Embracing the Cloud for Better Cyber Security

Craig A. Shue and Brent Lagesse
Cyberspace Science and Information Intelligence Research
Computational Sciences and Engineering
Oak Ridge National Laboratory
{cshue, lagessebj}@ornl.gov

***Abstract*— The future of cyber security is inextricably tied to the future of computing. Organizational needs and economic factors will drive computing outcomes. Cyber security researchers and practitioners must recognize the path of computing evolution and position themselves to influence the process to incorporate security as an inherent property.**

The best way to predict future computing trends is to look at recent developments and their motivations. Organizations are moving towards outsourcing their data storage, computation, and even user desktop environments. This trend toward cloud computing has a direct impact on cyber security: rather than securing user machines, preventing malware access, and managing removable media, a cloud-based security scheme must focus on enabling secure communication with remote systems. This change in approach will have profound implications for cyber security research efforts.

In this work, we highlight existing and emerging technologies and the limitations of cloud computing systems. We then discuss the cyber security efforts that would support these applications. Finally, we discuss the implications of these architectural changes, in particular with respect to malware and social engineering.

I. Cloud Computing's Impact

Organizations and users have proven willing to outsource their services and data to remote providers. These providers can offer a varying degree of service, from *software-as-a-service* to *platform-as-a-service* models. They can provide simple data storage or provide processing support. This outsourcing phenomenon fits under the broad paradigm of *cloud computing*. While some may use this term broadly, within this document, we use "cloud computing" only to refer to outsourced processing and storage, as in the case of software and platform services.

A. Cloud and Pervasive Computing

Cloud computing provides an excellent opportunity to expand pervasive computing. In particular, cloud computing can enable systems consisting of resource constrained devices to perform intense computations [1]. Further, cloud computing enables mobile devices to access a large store of information from nearly anywhere. Rather than developing pervasive systems that attempt to intelligently move or cache important data, these systems can now just store the data in the cloud and rely on it to be available when needed elsewhere.

We use an example scenario to demonstrate the relationship between cloud computing and pervasive computing. Consider a typical user, Alice, and her interactions with technology for a day. She may begin her day by waking up to her alarm clock and preparing for work. As she walks to her car, she reads the morning's news on her smart phone that was pre-fetched for her and distributed to her phone by her cloud services. If she has not finished her reading when she gets in her car, the cloud can perform a text-to-speech conversion of her news articles and read them to her as she drives. Upon completion of the news, she can stream her music repository to her car. Once at work, Alice can seamlessly move from her office to meeting rooms throughout the day while maintaing access to all of her information, applications, and sessions due to her cloud services being accessible from any machine. Even after returning home, she can relax with her favorite computer game through her desktop computer. If hosted by her cloud provider, she will not lose progress in her game even if the power fails; instead, she can continue playing through her smart phone.

To enable such a scenario, we must address new challenges in securing the information stored on the cloud and the access to it. The adoption of cloud computing as a part of pervasive systems will affect security in pervasive systems. By using the cloud as a processing and storage powerhouse for pervasive systems, the focus of security in these systems will shift to ensuring that the data and processing controlled by a third party is secure, and the transmission of data between the cloud and the pervasive system is secured. Further, since pervasive systems often enable users to log on from any number of devices, as demonstrated in the example above, authentication mechanisms will also be of high importance.

Cloud computing can advance the pervasive computing goal of making computers invisible to the user. Cloud systems can perform complex tasks while revealing a small, portable interface to the user. Likewise, the demand for small, mobile, and energy-efficient systems by pervasive computing applications will drive the progression of cloud computing. The benefits that cloud computing can bring to pervasive computing must not be negated by leading to insecure systems or obtrusive security mechanisms. Any attack on a system detracts from the invisibility of the pervasive computing vision [2] as it can result in a variety of damaging consequences. Furthermore, any obtrusive security mechanism in a system detracts from this vision and distracts the user from their goals.

B. Current Technologies

Cloud computing provides efficiencies for organizations. Rather than attempting to maintain patching, backup and recovery services, and software licensing for hundreds of machines, organizations employing cloud technologies can simply make users' machines into consoles that access network storage or terminals for remote computation. These cloud-based systems can be instantiated in various approaches:

- **Terminal Services:** Microsoft's Remote Desktop, Virtual Network Computing (VNC), and Citrix's terminal services clients allow users to access an entire desktop environment from a remote server.
- **Web-Based Productivity Tools:** The Docs Web application from Google and the Office Web Apps from Microsoft are two examples of software as a service. These tools provide users with remote data storage and provide software through a Web browser. Google's Chrome operating system is designed to provide only a Web browser, making the computer a terminal to Web-based applications.
- **Window Forwarding:** Unix and Linux-based systems have long had the ability to forward individual graphical windows to remote systems. This allows a client system to render the window and manage the interface while still executing the application on the server system, possibly in closer proximity to the data.
- **Remote Storage:** Amazon's Simple Storage Service and Mozy's online backup services allow a user to store data in remote systems. These services provide redundancy, allowing users to have greater reliability in case of disk or system failures.

