmail.com

3.3. Containment, Eradication, & Recovery

3.3.1. Containment Strategy

It is very likely that the cloud integration will affect the organization's current containment strategies. The IH team must understand what is necessary to implement containment related to the cloud integration. This is easier said than done. Can the IH team block cloud traffic for ingress or egress if needed? If the organization requires this, put it in the SLAs. Any process for containment that involves contacting the cloud vendor and engaging their personnel to implement access controls, introduces delay between discovery and removed access. During this time, the organization's systems or applications may be controlled and used to launch further attacks. Do delays caused by the cloud provider relationship expose the organization to unacceptable liability? Some cloud providers may have mature, quickly deployable mechanisms to assist the organization in these cases. Organizations should know what they are, how to use them, their costs and whether to ensure they are in the contracted list of services. What happens if the IH team determines that the organization needs to take a host, network or service offline, but cannot reach the self-service portal (e.g. in a multiple component attack that includes a DOS attack)? In this case, the customer would need to write the procedures to include steps for contacting cloud support. If the customer or the cloud provider determines that a cloud service is infected, can that service be isolated? Are there instances when a provider may choose to do so without first notifying the customer? What happens if that service is part of an interconnected chain of services that make up a larger production application? If needed for an investigation, can the organization get access to log or journal records regarding the measures taken by the cloud IH team? An organization's IH team will need to define the notification processes well ahead of implementation for this possibility. The organization will also want to ensure that the vendor provides them with clear reasoning supporting their decision and service restoration timetables for the IH team's notifications to senior management. A customer should also think the matter of containment through from theirs and the cloud provider's side of the integration. The organization has containment requirements and so does the cloud vendor. Does the cloud provider have containment strategies that could cut off the customer's service? Under what circumstances would this resource isolation occur?

Jeff Reed, jeff360@gmail.com

What level of notification does the vendor process provide to the organization? A customer should decide if this is an acceptable risk for them. How would such a cloud vendor containment decision impact the organization's ability to investigate the matter from their side? Such a move on the vendor's part could sever the IH team's visibility into a component of the enterprise architecture during an investigation. Containment, in cloud integration cases, creates challenges for vendor and customer alike. A company must carefully analyze the potential containment cases and negotiate mutually agreeable processes for containment decision and execution.

Containment can also have unwanted side effects, which vendor and customer must allow for. NIST describes a reasonably common scenario where the malicious code on one host watches another host, or the services on another, to detect discovery (Scarfone, *et al*, 2008). In their scenario, when one host is no longer visible on the network, the other proceeds to overwrite all the data on the host's hard drive. The intruder or malware may also take other unwelcome actions. This type of scenario highlights the need for a highly collaborative approach to incident handling between the customer and the provider. It implies a high level of communication and trust.

As an organization increases the number of cloud integrations, these relationships will become more difficult to develop, maintain and coordinate. As the organization steps through the process, analyze the impacts this will have on scaling the processes and relationships (refer to Section [4.4 Scaling](#)). Also, understand what the provider has in place to address these same scaling concerns.

3.3.2. Evidence Handling

With evidence collection and handling, we once again face the issue of many jurisdictions. The IH team must know where the organization's data is and what laws apply to such things as data collection, retention, forensic handling, and admissibility in court. Even if the organization is not serving a national or global user base, the cloud integration may take the organization across many jurisdictions. The organization will need to understand these jurisdictions, calculate their costs and risks, and accommodate them accordingly.

Jeff Reed, jeff360@gmail.com

NIST suggests that the IH team take an as-is system snapshot early during an incident investigation (Scarfone, *et al*, 2008). The team takes the snapshot prior to other incident handling activities so that the state of the system at the time of the incident is preserved. Is it possible take such a snapshot in the cloud integration scenario? Virtual machines can simplify the task. However, elastic scaling may make taking a snapshot difficult or impossible. Again, the various types of cloud integrations will yield different answers and thus the team must analyze each situation carefully when planning cloud evidence handling.

3.3.3. Attacker Identification

Though called out as a time-consuming and futile process by NIST, they also maintain that attacker identification is, nevertheless, an aspect of IH that often needs to be tackled (Scarfone, *et al*, 2008). Once again, cloud integration can make this difficult or impossible to do. It can be a complicated task just to identify the attacker's IP address due to a lack of visibility into logs and other intrusion monitoring data. If, as was suggested earlier, the organization is able to gain agreement from the cloud provider for this data, it will simplify the task. If not, then what? If the cloud deployment includes an Internet content delivery network (CDN) such as Akamai® or Internap®, the organization will have another layer of indirection to work through to identify the attacker. Nevertheless, it is important to consider how to perform attacker identification in a complex cloud environment.

3.3.4. Eradication and Recovery

This section describes some of the same issues a business will grapple with when addressing internal disaster recovery and administration (See also Section 4.2.1 Backups, Disaster Recovery (DR) and Business Continuity). However, if a cloud vendor owns the operating systems and applications, are they willing to work with the organization to harden them if the results of an incident investigation indicate that it is necessary to prevent further attack? If so, the organization will need to negotiate how long this will take as a part of the SLAs. The cloud vendor will typically be using standard images. The time it takes for their modification, test and deployment processes may put an organization's production IT systems at risk for some time. What is the difference

Jeff Reed, jeff360@gmail.com

between the cloud vendor's processes and the customer's processes for making changes to production systems and applications? The difference is that, the customer does not need to negotiate for resources with the cloud vendor while competing for priority with other national or international interests. A business must be certain that the processes and effects of cloud eradication and recovery processes are acceptable.

3.4. Post-Incident

If the organization has a post-incident lessons learned process, they may want the cloud vendor to be involved in this process. What agreements will the organization need with the cloud provider for the lessons learned process? If the cloud provider has a lessons learned process, does management have concerns regarding information reported or shared relating to the organization? The cloud vendor will not be able to see much of the company's processes, capabilities, or maturity. The company may have concerns regarding how much of its internal foibles to share. If there are concerns, get agreement internally first, then negotiate them, if possible, and have them written into the contract. If the vendor will not or cannot meet the customer's process requirements, what steps will the organization need to take?

