10-2011

Privacy in "The Cloud": Applying Nissenbaum's Theory of Contextual Integrity

Frances Grodzinsky
Sacred Heart University, grodzinskyf@sacredheart.edu

Herman T. Tavani
Rivier College

Recommended Citation

Grodzinsky, F. and H. Tavani. "Privacy in "The Cloud": Applying Nissenbaum's Theory of Contextual Integrity." ACM SIGCAS Computers and Society 41.1 (2011): 38-47.

Privacy in "the Cloud": Applying Nissenbaum's Theory of Contextual Integrity

F.S. Grodzinsky
Sacred Heart University
5151 Park Avenue
Fairfield, CT USA
grodzinskyf@sacredheart.edu

H. T. Tavani
Rivier College
420 Main St.
Nashua, NH USA
htavani@rivier.edu

Abstract
The present essay is organized into five main sections. We begin with a few preliminary remarks about "cloud computing," which are developed more fully in a later section. This is followed by a brief overview of the evolution of Helen Nissenbaum's framework of "privacy as contextual integrity." In particular, we examine Nissenbaum's "Decision Heuristic" model, described in her most recent work on privacy (Nissenabum 2010), to see how it enables the contextual-integrity framework to respond to privacy challenges posed by new and emerging technologies. We then apply that heuristic device to questions surrounding one aspect of cloud computing – viz., "cloud storage" technology. In particular, we focus on current practices affecting Google Docs as an instance of a cloud-storage system.

Keywords
Cloud computing, Google Docs, Privacy as Contextual Integrity, Cloud storage

Introduction
The brave new world of Cloud Computing offers many new benefits provided that the privacy and security risks are recognized and effectively minimized. (Ann Cavoukian, 2008: 26).

A right to privacy is neither a right to secrecy nor a right to control but a right to *appropriate* flow of information…but what this amounts to is a right to contextual integrity and what *this* amounts to varies from context to context. (Helen Nissenbaum, 2010: 127)

In this essay, we apply Helen Nissenbaum's theory of privacy as contextual integrity to "cloud computing." In a previous work (Grodzinsky and Tavani, 2009), we examined some implications of that privacy framework for another relatively recent technology: blogging. However, in that essay, we worked mainly from the account of contextual integrity in three of Nissenbaum's earlier works (1997, 1998, 2004). Here, we draw from a revised, expanded, and more robust account of the contextual-integrity framework included in Nissenbaum (2010).

An earlier version of this paper was presented at the Ninth International Conference on Computer Ethics - Philosophical Enquiry (CEPE 2011), Milwaukee, WI, May 30-June 2, 2011 (and printed in the CEPE 2001 Proceedings).

We begin with a brief analysis of what is meant by "the cloud" in the context of computing. Next, we examine Helen Nissenbaum's framework of "privacy as contextual integrity" and describe some ways that it has evolved in her recent book (Nissenabum 2010). In particular, we examine Nissenbaum's "Decision Heuristic" model, which is designed to respond to privacy challenges posed by new and emerging technologies, such as cloud computing. Before directly applying that model, however, we first differentiate among three distinct "delivery models" and three distinct "deployment models" of cloud computing; we then focus our attention on systems and practices surrounding cloud-storage technology in the context of a "software as a service (SaaS) Private Cloud." In the final section, we apply Nissenbaum's decision heuristic to some specific questions affecting Google Docs, as one example of a cloud-storage system.

Computing and "the cloud": some preliminary remarks

What, exactly, is *cloud computing*? Although it is difficult to give a precise definition of cloud computing,[11] some examples of this technology include: web-based email services such as Yahoo; photo storing services such as Google's Picassa; online computer backup services such as Mozy; and file transfer services such as YouSendit (Privacy Rights Clearinghouse: 2008). Major corporations that currently dominate the "cloud computing space" include IBM, Yahoo, Google, IBM, Amazon, Oracle, and Microsoft.

Knorr and Gruman (2010) point out that "the cloud" has been used as a "metaphor for the Internet." They also note that when applied to computing, the cloud can have two different meanings, one which is very narrow and one which is very broad. In its narrow sense, the cloud is often defined as an "updated version of utility computing," which is roughly equivalent to the kinds of "virtual servers available over the Internet." In the broader sense of cloud computing, the cloud refers to any computer resources that are used "outside the firewall," which can include conventional outsourcing of information-technology-related services. In both cases, users have very little "control over or direct knowledge about how their information is transmitted, processed, or stored" (Privacy Rights Clearinghouse (1). This, of course, raises concerns for personal privacy.

