

Security and Privacy Challenges in Cloud Computing Environments

Cloud computing is an evolving paradigm with tremendous momentum, but its unique aspects exacerbate security and privacy challenges. This article explores the roadblocks and solutions to providing a trustworthy cloud computing environment.



HASSAN
TAKABI AND
JAMES B.D.
JOSHI
*University of
Pittsburgh*

GAIL-JOON
AHN
*Arizona State
University*

Cloud computing has generated significant interest in both academia and industry, but it's still an evolving paradigm. Essentially, it aims to consolidate the economic utility model with the evolutionary development of many existing approaches and computing technologies, including distributed services, applications, and information infrastructures consisting of pools of computers, networks, and storage resources. Confusion exists in IT communities about how a cloud differs from existing models and how these differences affect its adoption. Some see a cloud as a novel technical revolution, while others consider it a natural evolution of technology, economy, and culture.¹

Nevertheless, cloud computing is an important paradigm, with the potential to significantly reduce costs through optimization and increased operating and economic efficiencies.^{1,2} Furthermore, cloud computing could significantly enhance collaboration, agility, and scale, thus enabling a truly global computing model over the Internet infrastructure. However, without appropriate security and privacy solutions designed for clouds, this potentially revolutionizing computing paradigm could become a huge failure. Several surveys of potential cloud adopters indicate that security and privacy is the primary concern hindering its adoption.³

This article illustrates the unique issues of cloud computing that exacerbate security and privacy challenges in clouds.⁴ We also discuss various approaches to address these challenges and explore the future work needed to provide a trustworthy cloud computing environment.

Cloud Computing: Definition and Features

Although several researchers have tried to define cloud computing, no single, agreed-upon definition exists yet. The US National Institute of Standards and Technology (NIST, <http://csrc.nist.gov>) defines it as follows:

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

To understand the importance of cloud computing and its adoption, we must understand its principal characteristics, its delivery and deployment models, how customers use these services, and how to safeguard them. The five key characteristics of cloud computing include on-demand self-service, ubiquitous network access, location-independent resource pooling, rapid elasticity, and measured service, all of which are geared toward using clouds seamlessly and transparently. *Rapid elasticity* lets us quickly scale up (or down) resources. *Measured services* are primarily derived from business model properties and indicate that cloud service providers control and optimize the use of computing resources through automated resource allocation, load balancing, and metering tools.^{1,2}

Applications running on or being developed for cloud computing platforms pose various security and privacy challenges depending on the underlying delivery and deployment models.

The three key cloud delivery models are software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS). In IaaS, the cloud provider supplies a set of virtualized infrastructural components such as virtual machines (VMs) and storage on which customers can build and run applications. The application will eventually reside on the VM and the virtual operating system. Issues such as trusting the VM image, hardening hosts, and securing inter-host communication are critical areas in IaaS.

PaaS enables programming environments to access and utilize additional application building blocks. Such programming environments have a visible impact on the application architecture, such as constraints on which services the application can request from an OS. For example, a PaaS environment might limit access to well-defined parts of the file system, thus requiring a fine-grained authorization service.

Finally, in SaaS, the cloud providers enable and provide application software as on-demand services. Because clients acquire and use software components from different providers, crucial issues include securely composing them and ensuring that information handled by these composed services is well protected.

Cloud deployment models include public, private, community, and hybrid clouds. Public clouds are external or publicly available cloud environments that are accessible to multiple tenants, whereas private clouds are typically tailored environments with dedicated virtualized resources for particular organizations. Similarly, community clouds are tailored for particular groups of customers.

Unique Security and Privacy Implications in Cloud Computing

Understanding the security and privacy risks in cloud computing and developing efficient and effective solutions are critical for its success. Although clouds allow customers to avoid start-up costs, reduce operating costs, and increase their agility by immediately acquiring services and infrastructural resources when needed, their unique architectural features also raise various security and privacy concerns.

Outsourcing Data and Applications

Cloud computing provides access to data, but the challenge is to ensure that only authorized entities can gain access to it. When we use cloud environments, we rely on third parties to make decisions about our data and platforms in ways never seen before in computing. It's critical to have appropriate mechanisms to prevent cloud providers from using customers' data in a way

that hasn't been agreed upon. It seems unlikely that any technical means could completely prevent cloud providers from abusing customer data in all cases, so we need a combination of technical and nontechnical means to achieve this. Clients need to have significant trust in their provider's technical competence and economic stability.

