

# UC Davis

## UC Davis Previously Published Works

### Title

Are Your Papers in Order? Developing and Enforcing Multi-tenancy and Migration Policies in the Cloud

### Permalink

<https://escholarship.org/uc/item/0dp8f206>

### ISBN

978-1-4577-1925-7

### Authors

Hay, Brian  
Nance, Kara  
Bishop, Matt  
et al.

### Publication Date

2012

Peer reviewed

# Are Your Papers in Order?

## Developing and Enforcing Multi-Tenancy and Migration Policies in the Cloud

Brian Hay  
University of Alaska  
Fairbanks  
brian.hay@alaska.edu

Kara Nance  
University of Alaska  
Fairbanks  
klnance@alaska.edu

Matt Bishop  
University of California  
Davis  
bishop@cs.ucdavis.edu

Lucas McDaniel  
University of Alaska  
Fairbanks  
lamcdaniel@alaska.edu

### Abstract

*As cloud usage continues to increase, new issues with respect to managing and securing resources in the cloud are becoming more apparent. While some people may believe that security and privacy in the cloud can be addressed without the consumer considering the physical location and internal structure of the cloud, we show that this is clearly not the case. Furthermore, we describe a mechanism by which cloud consumers can inform cloud providers of their requirements in a manner that still allows the cloud to remain dynamic and flexible. Specifically, this paper explores the analogy between human migration in the real world and virtual machine migration in an IaaS cloud environment. It addresses issues such as jurisdictional control, zone evolution, migration, and instantiation based on an examination of these analogous real-world scenarios and their applicability to the cloud.*

## 1. Introduction

Cloud computing is certainly a topic of significant current debate and research in academic and commercial environments. It is becoming an increasingly practical tool for a variety of tasks from scientific research to corporate data storage, processing, and management. While the form and definition of the cloud continue to evolve, the Infrastructure as a Service (IaaS) model of the cloud presents new challenges for users and service providers that will require changes in current operational models in order to address some of the unique issues that this dynamic environment offers.

As described in [4], the implementation model for most IaaS environments differs from the usual corporate computing model in several ways. The NIST definition of IaaS cloud computing describes five core characteristics [8]. Of those, the resource pooling, on-demand self-service, and rapid elasticity are particularly relevant to this discussion.

### 1.1 Resource Pooling

In the past, computing resources such as storage and hosts were generally used exclusively by a single resource owner. However, in the cloud model, resources allocated to one user are commonly instantiated on some physical infrastructure that also hosts resources allocated to other

users. A common way to accomplish this is to allocate resources as virtual machines (VMs), which are software containers that emulate the physical hardware of a computer [9]. For example, a virtual machine may be instantiated on a physical server hosting several other virtual machines, each allocated to a different corporation. In addition, a resource assigned to cloud consumer *A* may be reallocated to consumer *B* at a later time.

### 1.2 On-Demand Self Service

Cloud consumers must be able to “provision computing capabilities ... automatically ... without requiring human interaction with a service provider”. [8]

### 1.3 Rapid Elasticity

This characteristic addresses the need to scale resources dynamically to meet demand in real time, or at least near-real time.

These three important characteristics make clear that a given IaaS cloud is likely to be a dynamic environment in which the resources allocated to consumers, and indeed the state of the cloud as a whole, is in a constant state of flux. Consumers can manually or automatically deploy, redeploy, and release resources at will while being abstracted from the underlying implementation details.

In addition to consumer-driven actions, the cloud providers are also likely to influence the state of the cloud in an attempt to meet their service level agreements (SLAs) while utilizing their resources efficiently. One common technique utilized by cloud providers is virtual machine migration. Migration is an important concept in modern virtualization platforms, particularly as cloud environments become increasingly popular. The ability to move VMs between physical hosts is vital for load-balancing and to facilitate high availability operations. Other considerations may make such migration appealing, including:

1. The ability to more efficiently use physical hardware to meet consumer demand. For example, consider the simple case of a cloud provider that offers consumers the option to deploy small (25% of a physical server) or large (50% of a physical

server) VMs. If they have three physical servers with 3 small VMs each, they cannot deploy a large VM despite there being 75% of a physical server available in the resource pool. Migration of the small VMs would allow the same physical resources to be meet the demand.