According to the Privacy Rights Clearinghouse, a serious problem with cloud computing that needs to be answered from the perspective of personal privacy is: How does the host protect the user's data? Other, related questions, which affect both the integrity of, and users' access to, data in the cloud include: (i) Who owns the data? (ii) Can the host deny a user access to his/her own data? (iii) If the host company goes out of business, what happens to the users' data? (Privacy Rights Clearing House)

Although cloud computing clearly raises some serious concerns for users, we can also ask which kinds of advantages, if any, this technology may offer them. According to Cavoukian (2008), the Cloud offers flexibility and security to users who no longer have to worry about how to protect their data. It allows users to access their data via the

[11] Zeng and Cavoukian (2010: 3) describe "the Cloud" as "a broad, loosely-defined construct," which refers both to "services accessed via, and delivered through, the Internet and the hardware and system software in remote datacenters that provide those services." Zeng and Cavoukian also believe that cloud computing "changes the way we think about computing by decoupling data processing, data retention, and data presentation – in effect, divorcing components from location."

Internet and enables users to work on local, less expensive platforms. While this is very appealing to business owners, who are relieved of the burden of securing data, it can only be effective if users trust their cloud service providers. But is their data secure in the cloud? Pieters and van Cleeff (2009) worry about the ability to secure data, including a user's personal data, because of what they call the "deperimeterization" of information security, which, they believe makes it difficult to measure risk for achieving adequate security. A similar concern is also raised by Zeng and Cavoukian (2010) who worry about the "blurred security perimeter." And Cavoukian (2008) further argues that for cloud computing to be fully realized, users will have to have confidence that their personal information is protected and that their data in general is both secure and accessible.[12] We believe that this is a key challenge for the efficacy of cloud computing in general and cloud storage in particular.

The evolution of the contextual-integrity framework

Nissenbaum's privacy framework has evolved over the past fifteen or so years, as illustrated in a series of journal articles and other published works. In her essay "Privacy as Contextual Integrity," Nissenbaum (2004) expands upon some core concerns affecting (what she calls) "the problem of privacy in public," which she introduced in two earlier essays (Nissenbaum 1997, 1998). Her theory (as expressed in Nissenbaum 2004) is based on two principles: (i) the activities people engage in take place in a "plurality of realms" (i.e., spheres or contexts); and (ii) each realm has a distinct set of norms that govern its aspects. Nissenbaum argues that norms affecting these two principles both shape and limit our roles, behavior, and expectations by governing the flow of personal information in a given context.[13] There are two types of informational norms in Nissenbaum's privacy scheme: (a) norms of appropriateness, and (b) norms of distribution. Whereas the first of these determines whether a given type of personal information is either *appropriate or inappropriate* to divulge within a particular context, the second set of norms restricts the flow of information within and across contexts. When either of these norms is "breached," a "violation of privacy occurs" (Nissenbaum, 2004: 125).

We have argued elsewhere that one virtue of Nissenbaum's theory is that it illustrates why we must always attend to the *context* in which personal information flows, not the nature of the information itself, in determining whether normative privacy protection is needed.[14] Although we argued in a more recent work (Grodzinsky and Tavani 2009) that

[12] Cavoukian (11) believes that the "full potential of the cloud" will be realized when users can "seamlessly tap into and combine a wide range of online services" beyond those limited to laptop and desktop computers. For example, they could access data from cell phones, personal digital assistants, smart cards, etc.

[13] The contextual-integrity model proceeds on the assumption that there are "no areas of life are not governed by norms of information flow" (Nissenbaum 2004: 137).

[14] See, for example, Grodzinsky and Tavani (2005, 2008), where we argued that rather than focusing on the nature of the information included in a "P2P situation" or context – i.e., asking whether or not it should be viewed as private – we can ask whether P2P situations (or contexts in general) deserve protection as "normatively private situations." We also showed that "situations" (Moor 1997) are analogous to "contexts" in Nissenabum's scheme. Some similarities and differences between Nissenbaum's and Moor's context-based privacy theories are examined more fully in (Tavani, 2007, 2008a, b).

the contextual-integrity framework could also inform the privacy debate affecting a relatively recent technology, viz., blogs and the blogosphere, we also noted that critics might object that the contextual-integrity model is not *easily* adaptable to new technologies that emerge in which no clearly articulated practices, expectations, or norms govern the flow of personal information. Nissenbaum (2010) anticipates this kind of criticism affecting the application of her framework to new and emerging technologies by acknowledging a "blind spot" in her original privacy framework. In response to her potential critics, Nissenbaum puts forth a "Decision Heuristic" that addresses such attacks head-on.