An IH team collects and analyzes incident process metrics for trend and process improvement purposes. Like any other organization, the cloud provider will be collecting objective and subjective information regarding IH processes. As NIST points out, the use of this data is for a variety of purposes, including justifying additional funding of the incident response team (Scarfone, *et al*, 2008). Will the organization need this IH process metric data from the provider to enable a complete understanding of the integration area in case the organization ever has a need to bring the cloud function back in-house? Will the organization need this data for reporting and process improvement in general? The use of this data is also for understanding trends related to attacks targeting the organization. Would the lack of this attack trend data leave the organization unacceptably exposed to risk? Determine what IH process metric data is required by the team and write it into the contract.

The organization will need to decide if they require provisions with the cloud provider regarding their evidence retention policies. Will the vendor keep the evidence

Jeff Reed, jeff360@gmail.com

long enough to meet the organization's requirements? If not, will the organization need to bring the cloud vendor's evidence in-house? Will the vendor allow the customer to take custody of the evidence? If the vendor retains the evidence longer than the customer's policies dictate, does this create risk for the customer? If so, what recourse does the customer have? Legal counsel will need to provide direction in this area in order to ensure compliance with laws for all jurisdictions.

3.5. Incident Categories

In NIST's *Computer Security Incident Handling Guide*, Sections 4-8, various incident types are dissected for the purpose of understanding how IH processes and analysis should prepare to deal with them (Scarfone, *et al*, 2008). It is advisable that the IH team review each of these sections as it relates to the prospective integration. Though this section addresses some of the questions raised by a review of the categories outlined by NIST, it does not delve into each in detail.

Denial of Service (DOS)

Given the virtualized and multi-tenant characteristics of cloud computing, the organization will need to broaden their perspective when considering the impacts of DOS attacks. If someone has initiated a DOS or DDOS attack against another tenant on the same virtual server, what tools are available to detect the affect of this on the organization? If the cloud provider has the capability, they can put in rate limiting controls to mitigate some of the impacts. However, over provisioning is standard practice among cloud vendors which affects the extent of rate limiting's usefulness in the DOS case. From an IH perspective, if the vendor does not know, how will the organization know what is causing the service impact? If the organization attempts to scale with additional virtual resources, will this help or hurt (See Section 3.2.3 Analysis)? Again, how can an organization know whether to scale or constrain access to the cloud services or systems? It is important to find out these details prior to an incident.

DOS and DDOS IH add another aspect of access limiting to that described in an earlier section. As a preventative or impact-limiting step, does the cloud provider have a mechanism to allow the customer to limit inbound and outbound traffic to permitted services only? The answer the organization is looking for will be different depending on

Jeff Reed, jeff360@gmail.com

the cloud type. For example, if the customer is controlling some aspects of the virtual environment such as virtual firewalls, then the vendor will need to provide traffic rules at a higher level of environment abstraction (e.g. physical firewalls) to prevent overload of the organization's resources. This can limit some of the COS attack impacts the organization may otherwise see.

If the organization experiences a DOS attack that requires shutting off the source (IP addresses, ports, etc.) of the offenders, the organization will need cloud provider response SLAs that meet their needs. What happens if the traffic must be blocked at the cloud provider's ISP? This may seem silly because "Cloud providers are really large and surely have rapid response agreements with their ISPs, right?" While it is likely that this is true for large cloud providers, the organization should take nothing for granted. What about smaller cloud service providers who are, themselves, using another cloud vendor as a platform for their services? Have they contracted with their "ISP" (their cloud provider, in this case) in the same way? Engagement with cloud providers must start with no assumptions if the organization is to ensure the ability to respond to incidents according to business needs. (The issue of user adoption of a business' cloud services appearing to be a DOS/DDOS attack is addressed in Section

3.2.3	Analysis
-------	----------

)

Malicious Code

In the case of SaaS, if there is a covert insertion of malicious code into the provider's application, will the vendor ever notify the organization, when the vendor eventually discovers the event? (See Section

2.2	Policies, Plan, & Procedures
-----	------------------------------

 for a discussion of motivations regarding incident response and reporting.) Even if they follow a policy of disclosure regarding their own code, what of infections to other COTS code on those systems and servers functioning as the infrastructure for their applications? The possible intercept points in the code paths for the organization's sensitive data multiply when considering the virtual and elastic architecture in use. An IH team must consider the impacts of this on the organization's investigations and if they can put any mitigating controls in place. A complete discussion of the possibilities is beyond the scope of this paper. However, it is important that the IH team resolve these questions to their satisfaction before cloud integration.

Jeff Reed, jeff360@gmail.com

Blended attacks combine a number of attack vectors and transmission methods to extend their penetration impacts. Does the cloud vendor provide sufficient controls across all possible vectors? If they do, for IH planning, does the IH team have visibility into these controls and their logging and monitoring mechanisms? If not, does the vendor provide support for IH collaboration if or when the organization's IH team feels that it is the logical next step in the investigation? A combination of multiple attack vectors may well include components of the vendor's applications or infrastructure. Visibility into the vendor components, at least their log files, becomes critical to the IH process for blended attacks. The cloud vendor should also have the same concerns. An intrusion in the customer's systems can become an attack vector into the cloud infrastructure. The customer should ensure contracts contain support for the contingencies represented by blended attack scenarios.

Rootkits are always difficult to detect and thus they present a unique and difficult concern for cloud computing. In some cloud types (e.g. Hardware as a Service (HaaS) and IaaS) the organization may have some visibility into things like root compromise of a host, in others the organization may not. The further the organization gets away from the physical platform, the more difficult the task becomes. What level of support does the vendor provide for detecting rootkits? If the cloud provider discovers something, will the customer be informed? When? Can the organization have access to the rootkit that has penetrated the vendor's system for investigation purposes? The IH team will need to decide if they require access to the rootkit or if the cloud vendor's evaluation is sufficient. How quickly can the cloud vendor deliver an analysis to the IH team? Though the previous question may not matter ordinarily, if the vendor never delivers details, evidence or analysis to the customer, it might matter after all.

Unauthorized Access

Cloud services usually provide a self-service portal or other administration tool for allowing the customer to scale the environment, modify roles, add users, and other necessary cloud tasks. What happens if an intruder gains access to the organization's self-service portal? In the ENISA report, the editors highlight that customer access to an Internet accessible management interface for large sets of resources poses an increased

Jeff Reed, jeff360@gmail.com

risk (Catteddu, Hogben, 2009). This is especially true when combined with an understanding of web browser vulnerabilities. A business' IH team must work with the rest of their security team in understanding the capabilities, potential vulnerabilities, and audit logging of the self-service or administrative portal.