The benefits and lower costs of these cloud services enable systems and organizations to offload responsibilities that were previously handled internally. Many of these services scale well, allowing providers to leverage economies of scale to reduce costs while developing specialized services.

C. Emerging Technologies

Recent technological advances can offer greater support to cloud computing. Developers of modern Web browsers have optimized their Javascript engines, including just-in-time compilation and hardware acceleration, to expedite processing for clients. These optimizations allow Web sites to use more complicated client-side functionality without degrading client-side performance. Web browsers now natively support standards such as HTML5, allowing the integration of video and animated content without requiring third-party plug-ins, allowing sites to have rich interactions with the user.

Web browsers also have begun implementing isolation to prevent actions in one Web browser window from affecting other windows. This isolation allows the browser to continue functioning, even if one window malfunctions.

Combined, these Web browser optimizations allow Web sites to provide clients with code that will execute in an independent, isolated environment. This provides flexibility to Web-based applications without endangering the client.

D. Likely Progression

With the national focus on high-speed Internet deployment, increasingly complex cloud services will extend toward many new users and platforms. With evolutionary optimizations in Web browser technology, the continued adoption of rich, Web-based applications is likely. Accordingly, users will have decreased need to install or maintain host-based applications. The Google Chrome OS is an example of an operating system that has been greatly simplified, largely operating to support a Web browser.

With greater network availability, client devices are more likely to resemble terminals to remote resources. This approach will allow pervasive system interfaces to become smaller and lighter without sacrificing functionality. As an example, a recent demonstration showed a smartphone running a notoriously graphically/compute intensive game [3]. In the demonstration, a computation system simply pushed the rendered pixels for the smartphone to display, reducing the smartphone into a simple input/output device. With the pervasive computing researchers working on mechanisms to determine which computation to perform on the client and which to perform remotely, we are likely to see seamless execution across heterogenous devices.

II. Limitations of Cloud Computing

While likely to play an important role in the future of technology, cloud computing is not without its limitations. However, these limitations are also opportunities for cyber security researchers to influence technological advancement in a way that eliminates modern security problems. The following are some prominent limitations:

- **Lack of Control and Interoperability:** When users and organizations place their data in cloud systems and become reliant on a cloud provider, they lose some control and flexibility over their data and processing. Further, because providers may not be interoperable, the users and organizations may be unable to transfer their assets to another provider should the need arise.
- **Lack of Privacy:** To use cloud services, users and organizations expose their data to the cloud provider. In some contexts, this exposure may be unacceptable.
- **Safety of Cloud Servers:** Cloud service providers aggregate a great deal of data, making them an attractive target for attackers. These servers must be protected while still being able to flexibly support services for clients.
- **Client Authentication:** Clients must be able to regularly authenticate to a variety of remote service providers. The current client authentication approach, using passwords, does not scale and largely forces users into poor practices. A better approach must be developed.
- **Resource Allocation:** While cloud providers services are likely to perform most operations, some services may be delivered more efficiently if client support is integrated. Automatically detecting and utilizing client resources is a key area of research to support cloud computing.
- **Connectivity and Mobility:** Cloud services naturally require network connectivity between the client and the service provider; however, these services may be unavailable to users in remote locations or on airplanes.

We now discuss each of these limitations and their opportunities in greater detail.

A. Lack of Control and Interoperability

In out-sourcing their data and computation, organizations and users lose a degree of control over their data. If all of a user's data is stored by the provider and the provider is no longer available, the user cannot obtain his or her data. If the provider changes their terms of service in a way that is not acceptable to the user, the user may have no options to transport his data to another system.

In the case of software-as-as-service providers, organizations are restricted to the applications provided by the cloud service provider. If the cloud service provider does not offer a particular service or feature, organizations have no convenient way of adding the functionality. In traditional computing, users can use scripting or compiled languages to perform processing, giving them greater flexibility and control.