Extensibility and Shared Responsibility

Cloud providers and customers must share the responsibility for security and privacy in cloud computing environments, but sharing levels will differ for different delivery models, which in turn affect cloud extensibility:

- In SaaS, providers typically enable services with a large number of integrated features, resulting in less extensibility for customers. Providers are more responsible for the security and privacy of application services, more so in public than private clouds where the client organization might have stringent security requirements and provide the needed enforcement services. Private clouds could also demand more extensibility to accommodate customized requirements.
- In PaaS, the goal is to enable developers to build their own applications on top of the platforms provided. Thus, customers are primarily responsible for protecting the applications they build and run on the platforms. Providers are then responsible for isolating the customers' applications and workspaces from one another.
- IaaS is the most extensible delivery model and provides few, if any, application-like features. It's expected that the consumers secure the operating systems, applications, and content. The cloud provider still must provide some basic, low-level data protection capabilities.

Multi-tenancy is another feature unique to clouds, especially in public clouds. Essentially, it allows cloud providers to manage resource utilization more efficiently by partitioning a virtualized, shared infrastructure among various customers. From a customer's perspective, the notion of using a shared infrastructure could be a huge concern. However, the level of resource sharing and available protection mechanisms can make a big difference. For example, to isolate multiple tenants' data, Salesforce.com employs a query rewriter at the database level, whereas Amazon uses hypervisors at the hardware level. Providers must account for issues such as access policies, application deployment, and data access and protection to provide a secure, multi-tenant environment.

Service-Level Agreements

The on-demand service or utility-based economic

model necessitates the use of well-established service-level agreements.² An SLA is a part of a service contract between the consumer and provider that formally defines the level of service. It records a common un-

Cloud computing environments are multidomain environments in which each domain can use different security, privacy, and trust requirements.

derstanding about services, priorities, responsibilities, guarantees, and warranties. In cloud computing, SLAs are necessary to control the use of computing resources.

Therefore, the main issue for cloud computing is to build a new layer to support a contract negotiation phase between service providers and consumers and to monitor contract enforcement. Unfortunately, security, privacy, and trust are inherently non-quantitative and difficult to bargain, but there should still be ways to assure customers that services are provided according to what a service provider claims in the contract. The dynamic nature of the cloud necessitates continuous monitoring of attributes to enforce SLAs. Consumers might not completely trust measurements provided solely by a service provider, which might require agreed-upon third-party mediators to measure the SLA's critical service parameters and report violations.

Virtualization and Hypervisors

Virtualization is an important enabling technology that helps abstract infrastructure and resources to be made available to clients as isolated VMs.¹ A hypervisor or VM monitor is a piece of platform-virtualization software that lets multiple operating systems run on a host computer concurrently. Although this provides a means to generate virtualized resources for sharing, such technology's presence also increases the attack surface. We need mechanisms to ensure strong isolation, mediated sharing, and secure communications between VMs. This could be done using a flexible access control mechanism that governs the control and sharing capabilities of VMs within a cloud host.⁵

For some applications, it might be important to associate process outputs to specific hardware components because of the need to ensure authenticity of data generated (such as by sensor hardware) or to establish the use of authentic hardware components (for example, to ensure counterfeit components aren't used or for licensing purposes). In networked environments, hardware association could be used to establish traceback. However, virtualization might make such association difficult to establish.

Heterogeneity

Heterogeneity in clouds comes in different forms. First, cloud providers use various hardware and software resources to build cloud environments. To some extent, resource virtualization achieves high-level system homogeneity, but the same infrastructure being used to support different tenants with different protection and system requirements can generate difficulties. There's also a potential issue with vertical heterogeneity of cloud services. For instance, a client might subscribe to an IaaS from one provider, couple it with a PaaS from another cloud provider, and acquire various pieces of SaaS from a third cloud vendor. The assumptions that each of these cloud providers make in building the services can severely affect the emergent trust and security properties. For example, providers might have used the lowest denominator or generic assumptions, which might be inappropriate for the composed environments. Furthermore, heterogeneity exists in the level of security treatment each component provides, thus generating integration challenges.