Another aspect of unauthorized access to consider is how cloud security architecture affects it. NIST's *Computer Security Incident Handling Guide*, Section 6.2.2, Incident Prevention, discusses both process and technical actions for preventing unauthorized access such as restricting physical access, disabling unneeded services, and using DMZs (demilitarized zones) for publicly accessible Internet services (Scarfone, *et al*, 2008). Does the organization thoroughly understand the security architecture serving the cloud application? How much control does the organization have over the security architecture serving the cloud application? A complete understanding of the security architecture is essential to the IH team. Without a thorough understanding, the team will not know how to discern authorized access from unauthorized access. The team will not know what event data to review; neither will they know how to interpret it.

Vulnerability assessments are an essential component of an IH capability. What are the costs and risks of running vulnerability assessments against the cloud services? Will the provider even allow a customer to perform vulnerability assessments against the cloud services? If a customer makes the determination that vulnerability assessments affecting the cloud integration are required, then the cloud vendor's position on them is important to the customer's direction.

Reconnaissance activity is often the first indication a company will have that they are the target of an attack. In the case of cloud services, is it possible for the organization to have any visibility into potential reconnaissance activity? This is critical information for potentially stopping an intrusion before it happens. Without some visibility into reconnaissance activity, the customer is at risk. If cloud vendor will not provide reconnaissance data, a customer must quantify and find ways to mitigate the risk before proceeding with the integration. Unauthorized access is more likely than most other incidents to lead to prosecution, so having access to this data is important for evidentiary purposes if for no other reason (Scarfone, *et al*, 2008).

Jeff Reed, jeff360@gmail.com

Multiple Component

The issues represented by a multiple component attack are similar to those of a blended attack, as discussed in Malicious Code above. NIST defines a multiple component incident as, “a single incident that encompasses two or more incidents” (Scarfone, *et al*, 2008). NIST goes on to state that this complicates the incident analysis process. Introducing cloud into a multiple component scenario can complicate an already complicated process by imposing additional processes, procedures, approvals, and delays. It also introduces a visibility gap that the organization may or may not be able to bridge depending on the vendor’s policies regarding disclosure.

4. Additional Considerations

This section goes beyond the NIST framework to review some additional cloud IH considerations.

4.1. Structural Perspectives

This section will highlight a couple of unique cloud structural characteristics. They represent the breadth and depth of cloud computing. They also represent both the greatest opportunities and the greatest challenges of integrating with the cloud.

4.1.1. Elasticity Concerns

Cloud computing has been built on industry developments dating from the 1980s by leveraging outsourced infrastructure services, hosted applications and software as a service (Owens, 2010). The addition of elasticity to this aggregation of approaches makes it compelling to CIOs in companies of all sizes. NIST's definition of cloud computing v15, included in its entirety in Appendix Section 7.1 defines elasticity this way:

Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time. (Mell, Grance, 2009)

The cloud offers enterprises the ability to get their businesses running without having to pay significant upfront capital costs and the ability to scale the IT infrastructure up or down as the organization progresses without worrying about over or under provisioning. Elasticity, combined with pay-per-use, is arguably the most important innovation represented by the technology we are referring to here as cloud computing. Dustin Owens, a senior principal consultant with BT America's Business Innovation Group, in a Jan. 2010 Communications of the ACM article, *Securing Elasticity in the Cloud*, states:

Elasticity, in my very humble opinion, is the true golden nugget of cloud computing and what makes the entire concept extraordinarily evolutionary, if not revolutionary... When combined with on-demand self-service capabilities it could

Jeff Reed, jeff360@gmail.com

truly become a game-changing force for IT... elasticity could bring to the IT infrastructure what Henry Ford brought to the automotive industry with assembly lines and mass production: affordability and substantial improvements on time to market. (Dustin Owens, 2010)

With all of this positive press, one wonders why there is not a mad dash to embrace the cloud. The substantial benefits which this technology brings to the enterprise also carries with them significant security hurdles. Though this paper does not address the broader security issues, it does look at the IH team's concerns regarding the various characteristics comprising this core cloud attribute, namely elasticity. To emphasize the unaddressed concerns in the new but unstable space of the cloud, Owens goes on to say:

Enlightening as this realization has been, it has also become clear that several monumental security challenges (not to mention many monumental nonsecurity-related challenges, not the least of which are fully functional availability and how well an organization's environment is prepared to operate in a distributed model) now come into play and will need to be addressed in order for the elasticity element of cloud computing to reach its full potential. (Owens, 2010)

Owens rightly points out that most of the discussion about cloud focuses on the challenges related generally to IT outsourcing and which are not new to the industry (Owens, 2010). Such IT outsourcing topics as vendor access to customer's data, perimeter security considerations, DOS/DDOS attacks, resource exhaustion and compliance challenges are not new to the industry.

Some of the more significant issues raised by Owen in the rest of his article are broken down and addressed below.

Virtualization – Hypervisor Vulnerabilities

Enabling elasticity without the use of virtualization is difficult. Chasing incidents in a virtualized environment that is fully owned, housed and maintained by the organization is not new. However, pursuing incidents into a virtualized cloud deployment is less clear.

Jeff Reed, jeff360@gmail.com

Consider the following scenario. A larger or more experienced organization X has its resources deployed into a cloud provider's virtualized environment. A smaller or less experienced company Y deploys onto the same virtualized environment, on the same hypervisor. Although company X has thorough controls and processes for protecting their perimeter and even encrypting their local files, company Y does not. A hack through the hypervisor can make company X's protections moot. Company Y's posture now puts Company X's environment at risk.