Opportunities: With multiple providers and competition, users can regain control over their data. With a clear, universal application programming interface (API) to cloud services and the user's authorization, competing service providers would be able to access and manipulate the data on another service provider. Such an API would also enable pervasive systems to run seamlessly between providers and allow interaction between users on different providers. Another possible solution is the deployment of systems such as cloudlets [4].

Platform providers host virtual machine images which users must populate to perform their computation. These systems provide users more control, but may require user applications to be written for the system. Some languages and programs are sandbox-aware, such as Java, while others can be executed in `chroot` environments. However, some variants of attack code are able to break out of virtualized environments, allowing access to the system. Accordingly, research that can secure or provably verify the integrity of these environments will allow cloud providers to grant greater access and flexibility to their users without sacrificing security.

B. Lack of Privacy

In many cloud systems, the provider mines the user's data in order to provide advertising or collect aggregated data that can be used to offset the costs of providing the service. This lack of confidentiality may not be acceptable, especially when the cloud system supports a sensitive context-aware system [5]. In other instances, users may lose some degree of legal protection when their data is hosted by a third party.

Opportunities: Recent advances in homomorphic encryption have provided mechanisms for performing operations on the encrypted version of data and deterministically affecting the unencrypted version [6], [7]. For cloud computing applications, this would allow an organization to outsource computation while not exposing the actual information to the remote system. Unfortunately, these approaches often require specialized computation and are computationally expensive.

C. Safety of Cloud Servers

While resource constrained devices may not be able to store all their relevant information locally, they cannot completely offload this information to cloud providers without additional protections. Cloud computing systems have become an attractive target to attackers [8]. Further, while these providers may currently be trustworthy entities, they are not immune to business failures. If a cloud provider's business fails, it may sell off user data as an organizational asset. These providers can also be acquired by other organizations, which would obtain the user's information.

Opportunities: By leveraging support from resource constrained devices, users can reduce the attractiveness of stored cloud data to adversaries. Currently, most cloud data is stored in an unencrypted format or in an encrypted format in which the cloud provider holds the encryption/decryption keys. Instead, we can restrict cloud computing data stores to handling only encrypted data and rely upon the user's explicit approval to access the data. Rather than requiring the user to be involved in encryption/decryption, the user could use a third-party broker to provide unlocking support to the cloud, allowing the cloud to decrypt pieces of data needed to satisfy a user's request. By doing so, resource constrained devices would still be able to interact with the cloud while maintaining the security of their data. With such systems, an insider or external attacker at that organization would be unable to access user data that was not actively being processed by the cloud. When combined with homomorphic encryption schemes for Web application providers, only the user would ever have access to the unencrypted information.

To ensure availability in case of an attack or disaster, cloud data must also be stored redundantly at multiple physical locations. Data storage may be provided by multiple provider organizations to ensure no organization has exclusive control or responsibility over the user/organizational data.

D. Client Authentication

With so much remote execution, cloud computing requires robust credential management that enable secure logins to multiple cloud services from multiple devices in a seamless manner. The password schemes currently employed are a burden on users and have practically forced users into poor practices. Generally, users can remember a small number of passwords, yet each Web resource generally requires users to develop a unique set of credentials. Services such as OpenID, which allow users to have a single set of credentials for multiple sites, are powerful, but may be inappropriate for sensitive institutions such as banks or government sites. Users may instead be able to use one-time-password devices, but they would need to have a unique device for each remote site to prevent one site from being able to use the credentials to authenticate to another.

Opportunities: Rather than use passwords, users can instead leverage mobile devices. If properly isolated from the devices's general computing and communication hardware, a secure dedicated circuit could be created to manage hash chains for multiple sites. To further increase security of the approach, the hash chain could be manually seeded by the mobile device and a visual depiction of the seed value transmitted to the remote site via a camera capture of the mobile device's screen. In Figure 1, we provide a visual depiction of the authentication process. To authenticate, the user could then select the appropriate Web resource and then present the mobile device's screen to the computer's camera, allowing the Web site to obtain the device's hash value from the visual depiction and compare it with the locally derived hash, authenticating the user. This approach allows greater key sizes and entropy than passwords, making exhaustive enumeration attacks challenging while still making the technology usable for people.