Nissenabum (2010: 148) notes that the contextual-integrity framework can "guide an assessment" of a "problematic new practice resulting from the development of a novel technical device or system" by asking the question: "Does the practice in question violate any context-relative informational norms." She notes that these norms, in turn, are characterized by "four key parameters: contexts, actors, attributes, and transmission principles" (Nissenbaum: 140). Whereas contexts are the "backdrop of informational norms," actors can be understood in terms of three components: "senders of information, recipients of information, and information subjects" (141). "Attributes" can be understood as a "*type* or *nature* of information," while transmission principles serve as a "constraint on the flow of information…from party to party in a context" (143-145, Italics Nissenbaum).

Nissenbaum's "decision heuristic"

A key objective of Nissenbaum's "decision heuristic" is to provide an approach that enables us both to: (a) understand the "source or sources of trouble in new and emerging technologies," and (b) evaluate the "system or practice in question"(Nissenbaum, 2010: 181). Her decision heuristic includes a series of guidelines articulated in nine steps:

1. Describe the new practice in terms of information flows.
2. Identify the prevailing context… and identify potential impacts from contexts nested in it…
3. Identify information subjects, senders, recipients.
4. Identify transmission principles
5. Locate applicable entrenched informational norms and identify significant points of departure.
6. Prima facie assessment…A breach of information norms yields a prima facie judgment that contextual integrity has been violated because presumption favors the entrenched practice.
7. Evaluation I: Consider moral and political factors affected by the practice in question…
8. Evaluation II: Ask how the system or practices directly impinge on values, goals and ends of the context...
9. On the basis of these findings, contextual integrity recommends in favor of or against systems or practices under study.... (Nissenbaum: 182)

We will see that the criteria included in Nissenbaum's heuristic device are especially helpful in analyzing some privacy concerns that arise in the case of cloud computing. The first five steps or components in this model are "descriptive"; they help us to gain a clear understanding of the features in a new technology (such as cloud computing) that may have implications for privacy. Steps 6-9, on the contrary, are essentially normative in nature since they guide us in evaluating the features and practices associated with a new technology.[15]

Applying nissenbaum's "decision heuristic" to cloud computing

As noted above, the concept of "cloud computing" is multi-faceted and thus complex. This can make it difficult to analyze the cloud in terms of a single context or even in terms of a simple model. Zeng and Cavoukian (2010: 3-4) differentiate three distinct *delivery models* for cloud computing: (1) Software as a Service (or SaaS), which "delivers applications to consumers (either individuals or enterprises) using a multitenant architecture": (2) Platform as a Service (PaaS, which delivers "delivers development environments to consumers"; and (3) Infrastructure as a Service (IaaS), which delivers resources such as servers, connections, and related tools necessary to build an application from scratch." The authors also distinguish three *deployment models* for the cloud: (a) the "Public Cloud," which is "provided by an off-site third-party environment service provider who shares resources in a multitenant operating environment, and bills on a utility computing basis"; (b) the "Private Cloud," which is "provided by an organization or its designated service provider and offers a single-tenant operating environment:; and (c) the "Hybrid Cloud," which is "a composition of two or more Clouds (public or private) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability" (Zeng and Cavoukian, 2010:4). Using this model, the contexts of the Cloud are shown in Table 1.

SaaS—Public Cloud	PaaS—Public Cloud	IaaS—Public Cloud
SaaS—Private Cloud	PaaS—Private Cloud	IaaS—Private Cloud
SaaS—Hybrid Cloud	PaaS – Hybrid Cloud	PaaS—Hybrid Cloud

Table 1

We can now see why it is difficult to determine what, exactly, *the* "context" of the cloud would be. It might be that it is made up of several contexts, depending on the type of delivery model identified in the table above. However, we focus our analysis on the concept of cloud *storage, as the aspect of cloud computing that most significantly affects the privacy of personal information and accessibility of data*. In this context, the "cloud" is a repository in which one's data is stored – i.e., residing in a "space" or a "device" that is geographically remote from and independent of the hard drive or the storage device on (or connected to) one's physical personal computer. We examine Google Docs as an instance of a cloud-storage system. Using the categories articulated in Table 1, we would

[15] For some additional accounts of how normative features and values embedded in a particular technology can have privacy implications, see Brey (2000) and Friedman et al. (2008); Friedman et al. use cookies technology to illustrate this connection.

classify Google Docs as SaaS, because it: (a) provides tools to create applications, and (b) is a repository for the user's data. We further classify Google Docs as private – i.e., a private SaaS – because one is required to sign up for this service and consent to its policy; it is also password-protected and is under the rubric of Google, which is a corporation. However, we should note that Google Docs is also a free service that anyone can use.