Clearly, company X must involve the cloud vendor's incident team in the investigation if they are to be able to pursue an incident. Indeed, the company's first indication of a problem may come from the cloud vendor as this attack vector is nearly impossible to detect from the guest OS. The attack would be even more difficult to detect than an insider attack since it could be carried out without the need to log in or authenticate. Mr. Owens makes clear the issues vendors and organizations will need to address in his statement:

Cloud providers will need to be prepared to account for and show how their particular services are able to control vulnerabilities such as [client VM traversal through the hypervisor] and keep similar yet-to-be discovered vulnerabilities from having devastating impacts on their customers. Perhaps more importantly, *critical infrastructure* (see http://en.wikipedia.org/wiki/Critical_infrastructure for definition) could be subject to insurmountable risk and/or loss of sensitive information if providers lack the necessary controls. (Owens, 2010)

The capability of properly investigating incidents, particularly live incidents, is arguably one of the 'necessary controls' for limiting loss in the virtual environment. Can (or should) the organization negotiate with the cloud provider for a suite of virtualization security controls for this environment? It would not be a simple task just to decide what to require. It may not be possible to provide one tenant with this level of virtualization controls without putting other tenant's sensitive data at risk. If the vendor agreed to provide the organization with the requested security controls (or already carry such protections), will the organization need to access or integrate the log messages created from this infrastructure? Is it even possible for the vendor to provide this level of log

Jeff Reed, jeff360@gmail.com

integration in a distributed, elastic environment where there are multiple tenants? Will the vendor allow the customer to have the log data since this could put other tenant's sensitive information at risk? Would the customer want the vendor to provide it knowing the converse is also true? If the customer did want it, can they support the volume or cost of this additional service and the attendant log volume? If this level of monitoring is not available from the cloud vendor, is the gap in event visibility acceptable to the customer? A business' IH team must collaborate closely with the vendor's IH team in pursuing rootkit-related incidents. The business must negotiate the terms of such collaboration to meet IH requirements.

One possible solution is per-tenant hypervisor isolation. Can the cloud vendor provide the organization environment separation from other clients? If they can, though this may ameliorate the challenges presented by the earlier example, does it sufficiently mitigate the issues stated as questions above? The cloud provider can mitigate many of the concerns raised regarding hypervisor vulnerabilities by providing an isolated environment and by providing "sole hypervisor (or hypervisor-like) management access that connects only to the customer's virtual environment" (Owens, 2010). However, would such a service be financially feasible for either the customer or the cloud vendor? It not only undermines the cloud services business model for the vendor, but it also undermines the cost advantage to the customer. Would it be practical for either the organization or the vendor to support an isolated environment? For the cloud vendor and the customer, there are scaling concerns. The cloud vendor would have to maintain special environments for each customer that required isolation. For the customer, there would be additional administration and potentially additional skills required (e.g. different virtualization technology). Could or would the vendor allow the organization additional hypervisor-level controls with remote access? If so, would this limit the scalability/elasticity of the particular cloud type involved? The customer needs to determine if remote hypervisor access introduces other security concerns. Could the organization effectively tie the vendor's hypervisor-level logging to the enterprise logging? Can the hypervisor-level logging integrate with a third party logging vendor? As in the earlier discussion regarding third party logging, the customer will need to

Jeff Reed, jeff360@gmail.com

determine the costs of logging integration. The IH team will also need to modify their processes to include cloud vendor hypervisor-level investigations.

Fine-Grained Access and Predefined Security Controls

The cloud vendor must provide access and security controls sufficient for regulatory and best practice constraints such as separation of duties. To do this, network paths and server surfaces must be restricted to only those protocols, sources and destinations that are required. Additionally, an organization typically puts in place controls to ensure that, for instance, developers do not have access to production environments. The IH team will need to know; whether these controls are available, how they work, how to deploy them, what roles are defined, what audit capabilities are in place, and how they can integrate their tools and processes. What follows are some areas of particular importance to the IH function.

Various aspects of IH require an organization to constrain networks, hosts, and applications. Does the contract allow the customer to access the network devices (e.g. firewalls, routers, OS-level firewalls, self-service controls, and application controls) for reducing network attack surfaces as required? Some of these devices are possibly servicing other customers. Can the cloud provider allow the customer, or any customer for that matter, to have access? If not, can the vendor respond quickly enough for the IH team's requirements? Does the customer need specific SLA language in the contract? The customer will have to decide what risks the cloud offering introduces to their environment, caveat emptor. In any case, a review of required network, host, and application controls for IH process requirements is essential to the team's success.

As part of an IH investigation, an IH team will also need access to historical versions of configuration files. Is the cloud services provider maintaining historical configuration snapshots for investigative purposes? Does the contract allow the IH team to get this historical information if needed or, at least the rules or access control lists that pertain to accessing the organization's systems? Find out if the vendor has the information and the processes your team will need to follow to get it. The customer should ensure that response times written into the SLAs.

Jeff Reed, jeff360@gmail.com

When looking at separation of duties the organization will need to research what role framework the vendor provides. Going back to the self-service portal, does it provide the flexibility the organization needs for multiple and parallel environments (development, test, pre-production, production, post-production, multiple development and integration paths, and so forth) to meet enterprise requirements? Determine if the portal, given the organization's roles, provides sufficient control over access for each role by environment, at the appropriate level of platform (i.e. OS, virtual hardware) and access (e.g. by environment, file system hierarchy, specific programs, processes, databases, servers). The customer will need to know if portal administration provides audit logging for add, change and delete events at the level required for IH (and audit). Additionally, know if it provides control over who can expand or contract services for each environment defined (elasticity settings for the environment). Inappropriate modification of cloud scaling can appear to (and, in fact) be an incident that would require the IH team's response.

A cloud customer must find out if there is sufficiently granular control at the self-service portal to provide 'minimum necessary' access to the portal and other controls for the organization's environments. The customer will also need to know if there is an audit log that is recording who made portal and control changes where, when, and to what. What are the consequences to the organization and the integration if any of these access and security controls are not available to the level required by the customer? Are there ways to mitigate the gaps? If the answers to these questions are not positive, research whether the vendor can provide sufficient IH support processes, forensic data, legal support, and response SLAs to make up the difference.

Configuration and Change Management

Regarding configuration and change management, Mr. Owens says: "Even where a portal is capable of granular-access controls that control which actions a given user is able to perform, it also needs to enforce *when* and *under what circumstances* a user is allowed to perform certain actions" (Owens, 2010). An organization needs proper change management controls for their cloud service to prevent untested code or system changes contrary to policy. These events can be, and often are, interpreted as security

Jeff Reed, jeff360@gmail.com

incidents. The IH team will need log, journal and other data from the cloud change management controls to help in identifying the source of problems.