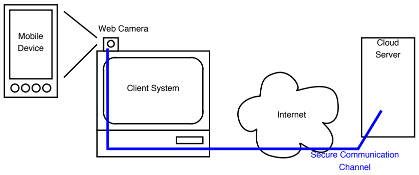

Fig. 1. Mobile device for remote authentication

Fortunately, the technology for this kind of authentication already exists. Software for reading QR-codes is common on many modern smartphones [9]. Likewise, these devices are able to display barcodes that can be read by simple barcode readers. Accordingly, the transmission of hash values from a mobile device to a computer's camera proves little barrier to adoption. The construction of hash chains and rendering them into visual depiction is likewise straight-forward. Adding this component to a mobile device in a secure way may require some separate hashing hardware and storage for initial hash values; however, device manufacturers may be able to reuse the device's screen and input hardware to reduce costs.

While hardware-based hashing operations may significantly improve the security and usability of remote authentication, there are important questions that must be addressed. For example, what happens if the device is lost or stolen? How would the credentials be revoked and how would a user be able to reestablish their identity and control? Research in addressing authentication management techniques is critical to cyber security.

E. Resource Allocation

To support cloud systems, providers often build efficient data centers with mass storage and computational resources. These centers can perform all the computation for resource-constrained devices, such as smart phones. However, for other clients, such as a user's desktop machine, the client machine may be able to perform some of the computation. Further, computing on the client may be more efficient, improve performance, and scale better. While not a direct security concern, the mechanisms to offload this computation must protect the computation from any attackers.

Opportunities: Improvements to resource allocation schemes and the dynamic determination of what computation to perform in the cloud and what to perform on the client will aid the cloud experience. In particular, these resource decisions must reflect a dynamic network environment: some systems may have low bandwidth connections while others may have excellent connectivity. These systems must be able to take these factors and any network congestion into data transfer and computation location strategies.

Distributed computing applications have leveraged volunteered resources from systems across the Internet to perform large-scale computations. Cloud computing systems may also be able to leverage distributed resources in clients in exchange for lower rates or better performance guarantees for their clients. In distributed computing environments, we must have mechanisms to ensure that remote systems are trustworthy and providing correct responses. In current systems, redundancy is used to detect dishonest systems and protect computation accordingly. Innovations to enable trusted distributed computation with minimal redundancy would be preferable.

F. Connectivity and Mobility

While cloud data storage relieves users of the burden of managing their own data, it also requires connectivity to the cloud. Users may need to use their applications and access their data without the opportunity to connect to their cloud service provider. In this case, the client should be able to prefetch data that the user will likely need during the period of time spent offline. Similar to offloading some computation to clients in Section II-E, selecting the data to offload for offline operation is not directly a security concern, but the transmission, storage, and computation of that data is a security concern.

Opportunities: Predictive and opportunistic pre-fetching by mobile and intermittently connected devices are essential to the user experience of cloud systems in several situations. Many factors must be taken into account to efficiently and successfully prefetch necessary data. The client must account for the tasks that the user will likely need to perform, the data that is related to those tasks, the ability of the client's hardware to store that data and to perform the necessary computations, and for mobile devices, the energy constraints of the device.

Given that the correct data is selected to be fetched from the cloud servers, these systems must be able to securely transmit and store the data. While some machines may be capable of performing operations on homomorphically encrypted data, it is unlikely that most mobile devices can do so within the constraints of available battery life. As a result, data must be unencrypted for operations, likely by a trusted third party that then re-encrypts it with a symmetric key prior to transmission to the mobile device.

III. Implications of Research and Adoption

Continued adoption and advances in security mechanisms for cloud computing may shift the use of computation, networking, and the cyber security battleground. In particular, this shift may affect malware usage, organizational auditing and insider threat detection, social engineering attacks, and traditional ISP revenue streams. We now discuss each of these implications in detail.

A. Rendering Malware Obsolete

With secure input/output and a secure connection to a remote destination, an compromised system could still be securely utilized, substantially affecting the utility of malware for adversaries. Input and output components with trusted platform modules (TPMs) and security-enabled interconnects, like HDMI, can allow users and remote resources to encrypt transmission before using a system's general resources (e.g., memory, processor), preventing traditional malware from being effective. The recently announced Intel Insider [10] technology allows a remote resource to stream encrypted video content to a system's GPU, where it is then decrypted and displayed. This approach is designed to prevent software on the system from being able to record the video content. Such an approach could be used for securing user access as well: a Web page could generate a random keyboard for a user to click on to enter their password. If an attacker were unable to see the user's screen, they would be unable to harvest the user's credentials, even if they had full control over the operating system and input devices.