The primary question associated with an evaluation of Google Docs vis-à-vis Nissenbaum's framework is: Which kinds of "context-relative informational norms" apply? Also, are there any clear bounds prescribing this context? It is in the application of Nissenbaum's heuristic to Google Docs that we can ultimately determine whether cloud-storage technology has changed our understanding of data storage from the perspective of personal privacy.

Google docs and the contextual-integrity framework

In addition to briefly examining Google Docs as one instance of a cloud- storage system, we also examine Google's privacy policy and consider how the privacy parameters governing that storage context have evolved thus far. We then ask: In which ways does storing one's personal data in the Google Docs cloud threaten personal privacy? In particular, we ask whether this specific cloud-storage context can satisfy the privacy requirements specified in Nissenbaum's contextual-integrity framework vis-à-vis her decision heuristic.

Google Docs can be viewed as an example of a Web-based application that allows users to create and edit documents online. It is also a document-sharing service that allows users to collaborate with others in real-time activities that include word processing, spread sheets, presentation, and data storage. Earlier, we saw why Google Docs can also be viewed as a (private) "software as a service" (or SaaS) version of cloud computing.

We begin our analysis by noting that Google Docs purports to be sensitive to the privacy concerns of its users. When, in March 2009, Google reported that a bug in Google Docs had allowed unintended access to some private documents, and it was estimated that "0.05% of documents stored via the service were affected by the bug, Google claimed that the bug had been fixed within a few days (Breitbart, 2010). However, in regards to "personal account" information, Google Docs includes the following statement with respect to "Privacy and Security: Personal Account Information":

> We may use information that is not personally identifiable to improve the quality of other Google services and provide you with a seamless experience when using multiple Google products. Please be assured that we don't rent or sell your personally identifying information to other companies or individuals without your explicit consent. (http://docs.google.com/support/bin/answer.py?answer=47585&ctx=sibling)

Does this policy satisfy the requirements specified in the contextual-integrity framework? First, we should point out that Google's policy does not state explicitly what is meant by, or included under the category, "personally identifiable information." Providing a precise

definition of what counts as information that is personally identifiable has not always proved easy, and perhaps it would be unfair to hold Google to a higher standard here. But even assuming that such a definition could, in principle, be achieved, it would not necessarily resolve the questions affecting the contextual-framework model. Recall that in Nissenbaum's scheme, protecting the privacy of individuals cannot necessarily be reduced to merely protecting "personally identifiable information." For example, information that may not be readily viewed as "personally identifiable," as well as information about persons that may seem innocuous needs to be protected in a specific context to fully enjoy "integrity." As Nissenbaum points out, it is not simply the kind of personal information – i.e. information that is intimate, confidential, or sensitive – that needs protection; it is also information pertaining to persons that easily shifts between contexts with little or no normative protection that can be problematic.

Google Docs in particular, and cloud-storage systems in general, are examples of what Nissenbaum calls "socio-technical" systems/practices: technical systems as they function in "social contexts" (184). These systems can affect a range of contexts in a variety of ways, much in the same way that Nissenbaum notes that telecommunications systems do, because they also depend on "a host of factors." She goes on to note that for technologies of this type, "the ideal is flexibility in how well they are adapted to particular contexts so the flows of information may be tailored according to the requirements of entrenched informational norms" (184-185). Nissenbaum uses the example of Caller ID as an instance of how the telecommunications industry was able to design into their systems some fine-tuned settings responding to challenges that threatened contexts of social life. In the same way, the cloud-storage industry may follow the telecommunications field in offering technical services that also respect the integrity of social life. Has it? We will use Nissenbaum's heuristic to approach this question.