Change and configuration management affects what the organization has control over (e.g. source code movement or configuration file changes) and those portions of the environment over which the organization has no control. The cloud provider may control updates of firewalls, routers, firmware, virtual environment, OS, application servers, host-level firewalls, host-level IDS, and more. Does the contract include notice of and historical information for these events? Both vendor and customer change management requirements written into the contract support successful interactions as critical changes occur.

There are issues with dormant (“switched off”) VM images, VM baseline images and image maintenance to consider in a cloud change management. The cloud provider is bringing on- and off-line VMs from image storage to provide the customer with the elasticity needed. When looking at historical events concerning VMs that the environment has brought online at any point in time, the organization needs to ask whether they have access to the patch history for all of the OS, and application software for their configurations. Determine whether the cloud vendor maintains audit log data regarding specific VM image changes made, when and by whom. Additionally, ask if they will make this data available to the IH team. If not, the organization will need to decide what risks this presents to them.

Audit Trails

Though already discussed several times in the context of elasticity, the IH team should carefully consider audit logging for the cloud deployment. Determine the full range of audit information that the integration requires for forensics, audit and compliance. If not provided to customers as a standard offering, ask if the cloud services vendor provides this data to the customers on special request. Also, ask if the cloud vendor will make audit log data available to the organization in the timeframes required for audit reviews and IH investigations.

Jeff Reed, jeff360@gmail.com

Data Remanence

Data remanence requires a quick definition for a better understanding of how it affects cloud integration. The American Heritage Dictionary of the English Language, Fourth Edition, defines remanence as “The magnetic induction that remains in a material after removal of the magnetizing field” (Houghton Mifflin Company, 2000). Data remanence deals with the magnetic “impression” of the data that is left on a hard disk or in memory after the data has been deleted from them. An organization will, typically, when disposing of old magnetic media put it through a rigorous process to eliminate all traces that would otherwise remain through remanence. As the cloud releases VM or database instances, what processes do the resources undergo to remove data written when the organization occupied that space? Are these removal processes sufficiently rigorous for the organization’s requirements? The IH team should ask if they have the ability to validate these removal processes. If another customer’s data remains in resources brought into the organization’s environment, and an employee or contractor is able to retrieve remanent data (e.g. ePHI, credit card data, or police investigation reports) from that resource, is the organization liable? How can one be sure? These questions need to be answered for all jurisdictions. If an organization must integrate with cloud services, this may force them to encrypt all of their data due to the data remanence issue. Does the organization’s internal application architecture allow for processing this encrypted data? Does the cloud service model allow for processing this encrypted data? Do the databases, if used, allow for the appropriate storage of the encrypted data? Does the cloud vendor contract allow for handling the encrypted data? If so, it may require that the cloud vendor retain copies of the encryption keys. Is it acceptable to the organization for the cloud vendor to have the encryption keys to your data? Doing so may make moot the company’s entire reason for encrypting the data in the first place. The IH team will need clearly defined processes regarding when the vendor may employ the use of the encryption keys if they are to retain copies of them. In addition, the organization will want the vendor to maintain and provide on request a journal of activities related to the use of these keys. If the cloud services are to provide the function and value they are being purchased for, can everything that is a concern to the organization be encrypted (e.g. configuration files, password files)? In the case of the client VM traversal or

Jeff Reed, jeff360@gmail.com

hypervisor attack, these encryption efforts may not be enough. This will depend on where and when the data is encrypted and decrypted. If it is not enough to encrypt the data, then what can the organization do? Is disk-based encryption an option in any cloud type? Probably not, but it may be worth investigating in some cases. Data remanence creates difficulties for an IH team in cloud scenarios that the customer should thoroughly analyze prior to deployment.

Cost-of-Service (COS) Attacks

Elasticity, as stated earlier, provides the ability to rapidly scale up or down as required with little vendor interaction. In the cases where this scaling is automated, it may allow for a certain level of DOS/DDOS mitigation. However, given the pay-per-use nature of scaling, this can also transform a mitigated DOS attack into a Cost of Service (COS) attack (aka eDOS, economic Denial of Service or Sustainability). The event that causes this can be incidental or intentional. The results are the same either way. How does or should the IH team investigate these? It will require the same tools used by and a close collaboration with the operations teams. This issue is somewhat ameliorated if there are measured service capabilities. Even this is not a silver bullet, however. This only allows a concerted attack to inflict both COS and DOS attacks in a single incident.

4.1.2. Clouds of Clouds

An important point to consider is that the cloud provider may be using services from yet another cloud provider. They, in turn, may be building on services from a third cloud provider whose platforms a fourth cloud vendor serves. As cloud adoption grows, this example of nested cloud services may prove to be the simple case.

Ensure that the organization has access to the full architecture map from the cloud vendor for their services offering. The map must include third party cloud vendor services upon which they are building their services. If they are using other cloud vendor's services for the processing they are providing, this possibly introduces more jurisdictions, legal liabilities, and a multiplication of the due diligence efforts on the customer's part. There will now be other support agreements to make or, at least, review. After all, the IH team's need to investigate an incident now follows a path through the additional cloud vendor's environments. A business is contracting with a cloud provider

Jeff Reed, jeff360@gmail.com

for services. That cloud provider is contracting with another vendor for cloud services. The business must ascertain whether the second contract meets the provisions required by the first contract. Does the business' contract with the cloud provider cover future cases (additional cloud services from other cloud vendors)? Does the customer know on what the second tier of cloud providers are building their services? And this continues ad infinitum. Remember, the discussion here is only about incident handling. Does the customer know where their data is? How can one be certain?

If the organization steps back to consider cloud depth, they will find that they must cover the ground discussed in this paper *again* to think through the scenario for just one additional layer of cloud provider. An analogy might be a picture of a picture within a picture. It may be more accurate to say that it is like pictures within pictures within pictures and so forth. How deep does the labyrinth of nested clouds go? Is there a reasonable threshold beyond which continued effort to understand the complete picture is too costly? This is clearly something that each organization must decide for itself. If regulation obligates the organization, these questions must be answered.