2. The ability to migrate VMs to alternate geographic locations (e.g., for disaster preparedness or recovery, to provide better network performance as user demographics change, or even to place the VM in a legal jurisdiction more compatible with the activities being performed).
3. The ability to organize and isolate VMs to prevent conflicts, such as when VMs assigned to two competing corporations are assigned to the same physical host.

The result is a cloud environment in which consumers are very likely to have assigned virtual resources that are co-located (i.e., instantiated on the same physical resources) with a variety of other unknown consumers, and that their virtual resources may be migrated between physical resources over their deployed lifetimes. While this migration and multi-tenancy are certainly not inherently bad, it is important to consider the potential issues that such a dynamic environment may raise.

Eran Feigenbaum, the Chief Security Officer for Google Apps, recently claimed that “[p]rofessionals should worry about security and privacy of data, rather than where it is stored” [7]. This is the idea behind cloud computing: make the computations and storage completely transparent to the user. As we will now show, this is a stunningly, although probably not surprisingly, naïve point of view with respect to security and privacy considerations, because those considerations clearly and inherently depend upon the physical characteristics, and specific physical locations, of the cloud and its components.

## 2. Security Implications

A simple security scenario is that of a single server which provides some service(s). In this environment the administrator must consider the security of the operating system, each deployed service, and the system software. This is a well-contained problem, and essentially all aspects of the environment are under the administrator’s control.

A slightly more complex scenario is that of a server dependent on some external trusted system, such as an authentication server. This problem is also well-contained, in that the administrator has control over everything other than the external host and the connection to it, and will presumably carefully choose the external trusted components (e.g., the authentication services).

More complex scenarios are in common use today (e.g., multi-tier architectures ranging from front-end load

balancers to backend databases), but the situation in the cloud is fundamentally different in that it introduces two new categories of potential vulnerability, both of which are very much out of the control of the resource administrator (who, in this context, is the cloud consumer).

The first such category is the cloud provider itself. In order to allow rapid deployment and configuration, cloud providers typically implement web-based management platforms that allow consumers to manage their resource allocation levels (and quite commonly the configuration of those resources). In addition the cloud providers have their own tools to manage their resource pools. These platforms are as vulnerable to attack as any other software platform, either due to vulnerabilities in the software itself, or by social engineering approaches such as the spear phishing attacks that have recently proved so effective against many large corporations such as Google (Gmail), RSA, and others. While the implementation of these platforms is beyond the control of the cloud consumer, they can at least select providers based on their observations of the available providers’ business practices and interface characteristics. For example, if provider *A* provides HTTPS login form access to their management interface from an HTTP page, and provider *B* strictly limits management interface access to VPN sessions authenticated using strong 2-factor authentication, the consumer may opt for provider *B* on that basis. Of course this does not provide a complete picture, but does at least provide a cloud consumer with some information on which they can base their choice of provider.

The second category is the co-location of resources with resources assigned to other consumers whose identities, and security postures, are unknown. This issue is likely to worry a security-minded administrator. While there are several virtualization platforms on the market today, all of them have a trusted virtualization layer that, if compromised, leads directly to full compromise of any of the virtual machines running on the physical host. This includes the ability to monitor all activity on the virtual machine, and almost certainly to alter the state of the virtual machine in an arbitrary manner (for example, to execute arbitrary code in the context of the virtual machine). Virtualization layers are complex software systems, and it is unreasonable to believe that they have no vulnerabilities that would allow a virtual machine user to gain control of the virtualization layer, and from there to gain control of all other virtual machines running on the same physical host. In this environment, the security of a virtual machine depends not only on the controls put in place on that VM, but also the controls put in place by the administrators of all other virtual machines running on that physical host. For example, consider the case where Alice runs well secured server on the same physical host as Bob, who has not updated his VM since 2005 and runs many services with default configurations. If Eve has access to a so-called VM breakout exploit and an exploit for one of

Bob's services, Eve can compromise Alice without there being a traditional vulnerability on Alice's server.