In a multi-tenant environment, the protection requirements for each tenant might differ, which can make a multi-tenant cloud a single point of compromise. In addition, each tenant could have different trust relations with the provider—and some tenants could actually be malicious attackers themselves—thus generating complex trust issues.

Compliance and Regulations

As we already mentioned, ensuring that cloud providers and clients comply with established SLAs and existing regulatory requirements such as Sarbanes-Oxley and HIPAA is a key issue.³ In existing environments, organizations typically have well-established processes for compliance monitoring and enforcement. Cloud computing also promises to be a global phenomenon by potentially harvesting widely dispersed computing and infrastructural resources, thus making cloud services accessible from anywhere and at anytime. This can potentially raise multiple jurisdiction issues with regard to protection requirements and enforcement mechanisms.

Security and Privacy Challenges

Cloud computing environments are multidomain environments in which each domain can use different security, privacy, and trust requirements and potentially employ various mechanisms, interfaces, and semantics. Such domains could represent individually enabled services or other infrastructural or application components. Service-oriented architectures are naturally relevant technology to facilitate such multidomain formation through service composition and orchestration.² It is important to leverage existing research on multidomain policy integration and the

secure-service composition to build a comprehensive policy-based management framework in cloud computing environments.

Authentication and Identity Management

By using cloud services, users can easily access their personal information and make it available to various services across the Internet. An identity management (IDM) mechanism can help authenticate users and services based on credentials and characteristics.⁶

A key issue concerning IDM in clouds is interoperability drawbacks that could result from using different identity tokens and identity negotiation protocols. Existing password-based authentication has an inherited limitation and poses significant risks. An IDM system should be able to protect private and sensitive information related to users and processes. How multi-tenant cloud environments can affect the privacy of identity information isn't yet well understood. In addition, the multi-jurisdiction issue can complicate protection measures.³ While users interact with a front-end service, this service might need to ensure that their identity is protected from other services with which it interacts.^{6,7} In multi-tenant cloud environments, providers must segregate customer identity and authentication information. Authentication and IDM components should also be easily integrated with other security components.

Access Control and Accounting

Heterogeneity and diversity of services, as well as the domains' diverse access requirements in cloud computing environments, demand fine-grained access control policies. In particular, access control services should be flexible enough to capture dynamic, context, or attribute- or credential-based access requirements and to enforce the principle of least privilege. Such access control services might need to integrate privacy-protection requirements expressed through complex rules.

It's important that the access control system employed in clouds is easily managed and its privilege distribution is administered efficiently. We must also ensure that cloud delivery models provide generic access control interfaces for proper interoperability, which demands a policy-neutral access control specification and enforcement framework that can be used to address cross-domain access issues.⁸ The access control models should also be able to capture relevant aspects of SLAs. The utility model of clouds demands proper accounting of user and service activities that generates privacy issues because customers might not want to let a provider maintain such detailed accounting records other than for billing purposes. The outsourcing and multi-tenancy aspects of clouds could accelerate customers' fears about accounting logs.

Hence, utilizing a privacy-aware framework for access control and accounting services is crucial, and it should be easily amenable to compliance checking.

Trust Management and Policy Integration

Although multiple service providers coexist in clouds and collaborate to provide various services, they might have different security approaches and privacy mechanisms, so we must address heterogeneity among their policies.^{2,9,10} Cloud service providers might need to compose multiple services to enable bigger application services. Therefore, mechanisms are necessary to ensure that such a dynamic collaboration is handled securely and that security breaches are effectively monitored during the interoperation process. Existing literature has shown that even though individual domain policies are verified, security violations can easily occur during integration.¹⁰ Hence, providers should carefully manage access control policies to ensure that policy integration doesn't lead to any security breaches.