Even if not obligated by law regarding their data, the risks represented by the obscurity buried in the cloud layers is still important for the organization to understand. The motivation to rapidly make profits can often drive companies to make premature cost-of-service-only decisions regarding their providers. Setting aside the general security implications of poorly examined cloud depth, will the organization have any visibility into these other cloud layers for IH investigations? Jurisdiction issues again become apparent here, as do IH procedural issues.

The cloud provider may have provisions for narrowing the list of involved jurisdictions. Making a suggestion regarding tightening security control Brodtkin quotes Gartner:

When you use the cloud, you probably won't know exactly where your data is hosted. In fact, you might not even know what country it will be stored in. Ask providers if they will commit to storing and processing data in specific jurisdictions, and whether they will make a contractual commitment to obey local privacy requirements on behalf of their customers, Gartner advises. (quoted by Brodtkin, 2008)

Jeff Reed, jeff360@gmail.com

If the cloud provider does have the ability to narrow the impacted jurisdictions, it is critical to know if they do so in a way that still allows for the full benefits envisioned for this service. In other words, to know if the act of narrowing the jurisdictions involved prevent the scaling, distributed resources, or other elasticity mechanisms that make cloud services attractive in the first place. In addition, if they can narrow the impacted jurisdictions while retaining these benefits, can they guarantee it throughout the entire cloud supply chain? The consequences of not understanding which jurisdictions are affected can be significant (also see Section

4.3.	Legal
------	-------

).

4.2. Operations

4.2.1. Backups, Disaster Recovery (DR) and Business Continuity (BC)

An IH team's processes will need to include conditions under which an incident invokes DR and BC processes and account for them in the contract. NIST says that, "Organizations should also ensure that incident response policies and procedures and business continuity processes are in sync" (Scarfone, *et al*, 2008). If the IH team requires access to the cloud vendor's BC and DR systems or facilities, the contract will need to provide for it. Legal will need to research whether the jurisdiction of the BC and DR sites cause issues for the organization. If the business anticipates that IH investigations will require access to BC and DR systems and sites, ensure that the contract includes coverage for capabilities, processes, SLAs, contacts, and costs.

Does the cloud provider have data retention policies that would remove or retain data in conflict with the organization's policies? Data retention requirements can vary widely. In cases such as instant messaging (IM), a company may require that no data is retained when the IM application exits. In other cases, data may have to be stored for years due to regulatory or other reasons. An organization will need to ensure that the cloud vendor can accommodate these retention requirements.

A more subtle BC and DR concern is that of the cloud vendor going out of business. For IaaS and PaaS (Platform as a Service), beyond their BC site considerations, a company will need to have plans for re-establishing their environment with either another cloud vendor or in-house. The SaaS case may prove more difficult. If a business integrates a SaaS service as a component of an enterprise application and the cloud

Jeff Reed, jeff360@gmail.com

vendor closes their doors, the business must recreate the integrated service. Can the business recreate the SaaS service or is there a compatible service provided by another SaaS provider? In either case, can the company restore the service quickly enough to avoid unacceptable losses? Considerable analysis may be necessary to understand all of the consequences resulting from the particular vendor closing their doors. There will most certainly be more to consider than simply having a copy of the affected data. The loss of specific business processing in a critical business application can be catastrophic if not rapidly restored.

Cloud integration complicates BC and DR for a customer. As in almost every other aspect of cloud integration there are many facets to consider if the business is to ensure resilience and the proper function of its IH processes.

4.2.2. Staff – Cloud Motivations & Impacts

A strong motivator for adopting cloud is the ability to eliminate IT staff and overhead. The cost savings and efficiencies gained in this aspect of cloud computing are compelling. However, an issue arising from this is the loss of operational and effective decision-making skills in the organization. This has been experienced in many IT outsourcing and off-shoring scenarios. If an organization has structured their IH capability to utilize skilled professionals from within the IT operations teams, this raises another set of concerns. Can the organization regain the lost visibility that the outsourced roles and their associated skill-sets provided to the IH team? Even if the organization is able to retain the team members by reassigning them to the IH team at the time of outsourcing, how can the organization maintain those skill sets without the operational environment that created them? One alternative is to plan for the education and the environment necessary to maintain the required skills. If this is possible, the business will need to factor the training as well as the setup and ongoing maintenance of the training environment into the cloud integration ROI. Nevertheless, can IH management afford to provide an environment (“sandbox”) for them and, if not can the organization maintain the motivation of a highly skilled team without an environment for practicing and sharpening those skills? The cost of not doing so will be an ever-diminishing team experience base with which to engage incidents. The effects of atrophying skill-sets have

Jeff Reed, jeff360@gmail.com

also been a recurring theme in books, television and movies for years. A sophisticated people advance to a level at which all their needs are comfortably met and which no longer requires of them certain activities or rudimentary survival skills. The people achieve their level of comfort by some fantastic technological advancement. Over a process of many years, the fundamental skills are lost to this people. The plot usually goes bad when the amazing technology fails due to some unforeseen catastrophe or condition encountered in the future leaving the progeny of this people facing annihilation if they cannot adapt. Creative script writing may solve the problem for these folks, but not for the IH team in the cloud environment.

Again, balance is required along with a reasonable level of research and analysis. Too much caution will prevent the necessary advancement for an organization to remain competitive. Too little caution can result in poor decision making with equal risk.

4.3. Legal

This paper deals with legal aspects at a cursory level. To get this right, the IH team will need to engage the legal team and, as necessary, educate them on the various technologies, business context, and other aspects of the cloud integration. The IH team should work with them to do further topical research as areas of concern are identified. While the IH team will understand technical nuances that legal will not, legal can guide the IH team regarding the legal implications of the integration cases.

As has been noted repeatedly, it is critical to their legal processes for the organization to know where its data is being stored at *all* points in time. Cloud computing encompasses global locations for data, processing and other shared resources. The organization must know where their data *could* be stored at *any* point in time so that the IH team can work with the legal team regarding the jurisdictions involved. Legal will need to review the laws in all applicable jurisdictions on a cyclical basis. Ask your cloud provider if they have processes or services to help their customers with this.