Since Eve and Bob could be the same person, Alice could be compromised without the need for an exploitable VM to be co-located. Eve could deploy a VM, and then perform the VM breakout exploit on her own VM. Deliberate co-location of VMs is a subject that has received some preliminary study [3], and as likely to receive more attention (as is the issue of VM breakout attacks).

Even without full VM breakout capability, co-location may provide useful information to an attacker. For example, Eve could attempt to use side-channel attacks involving cache latency, CPU utilization, and other such channels on a co-located web server to attempt to acquire private keys for SSL certificates.

### 3. Legal and Regulatory Implications

Beyond the potential security issues brought about by co-location, several significant multi-jurisdictional issues arise as VMs are deployed in or migrate among physical infrastructures. Resources in the cloud are often not fixed to any geographical location. The VMs may migrate between physical locations during their useful lives.

Suppose that one jurisdiction imposes severe penalties for exposure of unencrypted Personally Identifiable Information (PII), whereas another jurisdiction does not. Instantiation of cloud resources in these two jurisdictions exposes a cloud consumer to different (degrees of) liabilities. Even when one jurisdiction is selected at VM instantiation, the cloud provider may migrate that VM to the other jurisdiction, or instantiate clones of the original VM automatically in the second jurisdiction in response to elevated demand or as part of a disaster recovery plan.

A solution to this simple example is, of course, to encrypt all PII regardless of which jurisdiction the VM will be deployed in. However, that type of approach is not possible in the case of conflicting requirements. For example, in the United States there are very significant limitations on the ability to restrict speech (no matter how repugnant it may seem), whereas in other nations, certain types of speech (such as pro-Nazi propaganda) are limited. As such a service providing some public service may be limited in the ability to filter content when deployed in one physical location (the United State, for example), whereas the same service instantiated elsewhere in the cloud (France, for example), would be required to filter certain types of content. In such cases, there is no "most restrictive case", so it may not be possible for a cloud consumer to build a VM that meets the requirements of all jurisdictions.

### 4. Analogy to Real Life

Given these concerns, it is important to provide a mechanism to enable cloud consumers to describe deployment and migration practices for their resources, and label them as acceptable or unacceptable. Example use cases include:

- Company *A* does not want to have its resources physically co-located with competing companies.
- Company *B* wants to ensure that its restricted access systems (e.g., databases) are not deployed on physical resources that also host Internet facing systems such as web servers, whether those servers are owned by company *B* or some other entity.
- Company *C* is willing to deploy hosts in legal jurisdictions that do not have strong consumer data access rights.
- Company *D* is willing to co-locate their resources with other corporate consumers that have met some standard for security and identity, but is unwilling to co-locate with consumers that have not met that standard.

Today, some cloud consumers can deploy a private cloud that will address some of these issues, at the expense of the massive flexibility and scalability offered by large public clouds. In the public cloud, consumers have very little ability to specify any deployment or migration requirements beyond very basic information such as the initial deployment location (EC2, for example, allows the consumer to select the country of deployment).

### 5. Analogous Scenarios in the Physical World

This situation, while technological in foundation and motivation, has counterparts in other domains for which procedures exist. The experience from those domains can provide insight to this technological problem. For example, the movement of VMs through different deployment zones is analogous to the movement of humans through different jurisdictions. Governments and other entities have invested a significant amount of time, effort, thought, and money to address this issue for people moving through jurisdictions. Among the similarities in the two scenarios are the following:

- Rules within (and perhaps unique to) each jurisdiction.
- Evolution of jurisdictional boundaries and rules.
- Characteristics of the individual (VM or human).
- Migrations.
- Participation by zones and individuals.