Socio-technical systems and practices need to be analyzed not only from descriptive aspects (included in Steps 1-5) of the "decision heuristic," but also in terms of their relevant moral and political implications. This involves applying Steps 6-9 to the socio-technical practice under consideration – in this case, concerns affecting cloud storage practices on Google Docs. In some respects, we have already addressed key elements in the descriptive components of the heuristic contained in the first six steps in our earlier discussion of cloud storage. For example, we have described cloud storage in terms of its "information flows" and we have identified the "information subjects, senders, and recipients" involved. We have also identified the "transmission principles" and located the "applicable entrenched informational norms" that pertain to data storage. But we have not yet determined whether cloud-storage systems in general, or Google Docs in particular, either defy or breach any of the "entrenched norms"; so it is difficult for us to make the "prima-facie assessment" (described in Step 6) at this point. To determine whether contextual integrity is violated in cloud-storage practices affecting Google Docs, we will need to proceed directly to the evaluative criteria included in Steps 7-9 of the decision heuristic.

In Step 7, the first level of evaluation, Nissenbaum asks us to consider what kinds of harms – specifically, which kinds of threats to autonomy and freedom – the socio-technical system in question poses. For example, what implication might practices affecting Google Docs have for "justice, fairness, equality, social hierarchy, democracy,

and so on" (182). Nissenbaum notes that whereas in some cases, the implications for these factors may be clear cut, in many cases further evaluation is required; this, in turn, is carried out at Step 8, where Nissenbaum invites us to consider how a system or practice, such as storage of data on Google Docs, "directly impinges" on "values, ends, and goals of the context" (182). Finally, at Step 9, we are asked to make a recommendation for or against the socio-technical system/practice, based on the findings acquired from the two previous evaluative steps.

It is not clear to us, on the basis of what we have seen thus far that contextual integrity is violated in the Google Docs "cloud storage context". We distinguish Google Docs as a distinct context (SaaS-Private) that is based on a particular set of practices involving cloud- storage systems in general. We have seen that Google has a specific privacy policy that is arguably transparent and is designed to protect personally identifiable data that resides in its space. Its goal is to offer a free service, to all who sign up, and to enable users to have control over their own data. On our view, Google has met each of these goals. But we have also seen that Google Docs is not impervious to bugs (and from potential cyber attacks) that threaten the integrity of the data stored in its space. This concern, however, is as much a worry from the perspective of security as it is from privacy; and given Google's timely response in 2009 to fixing the bug in its system, it would seem that this company takes the privacy interests of its users (in its cloud storage space) seriously. So based on our analysis, thus far at least, of Google Docs vis-à-vis Nissenbaum's decision heuristic, it would seem that the contextual-integrity framework would "recommend in favor of" this context.

However, we have also seen that there are a wide variety of cloud-storage systems and practices (based on the scheme we articulated in Table 1, which identified nine distinct contexts). So, we cannot infer that every permutation of SaaS, or even every instance of Saas-Private, storage-device systems will satisfy the privacy requirements for the framework of contextual integrity. In fact, the decision heuristic would have to be applied to each distinct context and possibly extended to the applications that impact stored data. We should also note that Google Docs, while seeming to comply with the requirements of contextual integrity at present, could conceivably modify it practices for data storage at some future point in time in ways that would defy or violate one or elements in Nissenbaum's decision heuristic. So, any conclusion we reach about privacy in the context of cloud storage (systems and practices) must be guarded, as it is limited to our analysis of current practices involving Google Docs. On the other hand, however, it is encouraging for users who wish to turn to the Cloud as a repository for storing their personal data to know that Google Docs currently operates in ways that satisfy the criteria specified in the framework of privacy as contextual integrity.

Conclusion

We have examined some aspects of cloud computing from the perspective of personal privacy. In particular, we examined systems and practices affecting cloud storage from the framework of Nissenbaum's decision heuristic for her theory of privacy as contextual integrity. We applied Nissenbaum's model to Google Docs, as an example of a socio-technical system/practice involving cloud storage. We believe that Google Docs conforms to the requirements of the decision heuristic within the framework of contextual

integrity. However, we also saw that there are other variations of cloud storage in which the practices used may not necessarily comply with these standards. So, any conclusions drawn about privacy in the context of cloud-storage systems in general must be tentative, even if we can affirm the privacy practices involving Google Docs.

References

Breitbart (2010) "Google software bug shared private online documents," http://www.breitbart.com/article.php?id=CNG.54c3200980573ae4c... Retrieved March 25, 2011.

Brey, P. (2000) "Disclosive Computer Ethics," *Computers and Society*, Vol. 30, No. 4, pp. 10-16.

Cavoukian, A. (2008) Privacy *in the Clouds: A* White Paper *on Privacy and Digital Identity: Implications for the Internet*. Available at: http://www.ipc.on.ca/images/resources/privacyintheclouds.pdf.

Freidman, B., Kahn, P., and Borning, A. (2008) "Value Sensitive Design and Information Systems." In Himma, K. E. and Tavani, H. T., eds. *The Handbook of Information and Computer Ethics*. Hoboken, NJ: John Wiley and Sons, pp. 69-101.

Grodzinsky, F. S. and Tavani, H. T. (2005) "P2P Networks and the Verizon v. RIAA Case: Implications for Personal Privacy and Intellectual Property," *Ethics and Information Technology*, Vol. 7, No. 4, pp. 28-39.

Grodzinsky, F. S. and Tavani, H. T. (2008) "Online File Sharing: Resolving the Tensions between Privacy and Property," *Computers and Society*, Vol. 38, No. 4, pp. 28-39.

Grodzinsky, F. S. and Tavani, H. T. (2009) "Can the 'Contextual Integrity' Model of Privacy Be Applied to Blogs and the Blogosphere?" In M. Bottis, ed. *Eighth International Conference on Computer Ethics: Philosophy Enquiry*. Athens, Greece: Nomiki Bibliothiki, pp. 302-311. [Reprinted in slightly revised form in *International Journal of Internet Research Ethics*, Vol. 3, No. 1, 2010, pp. 38-47.]

Knorr, E., and Gruman, G. (2010) "What Cloud Computing Really Means," *InfoWorld*. Retrieved from http://www.infoworld.com/d/cloud-computing/what-cloud-computing-really-means-031?page=0,0 December 3, 2010.

Miller. K. and Voas, J. (2010) "Ethics in the Cloud," *IT Professional*, Vol.12, No.5, pp. 4-5.

Moor, J. (1997) "Towards a Theory of Privacy in the Information Age," *Computers and Society*, Vol. 27, No. 3, pp. 27-32.

Nissenbaum, H. (1997) "Toward an Approach to Privacy in Public: Challenges of Information Technology," *Ethics and Behavior*, Vol. 7, No. 3, pp. 207-219.

Nissenbaum, H. (1998) "Protecting Privacy in an Information Age," *Law and Philosophy*, Vol. 17, pp. 559-596.

Nissenbaum, H. (2004) "Privacy as Contextual Integrity," *Washington Law Review*, Vol. 79, No. 1, pp. 119-157.

Nissenbaum, H. (2010) *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Palo Alto, CA: Stanford University Press.

Pieters, W. and van Cleeff, A. (2009) "The Precautionary Principle in a World of Digital Dependencies," *IEEE Computer*, Vol. 42, No. 8, pp. 50-56.

Privacy Rights Clearing House. (2008) *The Privacy Implications of Cloud Computing*. Available at http://www.privacy rights.org/ar/cloud-computing.htm.

Tavani, H. T. (2007) "Philosophical Theories of Privacy: Implications for an Adequate Online Privacy Policy," *Metaphilosophy*, Vol. 38, No. 1, pp. 1-22.
Tavani, H. T. (2008a) "Informational Privacy: Concepts, Theories, and Controversies." In Himma, K. E. and Tavani, H. T., eds. *The Handbook of Information and Computer Ethics*. Hoboken, NJ: John Wiley and Sons, pp. 131-164.
Tavani, H. T. (2008b) "Floridi's Ontological Theory of Informational Privacy: Some Implications and Challenges," *Ethics and Information Technology*, Vol. 10, Nos. 2-3, pp. 155-166.
Zeng, K. and Cavoukian A. (2010) *Modeling Cloud Computing Architecture without Compromising Privacy: A Privacy by Design Approach*. Available at: www.privacybydesign.ca.

Biographies

Frances S. Grodzinsky is a professor of Computer Science and Information Technology at Sacred Heart University where she is co-chair of the Hersher Institute of Ethics. She is also a Visiting Scholar, Research Center on Computer Ethics and Social Responsibility, Southern Connecticut State University, New Haven, CT .

Herman T. Tavani is Professor of Philosophy at Rivier College, President of the International Society for Ethics and Information Technology (INSEIT), and a visiting scholar at the Harvard School of Public Health. He is the author, editor, or co-editor of five books on ethical aspects of information technology.