In cloud computing environments, the interactions between different service domains driven by service requirements can be dynamic, transient, and intensive. Thus, a trust framework should be developed to allow for efficiently capturing a generic set of parameters required for establishing trust and to manage evolving trust and interaction/sharing requirements.^{10,11} The cloud's policy integration tasks should be able to address challenges such as semantic heterogeneity, secure interoperability, and policy-evolution management. Furthermore, customers' behaviors can evolve rapidly, thereby affecting established trust values. This suggests a need for an integrated, trust-based, secure interoperation framework that helps establish, negotiate, and maintain trust to adaptively support policy integration.^{10,11}

Secure-Service Management

In cloud computing environments, cloud service providers and service integrators compose services for their customers. The service integrator provides a platform that lets independent service providers orchestrate and interwork services and cooperatively provide additional services that meet customers' protection requirements. Although many cloud service providers use the Web Services Description Language (WSDL), the traditional WSDL can't fully meet the requirements of cloud computing services description. In clouds, issues such as quality of service, price, and SLAs are critical in service search and composition. These issues must be addressed to describe services and introduce their features, find the best interoperable options, integrate them without violating the service owner's policies, and ensure that SLAs are satisfied. In essence, an automatic and systematic service

provisioning and composition framework that considers security and privacy issues is crucial.

Privacy and Data Protection

Privacy is a core issue in all the challenges we've discussed so far, including the need to protect identity information, policy components during integration, and transaction histories. Many organizations aren't comfortable storing their data and applications on systems that reside outside of their on-premise datacenters.⁵ This might be the single greatest fear of cloud clients. By migrating workloads to a shared infrastructure, customers' private information faces increased risk of potential unauthorized access and exposure. Cloud service providers must assure their customers and provide a high degree of transparency into their operations and privacy assurance. Privacy-protection mechanisms must be embedded in all security solutions.

In a related issue, it's becoming important to know who created a piece of data, who modified it and how, and so on. Provenance information could be used for various purposes such as traceback, auditing, and history-based access control. Balancing between data provenance and privacy is a significant challenge in clouds where physical perimeters are abandoned.

Organizational Security Management

Existing security management and information security life-cycle models significantly change when enterprises adopt cloud computing. In particular, shared governance can become a significant issue if not properly addressed. Despite the potential benefits of using clouds, it might mean less coordination among different communities of interest within client organizations. Dependence on external entities can also raise fears about timely responses to security incidents and implementing systematic business continuity and disaster recovery plans. Similarly, risk and cost-benefit issues will need to involve external parties. Customers consequently need to consider newer risks introduced by a perimeter-less environment, such as data leakage within multi-tenant clouds and resiliency issues such as their provider's economic instability and local disasters.

Similarly, the possibility of an insider threat is significantly extended when outsourcing data and processes to clouds. Within multi-tenant environments, one tenant could be a highly targeted attack victim, which could significantly affect the other tenants. Existing life-cycle models, risk analysis and management processes, penetration testing, and service attestation must be reevaluated to ensure that clients can enjoy the potential benefits of clouds.

The information security area has faced significant problems in establishing appropriate security metrics for consistent and realistic measurements that help risk assessment. We must reevaluate best practices and develop

standards to ensure the deployment and adoption of secure clouds. These issues necessitate a well-structured cyberinsurance industry, but the global nature of cloud computing makes this prospect extremely complex.

Security and Privacy Approaches

Here, we discuss various approaches to cope with the previously mentioned challenges, existing solutions, and the work needed to provide a trustworthy cloud computing environment. The approaches address security and privacy requirements of cloud service providers, service integrators, and cloud environments in general.

Authentication and Identity Management

User-centric IDM has recently received attention for handling private and critical identity attributes.⁷ In this approach, identifiers or attributes help identify and define a user. Such an approach lets users control their digital identities and takes away the complexity of IDM from the enterprises, thereby allowing them to focus on their own functions. Because users can access the cloud from various places such as home, office, school, or other public places, they must be able to export their digital identities and securely transfer them to various computers. User-centric IDM also implies that the system properly maintains the semantics of the context of users' identity information, sometimes constraining or relaxing them to best respond to a user request in a given situation.

Researchers are currently pursuing other federated IDM solutions that might benefit cloud environments.^{6,7} IDM services in the cloud should be able to be integrated with an enterprise's existing IDM framework.^{1,2} In some cases, it's important to have privacy-preserving protocols to verify various identity attributes by using, for example, zero-knowledge proof-based techniques.⁶ These techniques, which use pseudonyms and accommodate multiple identities to protect users' privacy, can further help build a desired user-centric federated IDM for clouds. IDM solutions can also be extended with delegation capabilities to address identification and authentication issues in composed services.