Instead of starting anew, let us consider how applying the lessons learned from this human travel can inform controls on deployment and migration in the cloud. We examine both successful and unsuccessful approaches to the problem of human migration. The following sections discuss the above issues and provide an example of how a system might be developed to being to approach this challenging problem.

## 6. Jurisdictional Control

Just as humans are born in some jurisdiction, but may migrate (move) to other jurisdictions through their lifetimes, VMs are instantiated in an area that is subject to a set of regulations, but may migrate to other zones. Analogies between the movement of VMs through jurisdictions in the cloud and the movement of humans through jurisdictions abound. The jurisdiction in which people are physically located defines what is legal for them to do. While not always strictly enforced, many border control offices have policies that put the burden of compliance on the visitor. For example, in the U.S., driving between states can change the rules and regulations with which the driver must comply, including the use of cell phones and radar detectors while driving, seatbelt requirements, adjustments in speed limits, and many other less obvious laws. Just as citizens crossing state lines or traveling abroad must be aware of and comply with the laws of the state or country they are visiting, VMs are restricted to actions permissible in the jurisdiction in which they are physically and logically located—and breaking those laws have associated consequences.

One way to handle the multi-jurisdictional issues is to establish a system to describe similarities across different jurisdictions and VMs. The cloud consists of servers, data centers, and other resources in locations we call *zones*, each with a potentially unique structure. By associating each zone with sets of migration characteristics that describe what the zone can provide and whose VMs they will accept,<sup>1</sup> similarities across the cloud can be identified. Similarly VMs with their own set of migration characteristics would state the requirements that the zone running the VM must have. By restricting the migration of VMs to those zones that meet the requirements, and *vice versa*, a more fluid transfer of VMs through the cloud can be achieved while complying with jurisdictional regulations.

Including legal requirements in this set of characteristics can prevent many potential problems. Restricting migration of VMs to zones with appropriate characteristics

---

<sup>1</sup> Here, “whose” means both an entity (for example, the U.S. government’s VMs), and a description of the characteristics of VMs that the provider will accept (such as VMs that encrypt all PII).

can allow for their movement across jurisdictions without worrying about violations or multi-jurisdictional issues. Better utilization of cloud resources can be achieved by migrating VMs to compliant zones that can meet their resources demands.

However in order to provide these capabilities, cloud consumers lose some of their privacy as information regarding the VM has to be revealed. They may have to describe certain characteristics of the VM (e.g., has encrypted PII) for the cloud provider to accept the VM; they may also find it advantageous (although optional) to describe other attributes. Although sacrificing some privacy by providing this optional information can be very beneficial (e.g., it may allow those VMs to be migrated to zones that are more suited to the type of computation they are performing), it might also be undesirable for corporations that wish to preserve the privacy of that information. This system provides many incentives for disclosing information as the cloud provider can use this knowledge to mitigate jurisdictional issues while increasing performance.

Just as humans struggle with the tradeoff between security and privacy, and are willing to relax aspects of privacy if properly incentivized, entities operating in the cloud face similar issues and may be willing to provide information (relax some privacy aspects) in exchange for incentives to comply with cloud provider requirements (or consequences if they do not). The options to encourage participation can involve positive incentives (carrots) and negative incentives (sticks). While beyond the scope of this paper, incentive models in the cloud could include high priorities associated with more comprehensive identity profiles, or even denial of service if compliance is deemed mandatory.


## 7. Zone Evolution

VMs may be moved across the cloud when the zone they are located in is modified. Among the modifications that may occur are:

- Physically moving the zone to a different jurisdiction
- Changes in laws governing that jurisdiction in which the zone is located.
- The addition or removal of physical equipment within the zone.
- The addition or removal of resources (e.g., a VM) within the zone.
- Modification of migration characteristics of the zone.

Modifications of the second type may be beyond the control of the IaaS service provider, as nation-state boundaries, regulations, and laws demonstrate their

ongoing fluidity. For example, new privacy laws in the EU have implications for any VMs running in datacenters physically located there. The service provider could also initiate the zone modifications. For example, we may want to redefine a TOP SECRET zone as SECRET, but the result is that any TOP SECRET VMs in the zone would no longer be in compliance, so before making the change to zone characteristics, we must identify non-compliant VMs and migrate them, suspend them, terminate them, or prevent the change to the zone characteristics. We may *have* to change these characteristics as a result of changes to the legal environment in which the zone resides (for example when an additional privacy-preserving requirement is created, or the relevant government requires that any cryptographic communication channels provide a “backdoor” for law enforcement monitoring). Again VMs not in compliance with the new characteristics of the zone need to be identified and dealt with (either by migration to other zones, by changing the characteristics of the VMs to meet the zone requirements, by terminating the VM, or by detaining the VM in a suspended state until a suitable zone can be found).



**Figure 1 - Potential Zone Evolution Ramifications**

Figure 1 depicts the result of a change to the characteristics of Zone A, which contains four VMs. VM 1 is found to be in compliance with the new Zone A characteristics, and remains in the zone. VM 2 and VM 4 are found to be out of compliance with the new Zone A characteristics and are forced to migrate to other compatible zones (forced migration). VM 3 is also found to be out of compliance with the new characteristics of Zone A, but a suitable alternate zone cannot be found, so the execution of VM 3 is suspended within Zone A. In effect, VM 3 becomes a detainee until the conflict can be resolved in a manner defined by the cloud provider. This resolution can take many forms, including locating an acceptable alternative zone, or adjusting the requirements of the VM to make it compliant with its existing zone.


New servers and resources may be added to zones to meet the ever-changing needs of the cloud. Knowledge of the characteristics and requirements of the VMs in the cloud allow for better utilization of cloud resources given the dynamic environment. This knowledge can also

indicate when a VM in a zone would be better suited elsewhere or when the VM is no longer in compliance with its current zone. In order to identify these VMs, their migration characteristics need to be known so that they can be compared to the new characteristics of the current zone, or to a new zone to which they will migrate. Once the emigration and immigration requirements have been met, these VMs will be able to migrate to a new zone in a manner similar to that of human migration.

## 8. Migration

During the creation of the Department of Homeland Security, the United States Border Patrol, along with the INS inspection division, the U.S. Customs inspection division, and the Department of Agriculture’s plant and animal inspection service, were merged into a new agency called U.S. Customs and Border Protection. It is one of the DHS's largest and most complex components with varied responsibilities, such as keeping terrorists out of the U.S., barring entry to terrorist weapons, securing and facilitating trade and travel, and enforcing hundreds of regulations [1, 2]. Its wide range of responsibilities has also made it one of the largest organizations in the government. Migration is one of its major concerns.

As with human migration, migration in the cloud may have a number of different motivational factors, not all of which have been identified at this time. Also analogous to human migration, the migration process in the cloud involves at least two zones, one from which the VM is departing, and one to which the VM is entering. Zone stakeholders have the dual goals of ensuring that both entering and departing VMs comply with restrictions on their movement.



**Figure 2 - Migration Path Between Two Zones**

As shown in figure 2, an Emigration Agent guards the exit from a zone. The agent is responsible for determining if an exiting VM meets the requirements for leaving a zone. For example, a VM with SECRET data can only leave the current zone if the destination zone can accept SECRET data. The Immigration Agent guards entry to a zone. That

agent is responsible for determining if a VM meets the requirements for entering a zone. For example, if the zone requires that a new VM may not be owned by a competitor of any VM already within the zone, the Immigration Agent must verify that the new VM meets this requirement.

From an immigration perspective, Immigration Agents must be able to ensure that VMs moving into a new zone meet the migration requirements of that zone. For example, a zone may specify that all VMs operating in the zone contain no classified data. Zones could be restricted to a particular set of corporate users (or alternatively to *not* belonging to a particular corporate user to prevent conflicts). Zones could restrict access to avoid situations they do not want to deal with (such as processing medical data in legal jurisdictions that do not protect PII). It would be reasonable to expect that any VM moving into that zone meet those criteria. Models could be developed to guide VM migration that mirror world travel such as reciprocal agreements between zones, zones requiring visas, trade agreements, etc.

Controlling VM immigration requires a technological equivalent of a Border Patrol with a formal process for crossing borders. Referring back to incentives, VMs that do not provide sufficiently detailed information to satisfy the Tech Border Patrol would be denied access to the zone. This brings up potential privacy issues, and there are significant research efforts addressing this issue for the human population that could likely be adapted to the cloud. VMs could store tags that reflect data attributes. Access to the individual tag items could be restricted to only those required to complete the migration transaction to maximize process privacy.

As with human zones, it is anticipated that some migration policies will be more strictly enforced in some zones than in others. Some zones will likely have rules that are not enforced and will thus attract a certain “class” of VMs as a result. This is analogous to the human world, where border patrol can be bribed to “look the other way” in exchange for some form of remuneration.

And again, as with the human system, zones must have the capability to deal with special circumstances. Such situations include the following:

- **Detainees:** VMs that have been denied visas or have other tags (e.g., a VM is no longer compliant with the current zone, and no new zone can be found)
- **Refugees:** Circumstances may arise where an environment is no longer available and something must be done with its VMs.
- **Forced Migration:** In a variety of situations, and for a variety of reasons, a VM will be forced to migrate against the wishes of its stakeholders. For example, a conflict of interest may be identified between owner A and owner B such that, if owner

A remains in the zone, all VMs for owner B must be migrated out.

- **Deportation:** In some situations, systems may be discovered to be operating illegally in an environment. For example, in a zone limited to hosting backend systems such as databases, a VM may be reconfigured after deployment to include an Internet accessible webserver. This would violate the rules for the zone, and the offending VM would then need to be deported to a zone more suited to its new role.

While not a focus of this paper, an important foundational component of any successful migration system, whether human or VM, is the careful consideration of identity attributes and privacy. Nissinbaum observes that “privacy has been the rallying cry against ... computer-based, digital electronic technologies that have hugely magnified the power of human beings over information” [5]. Just as humans have been questioning the lack of informational self-determination in many countries, organizations will make the same rallying cries about the virtual machines. What information is a VM required to provide at checkpoints and what are the consequences of non-compliance? Many questions will need to be answered and these issues will need to be addressed. Will there be standard identity “documents” similar to passports issued by zones according to a set of guidelines? Will checkpoints in the cloud query “Do you have anything to declare?”; will they immediately conduct an invasive search of the VM; or will some middle ground be found?

The issue of whether migration should or should not be used is certainly important for both cloud consumers and providers to consider. However, migration is a commonly used capability in today’s cloud environments. In this work we aim to describe a mechanism by which the cloud provider and consumer can communicate and enforce their requirements for deployment and migration (assuming it is used). Making a determination about the use of migration in any given cloud environment is therefore beyond the scope of this project.

## 9. Resource Instantiation

While the physical world analogy does not quite extend to the creation of new resources (because there is no agency that enforces where people are born), the migration approach can be applied to the instantiation of new resources. In effect, a new VM can be considered as a special migration case, coming from a “null zone” that allows all operations and has an open emigration policy. Under this view, instantiation and migration become essentially the same problem, implemented by the same code.

As with most complicated scenarios, a one-size-fits-all solution will not likely be feasible in the cloud. A proof-of-concept test using the Xen platform is being implemented by researchers at the University of Alaska Fairbanks. While still in the initial design phase, preliminary results are promising.

## 10. Future Considerations

This paper limits the discussion of the analogy to IaaS and the migration of VMs in the cloud. The role is functionally analogous to the INS Inspection Division prior to the combination of the roles into the DHS. It could easily be extended to address data migration (analogous to the duties of the U.S. Customs Inspection and Department of Agriculture's plant and animal inspection before the formation of the DHS). While the organizations have been combined, the individual and unique roles still exist and could provide insight into the development of similar technological solution that would work in the cloud.

Indeed, although this paper focused on VM migration, the same issues arise with data migration. In some legal jurisdictions, PII must have very strict protections; in others, less or no protection is mandated. The owners of the PII may bind requirements for its processing to the data, and the Emigration and Immigration Agents discussed earlier would need to enforce those requirements. Indeed, in some sense a migrating VM is simply data, so the ideas in this paper readily translate to data protection as well.

One issue not explored is the idea of a common framework and language to express the characteristics that affect immigration and emigration. As these describe attributes of both systems (or data) and clouds, one may view this paper as calling out for an attribution system to enable entities to take advantage of cloud services. Perhaps a framework such as is being developed for GENI [6] would prove useful here.

## 11. Acknowledgements

Kara Nance and Brian Hay were supported in part by a grant from the National Science Foundation (DUE-1023125) well as Project Number 1773 of the GENI Projects Office. Kara Nance and Lucas McDaniel were supported in part by the Department of Defense through the Juneau Economic Development Council. Matt Bishop was supported by grants from the National Science Foundation (CCF-0905503 and CNS-1049738), and Project Number 1776 of the GENI Projects Office. This work was facilitated by the Honeynet Project ([www.honeynet.org](http://www.honeynet.org)), a leading international security research organization, dedicated to investigating the latest attacks and developing open source security tools to improve Internet security.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsoring agencies.

## 12. References

- [1] U.S. Customs and Border Patrol. Retrieved September 20 from <http://www.cbp.gov/xp/cgov/about/>
- [2] U.S. Border Patrol. Retrieved September 15 from [http://www.usborderpatrol.com/Border\\_Patrol90.htm](http://www.usborderpatrol.com/Border_Patrol90.htm)
- [3] Ristenpart, T., Tromer, E., Shacham, H., and Savage, S. 2009. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In Proceedings of the 16th ACM Conference on Computer and Communications Security (Chicago, Illinois, USA, November 09 - 13, 2009). CCS '09. ACM, New York, NY, 199-212. DOI=<http://doi.acm.org/10.1145/1653662.1653687>
- [4] Hay, B., K. Nance, and M. Bishop. *Storm Clouds Rising: Security Challenges for IaaS Cloud Computing*. Cloud Computing Minitrack of the Software Technology Track of the 44<sup>th</sup> Hawaii International Conference on Systems Sciences. January 2011.
- [5] Nissinbaum, H. (2010) *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford law Books. Stanford University Press.
- [6] M. Bishop, C. Gates, and J. Hunker, "The Sisterhood of the Traveling Packets," *Proceedings of the 2009 Workshop on New Security Paradigms* pp. 1–12 (Sep. 2009).
- [7] Pauli, Darren, "Google: Who Cares Where Your Data Is?," SC Magazine, Australian Edition, (June 9 2011). Retrieved from [http://www.scmagazine.com.au/News/260041\\_google-who-cares-where-your-data-is.aspx](http://www.scmagazine.com.au/News/260041_google-who-cares-where-your-data-is.aspx)
- [8] Mell, Peter, and Timothy Grance, "The NIST Definition of Cloud Computing (Draft)", National Institute of Standards and Technology. Retrieved from [http://csrc.nist.gov/publications/drafts/800-145/Draft-SP-800-145\\_cloud-definition.pdf](http://csrc.nist.gov/publications/drafts/800-145/Draft-SP-800-145_cloud-definition.pdf).
- [9] VMware, "What is a Virtual Machine?". Retrieved August 30, 2011 from <http://www.vmware.com/virtualization/virtual-machine.html>