Many countries require a business to report to or provide archive for law enforcement in some serious crime scenarios including intended murder and penetration of government systems (West-Brown, *et al*, 2004). If the organization's processes cannot detect content such as that described by West-Brown, *et al*, or cannot afford the level of

Jeff Reed, jeff360@gmail.com

effort required to comply legally, then the organization has a dilemma. The IH team in particular will require very detailed access to private data to accomplish intent detection such as this. For each jurisdiction, legal counsel will need to direct the IH team regarding what the IH team can do legally in the course of the investigation. Here are some of the questions legal counsel should provide answers for about each jurisdiction:

- Can the organization view and retain information on individuals under any circumstances?
- Can the organization view and retain information on individuals who are children? (In some jurisdictions, different laws apply to children's data.)
- How does the IH team act if the data under investigation is stored in one country, is processed in another and has its point of access in a third (Scarfone, *et al*, 2008)?
- What must the IH team do if summoned to court?

To understand what an organization is facing when adopting cloud services they must multiply these questions by the number of cloud integrations they have or are considering and the number of cloud layers in each integration (refer to Section 4.1.2 Clouds of Clouds). This could leave the organization in an untenable position. Do the laws of one jurisdiction require the organization to take an action and the laws in another prohibit that action? The more legal jurisdictions that are crossed, the more likely that this type of situation will arise. The difficulties in understanding all of the legal issues multiplies the further one progresses in adopting cloud services.

There are costs associated with the legal processes required in cloud integration. Have the costs of legal research, analysis and accounting related to the cloud integration concerns expressed in this section been fully considered? The organization may have enough information to calculate their own costs for their current architecture. However, does the business understand enough about the proposed new cloud integration to forecast the legal research and analysis costs going forward? An organization should ask the cloud vendor to provide the information necessary to assist with these calculations. Find out what support the vendor provides as their operations change (e.g. new sites, new jurisdictions, new services, and so forth). If necessary, add support to the contract for

Jeff Reed, jeff360@gmail.com

obtaining the data and support mentioned here. A business needs to consider the costs for obtaining this data and support before making their final decisions about the integration.

Frequently, the IH team has very little time during the cloud integration process to render answers, guidance and other feedback to the procurement department. The IH team is often not even invited to the strategic architecture or solution engineering table. Forward thinking and planning, then, are essential. The IH team needs to be actively developing relationships in order to obtain early information regarding strategic direction or thinking. This will enable the team to work through these cases before they get a call from procurement, or worse, from the help desk.

There is also the issue of reporting information regarding victims. Does the organization know their responsibilities for all jurisdictions that apply? Again West-Brown, *et al*, specify the following:

Liability exposure here depends on who is requesting the information. You may be liable if you reveal the identity (without prior consent) of victim sites to other victims, law enforcement, or the media. But you may not be liable if you are required to report the same information to an internal audit. (West-Brown, *et al*, 2004)

Further complicating the issue of reporting, the cloud provider may additionally be required to report the information to the customer. The organization needs to know what it is willing to accept here. Brodtkin warns:

Customers are ultimately responsible for the security and integrity of their own data, even when it is held by a service provider. Traditional service providers are subjected to external audits and security certifications. Cloud computing providers who refuse to undergo this scrutiny are "signaling that customers can only use them for the most trivial functions," according to Gartner. (quoted by Brodtkin, 2008)

Each business must analyze the cloud characteristics introduced by their integration to create documented legal requirements for the organization and for the vendor contract.

4.4. Scaling

This paper has repeatedly mentioned the issues surrounding integrating with more than one cloud provider. The topics have been discussed in light of specific IH documentation, planning and operational concerns. The concerns apply to other third party integrations as well. However, it is also important to address the issue of organizational scaling as it relates to cloud integration.

Can the organization sustain the processes and decision matrices that multiple cloud providers introduce to the organization? They can probably do so to some extent. The point is that, each new cloud integration introduces additional organizational overhead and complexity. There is also increased distraction for the team. Can the organization afford that distraction in the middle of a critical investigation? The effect can be multiplicative. Each business will need to know their threshold for cloud scaling, as well as with the associated costs and risks.

5. Conclusion

Cloud computing, in its various forms, offers considerable benefits to industry. It does so by providing very complex, scalable computing infrastructures on which the organization can build its enterprise architecture. The organization needs to understand and account for the characteristics of these offerings in their IH policies, processes, personnel and cloud services contracts. Scalability and cloud depth present the IH and legal teams with serious challenges. As core cloud capabilities, an early analysis of scalability and cloud depth by the organization will enable them to make critical and timely decisions when the enterprise commits itself to cloud integration.

It cannot be over emphasized that, with cloud integration, there is no ‘one size fits all’. If an organization makes the mistake of adopting this mindset, the results can be devastating. One SaaS integration will not be the same as another SaaS integration. Neither will the IH concerns be the same for a PaaS integration as for an IaaS integration. An organization must thoroughly examine each cloud integration in its own context.

The organization will want to undertake a systematic approach to analyzing their IH capabilities and concerns in light of each new cloud integration. The IH team should consult other stakeholders throughout this process to ensure a sufficiently broad perspective of the issues, and to identify opportunities for collaboration and consolidation of tasks. By either utilizing the framework with which the organization started their IH capability or by adopting another well-established framework as a guide the organization can ensure that it is addressing all critical areas.

Jeff Reed, jeff360@gmail.com

6. References

- Brodkin, J. (2008). Assessing the security risks of cloud computing. *Network World*, 12(3). Retrieved September 17, 2010 from <http://www.infoworld.com/d/security-central/gartner-seven-cloud-computing-security-risks-853>
- West-Brown, M.J., Stikvoort, D., Kossakowski, K-P., Killcrece, G., Ruefle, R., Zajicek, M. (2004). *The Handbook for Computer Security Incident Response Teams, Second Edition*. Retrieved September 18, 2010, from <http://www.cert.org/archive/pdf/csirt-handbook.pdf>
- Scarfone, K., Grance, T., Masone, K. (2008). *NIST Computer Security Incident Handling Guide*. Retrieved September 18, 2010, from <http://csrc.nist.gov/publications/nistpubs/800-61-rev1/SP800-61rev1.pdf>
- Mell, P. & Grance, T. (2009). *The NIST definition of cloud computing (v15)*. Paper presented at the Cybersecurity and Standards Acceleration to Jumpstart Adoption of Cloud Computing (SAJACC) session of the NIST Cloud Computing Forum & Workshop, May 20, 2010. Retrieved October 9, 2010 from <http://csrc.nist.gov/groups/SNS/cloud-computing/>
- UK Centre for the Protection of National Infrastructure. (2010, March) *Security briefing 01/2010, cloud computing*. Retrieved September 18, 2010, from <http://www.cpni.gov.uk/Docs/cloud-computing-briefing.pdf>
- European Network and Information Security Agency (2009, November) *Cloud computing, benefits, risks and recommendations for information security*. (D. Catteddu & G. Hogben, Eds.) Retrieved September 18, 2010, from http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport
- Youseff, L., Butrico, M., & Da Silva, D. (2008). *Toward a Unified Ontology of Cloud Computing*. Retrieved September 18, 2010, from <http://freedomhui.com/wp-content/uploads/2010/03/CloudOntology.pdf>
- Chambers, J. (Speaker). (2009, April 20). *Collaborate with Confidence: Securely connect, Communicate, Conduct Business in Decentralized / Highly Collaborative Environment* (Online only recording of a speech presented at the RSA