While an attacker may be unable to intercept credentials, there may be benefit to compromising systems if they could be used for attacks. Part of cloud computing revolves around clients executing arbitrary code from their server in an isolated environment, preventing the code from affecting other sites or resources. Continued advancements in isolation will prevent Web applications from being able to attack other applications. Accordingly, malware from within a Web application is unlikely to be a significant risk factor.

Other malware attacks may be launched from outside Web applications. However, on operating systems that solely support a Web browser, no new applications are likely to be needed. This makes it possible to filter malware. Only operating system updates and updates to the Web browser need be supported on these systems. With proper code signing techniques, only the updates to these critical systems would be authorized, preventing the introduction of any third-party applications. Further, these systems may simply prohibit non-system local storage, since all data and applications would be available from cloud-based providers. Accordingly, on these systems, there would be no file-based attack vector.

On other systems, malware and exploits may target vulnerabilities in the operating system or Web browser. Even if the attacker were in complete control of the system, they would not have access to the user's data, since it would be housed on remote systems. Instead, attackers may simply attempt to monitor communication between the user and the Web server. Attackers attempting to obtain login or account credentials can be thwarted simply by using one-time password tokens, or the approach proposed in Section II-D, as the information they gain will not be usable.

B. ISPs as Cloud Providers

Internet Service Providers (ISPs) often lament their continued transition to simple data carriers. They continue to look for mechanisms to compete based on the services they offer while decreasing their overhead costs. For example, services like Akamai serve as content distribution networks and provide caching for popular Web pages. This approach allows an ISP to provide faster service to their clients while reducing the amount of traffic they transmit to peers, reducing costs. By embracing cloud services, these ISPs could provide lower-latency cloud experiences for their users and likewise reduce peering costs. ISPs could then distinguish themselves from their competitors by offering greater services on their clouds. Rather than simply offering connectivity, these organizations could blur the line into becoming content and software providers.

ISPs have embraced such models in the past. Cable television providers have explored remote digital video recorder (DVR) services, allowing their users to control a DVR hosted at the cable provider. This allows the cable provider to reduce the amount of customer premise equipment and costs while providing the same quality of service to which users are accustomed. ISPs have also placed themselves in the position of offering software to users. Many high-speed ISPs offer complimentary security software, such as anti-virus, to their users to reduce incidents of infection. Typically, these ISPs contract third-parties to provide this software.

By becoming a cloud provider, ISPs can become a hosting provider for third-party software. Rather than become involved in application development, the ISP can simply host software that might otherwise be written for a user's client. The ISP can then market their support for the latest high-end games and applications while indicating they will work on a customer's machine, regardless of processing and video card capabilities while connected to the ISP network.

1) Thin Clients: If ISPs decide to provide their own cloud services as a benefit to subscribers (in much the same way that they have previously used web hosting and email services to attract subscribers over the past couple of decades), one natural progression is that ISPs could provide thin clients for free or discounted with a service contract. These thin clients would be designed to perform basic processing and to display information from the ISP's cloud services. Cloud services could also act as an interface between the user and the rest of the Internet. Data that the user requests could be acquired by the ISP's cloud and presented in a manner that requires minimal effort by the user's thin client since most of the computation will be performed on the cloud.

2) Opportunistic Systems: Cloud systems can provide further benefit to mobile users through pre-distributing content to the users during off-peak times. Users have a tendency to charge their mobile devices such as smart phones and laptops overnight and are likely to continue doing so given the lower cost of energy during off-peak hours. This fact can be leveraged to predictively pre-distribute content to mobile users. By pre-distributing content while mobile devices are not using energy from the battery, we believe we can greatly extend the life of mobile devices. For example, ISPs can provide a service to acquire data that is commonly accessed by the user, such as morning news websites or music, and push them to the user's mobile device shortly before their alarm goes off. Further, the user's mobile device could contact the ISP's cloud services any time it is plugged into a power source and opportunistically utilize cloud services and acquire data to extend battery life.

C. Higher-Level Semantic Logging

Organizations traditionally gather low-level behaviors of systems to detect the presence of compromise or other attacks on the network. Other systems use application-layer auditing, which may not detect actions that take place outside the application. Finally, insider threat detection systems often rely on low-level system call analysis to profile users and their actions. This direction can be exceedingly difficult: it can be a challenge to relate system calls to their higher-level functions, let alone determine if the action is legitimate. With cloud-based software-as-a-service systems, the software and data are stored on the cloud. Accordingly, the client must routinely synchronize with the cloud. Cloud systems can use this synchronizing to construct an audit log for all access and alterations. This information can be supplied to anomaly detection software and insider threat identification systems. This feedback can also be used to help application designers adapt their software to enhance user productivity.

D. Rise in Social Engineering

Attackers typically focus on the weakest link in a system. With technical advancements in computer and network security, attackers have focused on deceiving the users of systems. Phishing attacks, in which the attackers impersonate a legitimate institution to gain the user's trust, are constantly present on the Internet. Typically, these phishing attacks are successful even with little sophistication. However, dedicated phishers can craft messages that are indistinguishable from legitimate messages and lure users into unsafe practices.

With cloud systems, these attacks become more potent. Rather than simply gaining access to a checking account number, phishers in cloud-based systems can gain direct access to all the user's data from the cloud. Modern browsers and software are often unable to effectively explain the problems with untrustworthy systems to users. Further, these programs often produce warnings when a system is actually legitimate (such as a certificate that is recently expired). These warnings desensitize users to actual threats. Instead, systems must increasingly make decisions without relying on the user to determine if the action is safe.

E. Barriers to Research

Clouds can also make examinations of human behavior and systems more difficult because of their closed nature. As an analogy, we consider email and social networking sites. Security researchers could previously analyze email corpa to find examples of spam attacks while social networking sites are closed and make research difficult. This lack of openness prevents researchers from detecting Internet-wide threats and understanding the mechanisms attackers are using. This lack of information sharing gives attackers the advantage: they can discover an exploit and use it on multiple sites. Some sites may discover and address the exploit while others may remain oblivious. Cloud systems may be similar: since the interaction between the cloud and its users are private, researchers will be unable to generalize threats and interactions in a way that can benefit everyone.

IV. Conclusion

The continued movement towards cloud computing will have a direct impact on cyber security research. Since pervasive computing has generally focused more on improving functionality and reliability, we see a transition to using a cloud computing backbone in pervasive systems as an opportunity to bring stronger security to pervasive systems. If cyber security researchers are involved in the evolution of this process, they can influence the process and change the current cyber security battlefield to one more amenable to the defenders. Advancements in virtual machine isolation, homomorphic encryption, client authentication, resource management, and secure opportunistic computing will facilitate the adoption of cloud computing while ensuring greater security and privacy for users.

Acknowledgements

This work was sponsored by a contractor of the United States Government under contract DE-AC05-00OR22725 with the United States Department of Energy. The United States Government retains, and the publisher, by accepting this submission for publication, acknowledges that the United States Government retains, a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this submission, or allow others to do so, for United States Government purposes.

References

[1] B.-G. Chun and P. Maniatis, "Augmented smartphone applications through clone cloud execution," in *Proceedings of the 12th conference on Hot topics in operating systems*. Berkeley, CA, USA: USENIX Association, 2009.

[2] M. Weiser, "The computer for the 21st century," *Scientific American*, vol. 265, no. 3, pp. 66–75, January 1991. [Online]. Available: http://www.ubiq.com/hypertext/weiser/SciAmDraft3.html

[3] J. Stokes, "AMD's next-gen GPU powers Crysis on an iPhone," 2009. [Online]. Available: http://arstechnica.com/hardware/news/2009/09/amds-next-gen-gpu-powers-crysis-on-an-iphone.ars

[4] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," *IEEE Pervasive Computing*, vol. 8, pp. 14–23, October 2009. [Online]. Available: http://portal.acm.org/citation.cfm?id=1638591.1638731

[5] X. Jiang and J. A. Landay, "Modeling privacy control in context-aware systems," *IEEE Pervasive Computing*, vol. 1, pp. 59–63, July 2002. [Online]. Available: http://dx.doi.org/10.1109/MPRV.2002.1037723

[6] C. Gentry and S. Halevi, "A working implementation of fully homomorphic encryption," in *EUROCRYPT*, 2010.

[7] C. Gentry, "Computing arbitrary functions of encrypted data," *Communications of the ACM*, vol. 53, no. 3, pp. 97–105, 2010.

[8] A. Jacobs and M. Helft. (2010) Google, citing attack, threatens to exit china. [Online]. Available: http://www.nytimes.com/2010/01/13/world/asia/13beijing.html

[9] J. Rouillard, "Contextual QR codes," in *Computing in the Global Information Technology, 2008. ICCGI '08. The Third International Multi-Conference on*, Aug. 2008, pp. 50 –55.

[10] N. Knupffer, "Intel insider - what is it?" Jan. 2011. [Online]. Available: http://blogs.intel.com/technology/2011/01/intel_insider_-_what_is_it_no.php