Access Control Needs

Among the many methods proposed so far, role-based access control (RBAC) has been widely accepted as the most promising model because of its simplicity, flexibility in capturing dynamic requirements, and support for the principle of least privilege and efficient privilege management.⁸ Furthermore, RBAC is policy neutral, can capture various policy requirements, and is best suited for policy-integration needs. Due to the highly dynamic nature of clouds, obligations and conditions are crucial decision factors for richer and finer controls on usage of resources provided by the cloud.

Recent RBAC extensions—such as credential-based RBAC, generalized temporal RBAC (GTRBAC),⁸ and location-based RBAC models—provide necessary modeling constructs and capabilities to capture context-based fine-grained access requirements. In clouds, service providers usually do not know their users in advance, so it is difficult to assign users directly to roles in access control policies. Therefore, using credential- or attribute-based policies might enhance this capability. However, little work exists in employing RBAC and extensions within intensely service-oriented environments such as clouds.

Secure Interoperation

Several recent works have focused on multidomain access control policies and policy integration issues, which can be adopted to build a comprehensive policy management framework in clouds.^{2,6} Researchers have addressed secure interoperation and policy engineering mechanisms to integrate access policies of different domains and define global access policies.^{9,10} A centralized approach creates a global policy that mediates all accesses and is appropriate for a cloud application that consists of various services with different requirements and is more or less fixed. In a more dynamic environment, the domains are transient and might need to interact for a specific purpose, making centralized approaches inappropriate and demanding decentralized approaches. We also need specification frameworks to ensure that the cross-domain accesses are properly specified, verified, and enforced. Security Assertion Markup Language (SAML), Extensible Access Control Markup Language (XACML), and Web services standards are viable solutions toward this.¹⁰

Policy engineering mechanisms can help define global policies to accommodate all collaborators' requirements. Emerging role-mining techniques can support this. Role mining uses the existing system configuration data to define roles. It first considers the existing users' permissions and aggregates them into roles.⁸ In a cloud, users acquire different roles from different domains based on the services they need. To define global policies, we can utilize these RBAC systems' configurations from different domains to define global roles and policies. Each global role can include roles from different domains that have been assigned to the same groups of users.

Several new approaches have been proposed for role engineering that could be adopted in clouds for policy engineering purposes. Changes to the existing role set might cause disruptions to the organization and prevent it from functioning properly. Therefore, role mining should look for a set of roles as close as possible to both the existing and optimal sets of roles. One possible solution is the StateMiner approach,¹² which

introduces new measures for optimality and presents a heuristic solution to find an RBAC state with the smallest structural complexity and that's as similar as possible to both the existing and optimal state.

Secure-Service Provisioning and Composition

To optimize resource utilization, cloud service providers often use virtualization technologies that separate application services from infrastructure. In the cloud, service providers and service integrators need to collaborate to provide newly composed services to customers. This sort of activity requires automatic service provisioning and composition frameworks that allow cloud service providers and service integrators to describe services with unified standards to introduce their functionalities, discover existing interoperable services, and securely integrate them to provide services. Such frameworks must include a declarative language to describe services, features, and mechanisms to provision and compose appropriate services.

The Open Services Gateway Initiative (OSGi) service platform provides an open, common architecture for service providers, developers, software vendors, gateway operators, and equipment vendors to cooperatively develop, deploy, and manage services.¹³ Researchers have developed ways to configure and map the OSGi authorization mechanism to RBAC.¹³ Declarative OWL-based language can be used to provide a service definition manifest, including a list of distinct component types that make up the service, functional requirements, component grouping and topology instructions, and so on. OSGi can also be adopted to develop an agent-based collaboration system for automatic service provisioning.

The challenges of such collaboration systems include dynamic access control to resources shared by agents and controlling collaborative actions that are geared toward a collaboration goal.

Trust Management Framework

To facilitate policy integration between various domains in cloud environments, a trust-based framework that facilitates automated trust-based policy integration is essential. In doing so, we must answer several questions: How do we establish trust and determine access

How do we manage and maintain dynamically changing trust values and adapt access requirements as trust evolves?

mapping to satisfy interdomain access requirements, and how do we manage and maintain dynamically changing trust values and adapt access requirements

as trust evolves? Existing trust negotiation mechanisms primarily focus on credential exchange^{9,10} and don't address the more challenging need of integrating requirements-driven trust negotiation techniques with fine-grained access control mechanisms.^{10,11} One possible approach is to develop a comprehensive trust-based policy integration framework that facilitates policy integration and evolution based on interdomain- and service-access requirements.

Because service composition dynamics in the cloud can be complex, trust and access control frameworks should include delegation primitives.¹¹ Existing work related to delegation, including role-based delegation, has focused on issues related to delegation of privileges among subjects and various levels of controls with regard to privilege propagation and revocation. Efficient cryptographic mechanisms for trust delegation involve complex trust-chain verification and revocation issues, raising significant key management issues. These approaches must be incorporated in service composition frameworks.⁹

Data-Centric Security and Privacy

Data in the cloud typically resides in a shared environment, but the data owner should have full control over who has the right to use the data and what they are allowed to do with it once they gain access. To provide this data control in the cloud, a standard-based heterogeneous data-centric security approach is an essential element that shifts data protection from systems and applications. In this approach, documents must be self-describing and defending regardless of their environments.

Cryptographic approaches and usage policy rules must be considered. When someone wants to access data, the system should check its policy rules and reveal it only if the policies are satisfied. Existing cryptographic techniques can be utilized for data security, but privacy protection and outsourced computation need significant attention—both are relatively new research directions. Data provenance issues have just begun to be addressed in the literature. In some cases, information related to a particular hardware component (storage, processing, or communication) must be associated with a piece of data.

Managing Semantic Heterogeneity

One key aspect of complex cloud computing environments is semantic heterogeneity among policies. Researchers have given little attention to automatic detection of semantic conflicts among different service providers' policies. Although XML has been adopted as the preferred language for information sharing, research has found it inadequate for describing information semantics.⁶ RDF, on the other hand, provides a facility for describing semantics by supporting ele-

ment attributes and properties description.¹⁰ Although we can capture semantics using RDF, representing relations between the various concepts that the elements represent is essential for facilitating semantic integration of policy information in interacting domains.

Use of an ontology is the most promising approach to addressing the semantic heterogeneity issue.¹⁴ To support ontology development, we can use both XML Schema and Resource Description Framework Schema (RDFS) to accommodate the domain-specific concepts.⁸ However, although RDF is based on XML syntax and OWL is based on the RDFS representation of concepts, neither of these technologies is likely to completely subsume the lower technology in clouds. An OWL-based framework is desirable to support semantic heterogeneity management across multiple providers within a cloud. For such a framework, a system-driven policy framework to facilitate managing security policies in heterogeneous environments and a policy enforcement architecture are essential.^{14,15} Several inference engines are available for inferring policy semantics.

Although security and privacy services in the cloud can be fine-tuned and managed by experienced groups that can potentially provide efficient security management and threat assessment services, the issues we've discussed here show that existing security and privacy solutions must be critically reevaluated with regard to their appropriateness for clouds. Many enhancements in existing solutions as well as more mature and newer solutions are urgently needed to ensure that cloud computing benefits are fully realized as its adoption accelerates. Cloud computing is still in its infancy, and how the security and privacy landscape changes will impact its successful, widespread adoption. □

Acknowledgments

The work of Hassan Takabi and James Joshi is partially supported by US National Science Foundation grants NSF-IIS-0545912 and NSF-CCF-0720737. The work of Gail-Joon Ahn is partially supported by US National Science Foundation grants (NSF-IIS-0900970 and NSF-CNS-0831360) and Department of Energy grants (DE-SC0004308 and DE-FG02-03ER25565).

References

1. Cloud Security Alliance, "Security Guidance for Critical Areas of Focus in Cloud Computing V2.1," <http://www.cloudsecurityalliance.org/csaguide.pdf>.
2. D. Catteddu and G. Hogben, "Cloud Computing: Benefits, Risks and Recommendations for Information Security," ENISA, 2009; www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport.
3. P.J. Bruening and B.C. Treacy, "Cloud Computing: Pri-

- vacuity, Security Challenges,” Bureau of Nat’l Affairs, 2009; www.hunton.com/files/tbl_s47Details/FileUpload265/2488/CloudComputing_Bruening-Treacy.pdf.
4. H. Takabi, J.B.D. Joshi, and G.-J. Ahn, “SecureCloud: Towards a Comprehensive Security Framework for Cloud Computing Environments,” *Proc. 1st IEEE Int’l Workshop Emerging Applications for Cloud Computing (CloudApp 2010)*, IEEE CS Press, 2010, pp. 393–398.
 5. Y. Chen, V. Paxson, and R.H. Katz, “What’s New About Cloud Computing Security?” tech. report UCB/EECS-2010-5, EECS Dept., Univ. of California, Berkeley, 2010; www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-5.html.
 6. E. Bertino, F. Paci, and R. Ferrini, “Privacy-Preserving Digital Identity Management for Cloud Computing,” *IEEE Computer Society Data Engineering Bulletin*, Mar. 2009, pp. 1–4.
 7. M. Ko, G.-J. Ahn, and M. Shehab “Privacy-Enhanced User-Centric Identity Management,” *Proc. IEEE Int’l Conf. Communications*, IEEE Press, 2009, pp. 998–1002.
 8. J. Joshi et al., “Access Control Language for Multi-domain Environments,” *IEEE Internet Computing*, vol. 8, no. 6, 2004, pp. 40–50.
 9. M. Blaze et al., “Dynamic Trust Management,” *Computer*, vol. 42, no. 2, 2009, pp. 44–52.
 10. Y. Zhang and J. Joshi, “Access Control and Trust Management for Emerging Multidomain Environments,” *Annals of Emerging Research in Information Assurance, Security and Privacy Services*, S. Upadhyaya and R.O. Rao, eds., Emerald Group Publishing, 2009, pp. 421–452.
 11. D. Shin and G.-J. Ahn, “Role-Based Privilege and Trust Management,” *Computer Systems Science & Eng. J.*, vol. 20, no. 6, 2005, pp. 401–410.
 12. H. Takabi and J. Joshi, “StateMiner: An Efficient Similarity-Based Approach for Optimal Mining of Role Hierarchy,” *Proc. 15th ACM Symp. Access Control Models and Technologies*, ACM Press, 2010, pp. 55–64.
 13. G.-J. Ahn, H. Hu, and J. Jin, “Security-Enhanced OSGi Service Environments,” *IEEE Trans. Systems, Man, and Cybernetics-Part C: Applications and Reviews*, vol. 39, no. 5, 2009, pp. 562–571.
 14. L. Teo and G.-J. Ahn, “Managing Heterogeneous Network Environments Using an Extensible Policy Framework,” *Proc. Asian ACM Symp. Information, Computer and Communications Security*, ACM Press, 2007, pp. 362–364.
 15. H. Takabi et al., “An Architecture for Specification and Enforcement of Temporal Access Control Constraints using OWL,” *Proc. 2009 ACM Workshop on Secure Web Services*, ACM Press, 2009, pp. 21–28.

Hassan Takabi is a PhD student in the School of Information Sciences and a member of the Laboratory of Education and Research on Security Assured Information Systems (LERSAIS) at the University of Pittsburgh. His research interests include access control models; trust management; privacy and Web


security; usable privacy and security; and security, privacy, and trust issues in cloud computing environments. Takabi has an MS in information technology from Sharif University of Technology, Iran. He is student member of IEEE and the ACM. Contact him at hatakabi@sis.pitt.edu.

James B.D. Joshi is an associate professor and the director of the Laboratory for Education and Research on Security Assured Information Systems (LERSAIS) in the School of Information Sciences at the University of Pittsburgh. His research interests include role-based access control, trust management, and secure interoperability. Joshi has a PhD in computer engineering from Purdue University. He is a member of IEEE and the ACM. Contact him at jjoshi@sis.pitt.edu.


Gail-Joon Ahn is an associate professor and the director of the Security Engineering for Future Computing (SEFCOM) Laboratory in the School of Computing, Informatics, and Decision Systems Engineering at Arizona State University. His research interests include information and systems security, vulnerability and risk management, access control, and security architecture for distributed systems. Ahn has a PhD in information technology from George Mason University. He is a recipient of the US Department of Energy Career Award and the Educator of the Year Award from the Federal Information Systems Security Educators Association (FISSEA). Ahn is a senior member of IEEE and the ACM. Contact him at gahn@asu.edu.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.



LISTEN TO GRADY BOOCH “On Architecture”

podcast available at  <http://computingnow.computer.org>