Jeff Reed, jeff360@gmail.com

- Conference 2009, San Francisco, CA). Retrieved November 20, 2010, from http://media.omegiaweb.com/rsa2009/webcast_exclusive.htm?id=2_3
- Greene, T. (2010, March 8). Cloud security, cyberwar dominate RSA Conference. Network World.
- Owens, D. (2010, June). Securing Elasticity in the Cloud. *Communications of the ACM*, 53(6), 46-51. doi:10.1145/1743546.1743565
- Blanton, S., & Schiller, C. (2010). *Is There Safety in the Cloud?* Retrieved September 18, 2010, from <http://www.educause.edu/EDUCAUSE+Quarterly/EDUCAUSEQuarterlyMagazineVolum/IsThereSafetyintheCloud/206543>
- Anderson, K. (2005). *Intelligence-based Threat Assessments for Information Networks and Infrastructures*. Retrieved November 20, 2010, from http://www.aracnet.com/~kea/Papers/threat_white_paper.shtml
- Tipton, H. (2007). *Official (ISC)2 guide to the CISSP CBK*. (K. Henry, Eds.) [Books24x7 version] Available from http://common.books24x7.com/book/id_30425/book.asp
- Tipton, H. & Krause, M. (2007). *Information Security Management Handbook, Sixth Edition*. [Books24x7 version] Available from http://common.books24x7.com/book/id_26438/book.asp
- Houghton Mifflin Company (2000). *The American Heritage Dictionary for the English Language, Fourth Edition*. Boston, Massachusetts: Author.

7. Appendices

7.1. NIST Cloud Definition v15

The NIST Definition of Cloud Computing

Authors: Peter Mell and Tim Grance

Version 15, 10-7-09

National Institute of Standards and Technology, Information Technology Laboratory

Note 1: Cloud computing is still an evolving paradigm. Its definitions, use cases, underlying technologies, issues, risks, and benefits will be refined in a spirited debate by the public and private sectors. These definitions, attributes, and characteristics will evolve and change over time.

Note 2: The cloud computing industry represents a large ecosystem of many models, vendors, and market niches. This definition attempts to encompass all of the various cloud approaches.

Definition of Cloud Computing:

Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.

Essential Characteristics:

On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service's provider.

Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, network bandwidth, and virtual machines.

Jeff Reed, jeff360@gmail.com

Rapid elasticity. Capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured Service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported providing transparency for both the provider and consumer of the utilized service.

Service Models:

Cloud Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Cloud Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Cloud Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models:

Private cloud. The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.

Community cloud. The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be

managed by the organizations or a third party and may exist on premise or off premise.

Public cloud. The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

Note: Cloud software takes full advantage of the cloud paradigm by being service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability.



Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS San Diego 2017	San Diego, CAUS	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Seattle 2017	Seattle, WAUS	Oct 30, 2017 - Nov 04, 2017	Live Event
SANS Gulf Region 2017	Dubai, AE	Nov 04, 2017 - Nov 16, 2017	Live Event
SANS Milan November 2017	Milan, IT	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Amsterdam 2017	Amsterdam, NL	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Miami 2017	Miami, FLUS	Nov 06, 2017 - Nov 11, 2017	Live Event
SANS Paris November 2017	Paris, FR	Nov 13, 2017 - Nov 18, 2017	Live Event
Pen Test Hackfest Summit & Training 2017	Bethesda, MDUS	Nov 13, 2017 - Nov 20, 2017	Live Event
SANS Sydney 2017	Sydney, AU	Nov 13, 2017 - Nov 25, 2017	Live Event
GridEx IV 2017	Online,	Nov 15, 2017 - Nov 16, 2017	Live Event
SANS San Francisco Winter 2017	San Francisco, CAUS	Nov 27, 2017 - Dec 02, 2017	Live Event
SANS London November 2017	London, GB	Nov 27, 2017 - Dec 02, 2017	Live Event
SIEM & Tactical Analytics Summit & Training	Scottsdale, AZUS	Nov 28, 2017 - Dec 05, 2017	Live Event
SANS Khobar 2017	Khobar, SA	Dec 02, 2017 - Dec 07, 2017	Live Event
SANS Austin Winter 2017	Austin, TXUS	Dec 04, 2017 - Dec 09, 2017	Live Event
SANS Munich December 2017	Munich, DE	Dec 04, 2017 - Dec 09, 2017	Live Event
European Security Awareness Summit & Training 2017	London, GB	Dec 04, 2017 - Dec 07, 2017	Live Event
SANS Bangalore 2017	Bangalore, IN	Dec 11, 2017 - Dec 16, 2017	Live Event
SANS Frankfurt 2017	Frankfurt, DE	Dec 11, 2017 - Dec 16, 2017	Live Event
SANS Cyber Defense Initiative 2017	Washington, DCUS	Dec 12, 2017 - Dec 19, 2017	Live Event
SANS Security East 2018	New Orleans, LAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS SEC460: Enterprise Threat Beta	San Diego, CAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, NL	Jan 15, 2018 - Jan 20, 2018	Live Event
Northern VA Winter - Reston 2018	Reston, VAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SEC599: Defeat Advanced Adversaries	San Francisco, CAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Berlin 2017	OnlineDE	Oct 23, 2017 - Oct 28, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced