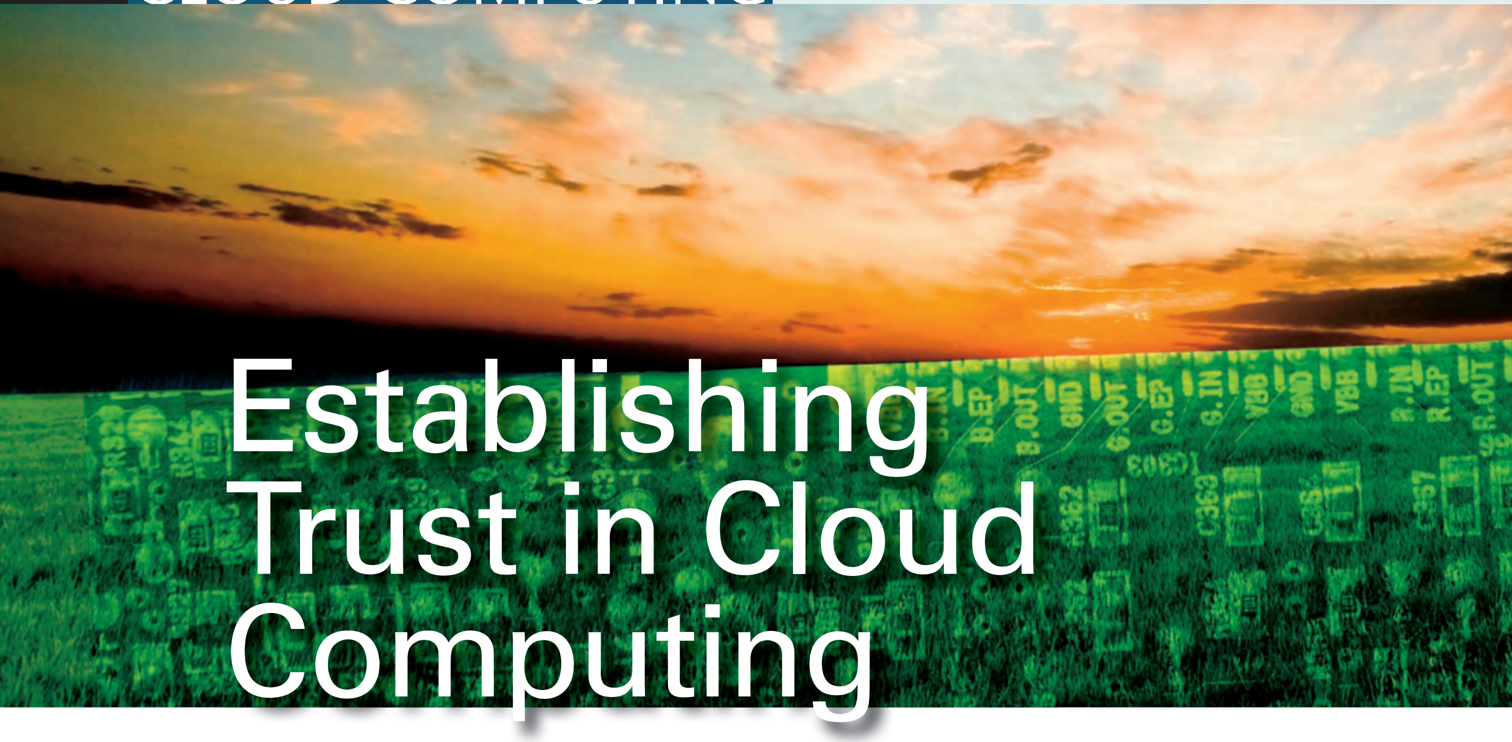

Establishing Trust in Cloud Computing

Khaled M. Khan and Qutaibah Malluhi, *Qatar University*

How can cloud providers earn their customers' trust when a third party is processing sensitive data in a remote machine located in various countries? Emerging technologies can help address the challenges of trust in cloud computing.

Cloud computing provides many opportunities for enterprises by offering a range of computing services. In today's competitive environment, the service dynamism, elasticity, and choices offered by this highly scalable technology are too attractive for enterprises to ignore. These opportunities, however, don't come without challenges.

Cloud computing has opened up a new frontier of challenges by introducing a different type of trust scenario. Today, the problem of trusting cloud computing is a paramount concern for most enterprises. It's not that the enterprises don't trust the cloud providers' intentions; rather, they question cloud computing's capabilities.

Yet the challenges of trusting cloud computing don't lie entirely in the technology itself. The dearth of customer confidence also stems from a lack of transparency, a loss of control over data assets, and unclear security assurances.

Unfortunately, the adoption of cloud computing came before the appropriate technologies appeared to tackle the accompanying challenges of trust. This gap between adoption and innovation is so wide that cloud computing consumers don't fully trust this new way of computing. To close this gap, we need to understand the trust issues associated with cloud computing from both a technology and business perspective. Then we'll be able to determine which emerging technologies could best address these issues.

What Is Trust?

Broadly speaking, trust means an act of faith; confidence and reliance in something that's expected to behave or deliver as promised.[1,2] It's a belief in the competence and expertise of others, such that you feel you can reasonably rely on them to care for your valuable assets.[3]

We trust a system less if it gives us insufficient information about its expertise. Mere claims such as "secure cloud" or "trust me" don't help much to boost the trust level of consumers

Published by the IEEE Computer Society

1520-9202/10/$26.00 © 2010 IEEE

unless sufficient information is presented with the services.[4]

Control

Control is another important issue in trust. We trust a system less when we don't have much control over our assets.

For example, when we withdraw money from an ATM, we trust that the machine will give us the exact amount because it's under our control—we receive ("control") the money. When we make a deposit using the same ATM, we usually don't have the same level of trust because we're losing control over our money—we don't know what happens after the ATM consumes it. Similarly, the more control consumers have over the data consigned to a cloud, the more they'll trust the system.

Ownership

We can also see a variation of trust, depending on the ownership of data assets. Alice might trust an online payment system when she pays with her credit card, but she might have less trust in the same system when using her client's card, because preserving her client's interest is one of her business objectives.

Similarly, when enterprises consign their data to cloud computing (data representing both their own interests and those of their clients), it creates two folds of a complex trust relationship. First, the enterprise must trust the cloud provider. Second, the enterprise must ascertain that its clients have enough reason to trust the same provider.[5]

Prevention

Contractual relationships are often used to establish trust. In a typical business environment, an organization is compensated if the service isn't delivered as expected. Cloud providers similarly use service-level agreements (SLAs) to boost consumers' trust. Unfortunately, these might not help in cloud computing.

Trust in cloud computing is related more to preventing a trust violation than to guaranteeing compensation should a violation occur. For most enterprises, a security breach of data is irreparable—no amount of money can guarantee to restore the lost data or the enterprise's reputation. The cloud computing trust model thus should focus more on preventing failure than on post-failure compensation.

Security

Security plays a central role in preventing service failures and cultivating trust in cloud computing. In particular, cloud service providers need to secure the virtual environment, which enables them to run services for multiple clients and offer separate services for different clients.

In the context of virtualization, the key security issues include identity management, data leakage (caused by multiple tenants sharing physical resources), access control, virtual machine (VM) protection, persistent client-data security, and the prevention of cross-VM side-channel attacks.

Vendors and research communities are working to address these cloud-specific security concerns. For example, Intel's SOA Expressway claims to enforce persistent security on client data by extending the perimeter of enterprises into the cloud provider (so the enterprises retain a certain amount of control over the computing tasks and data consigned to cloud).[6] The VMsafe API provides VM security protection at the host level.[7] Its VMotion capabilities can dynamically move VMs between physical devices as required.

To ensure integrity and authenticity, and to address access control in a cloud-enabled system, some have proposed using claim-based access control, a security assertion markup language, a security token service, and federated identity approaches.[8] Undoubtedly, these low-level security concerns are important, but to understand the issues related to consumer-level trust, we need to take a closer look at cloud computing.

A Cloud Computing Example

Imagine a company called SoftCom that handles thousands of healthcare-related digital images of its clients. The images are sensitive and should remain private and confidential. SoftCom decides to use CloudX, a public cloud provider located in Boston, for

- image processing—using SoftCom's ImagePro software on a remote application server,
- additional image-processing tasks (filtering and searching) that ImagePro doesn't support but that CloudX's iFilter and iSearch systems can perform, and
- image archiving.

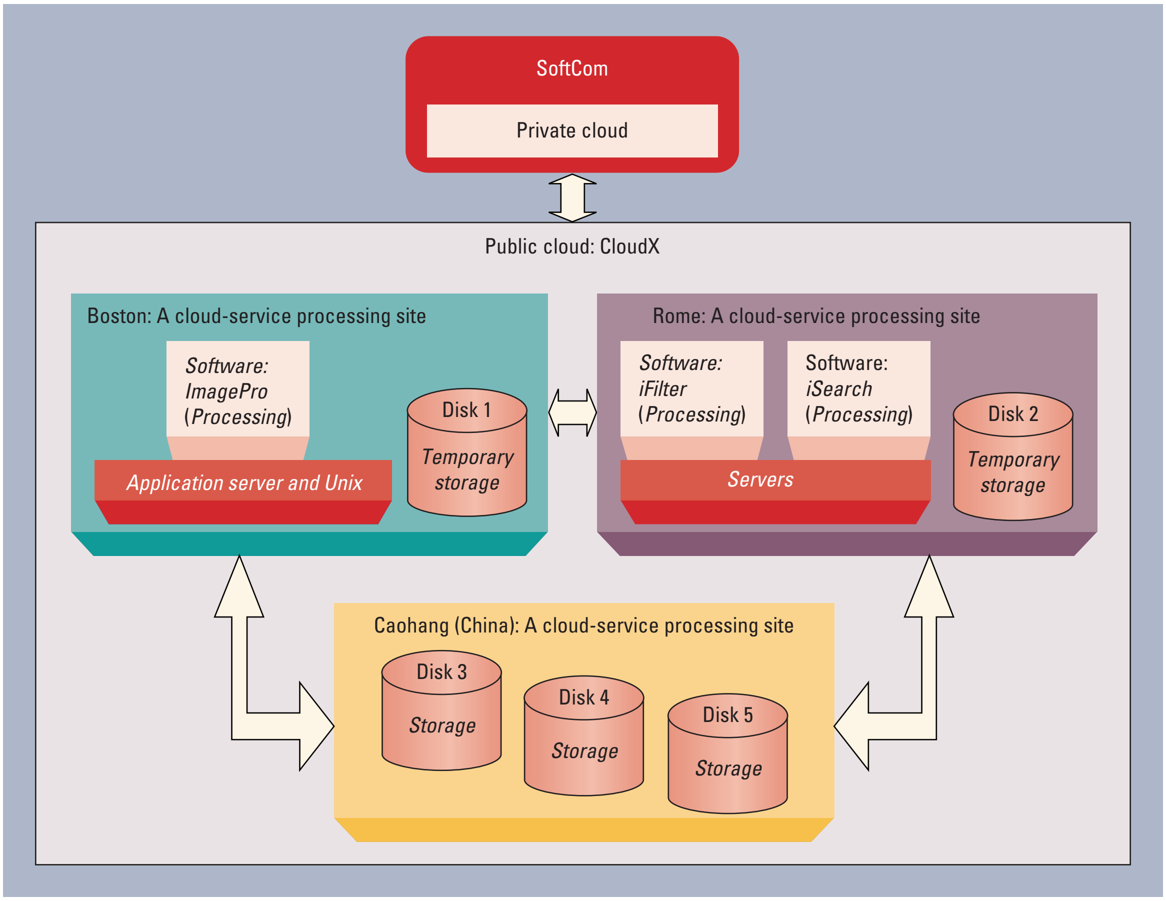
SoftCom
Private cloud
Public cloud: CloudX
Boston: A cloud-service processing site
Software: ImagePro (Processing)
Application server and Unix
Disk 1
Temporary storage
Rome: A cloud-service processing site
Software: iFilter (Processing)
Software: iSearch (Processing)
Servers
Disk 2
Temporary storage
Caohang (China): A cloud-service processing site
Disk 3
Storage
Disk 4
Storage
Disk 5
Storage

Figure 1. A hybrid cloud computing architecture. SoftCom retains a private cloud for sensitive research activities but employs a public cloud for other services.

Note that in a public cloud, an enterprise can offload its computing tasks to the external cloud provider. In a private cloud, the computing services and resources remain within the perimeters of the enterprise's private network, so the enterprise retains control of the computing tasks.[6] A hybrid cloud is a combination of private and public computing.

In this example, SoftCom uses the hybrid model. It retains a private cloud for sensitive research activities to develop new image-processing and data-mining algorithms. Yet it also uses CloudX for other services.

At the CloudX site in Boston, ImagePro—hosted on an application server running in a Unix environment—processes images and stores them temporarily on a disk (Disk 1). CloudX then transmits the images to another cloud site located in Rome for additional processing by iFilter and iSearch. Next, it stores the images on another temporary disk (Disk 2). CloudX archives the processed images on Disks 3, 4, and 5, physically located in Caohang, Shanghai. Its cloud infrastructure division manages these archives.

This scenario suggests that SoftCom consumes three types of services (see Figure 1): platform as a service (PaaS), software as a service (SaaS), and infrastructure as a service (IaaS).

In PaaS, consumers can build and deploy their applications on the cloud provider's platform as needed. In this case, SoftCom uses CloudX's application server and Unix platforms (in Boston) to deploy its ImagePro software.

In SaaS, consumers use software services provided by cloud providers, such as email, payroll processing, and invoice generation. In this case, SoftCom uses CloudX's iFilter and iSearch systems.

IaaS provides SoftCom with computing power and disk storage via CloudX's virtual environments. SoftCom can access the virtual servers and storage provisioned on CloudX's physical infrastructure.

The Challenges of Trust

Figure 1 illustrates how CloudX processes and stores SoftCom's images, transmitting them between various hardware and software devices located in Boston, Rome, and Caohang. This extensive sharing of computing resources from multiple sites includes additional communication links and involves several remote computing sites in the chain of services.

These additional links require SoftCom to entrust its images to devices and systems located in remote locations, managed by others, and regulated by the laws of other countries (Italy and China). Yet SoftCom doesn't know whether the security profiles of those sites are the same as at the site in Boston or whether the regulatory compliances such as the Health Insurance Portability and Accountability Act (HIPAA) hold in all those sites. Although CloudX provides SoftCom with a comprehensive SLA, two major trust-related factors are a concern in this scenario.

Diminishing Control

SoftCom finds that the moment its images leave its perimeter, it doesn't have much control over them or the processes that manipulate them. It doesn't know who can access the images—which are stored on various disks in multiple locations (Boston, Rome, and Caohang) and possibly managed by third-party providers.

In cloud computing, this lack of control over the data and processes triggers the risk of losing data confidentiality, integrity, and availability. Cloud computing virtually requires consumers to relinquish control of running their applications and storing their data.

The degree of lost control over the data and processes depends on the cloud service model. For example, in IaaS and PaaS, the provider usually has complete control of the server, storage facility, and network. It's the same with SaaS, but the provider also controls the applications. Enterprises retain only partial control of their data,[9] which they often find quite alarming.

Lack of Transparency

The consumer's perception is that a cloud is generally less secure than an in-house system,[10] but better transparency could help address this issue.

Data stored in a cloud provider's devices isn't located on a single machine in a single location or country. Rather, the data is stored and processed across the entire virtual layer. There are two issues involved in transparency: one is the physical location of the storage and processing sites, and the other is the security profiles of these sites.

In our example, SoftCom has lost visibility of its applications and storage sites. It should know where its images are processed and stored, because in some countries, the laws might not support SoftCom should a data breach or loss occur. In this highly fluid distributed environment, SoftCom needs to know how its images are protected while being moved within the system or across multiple sites owned by multiple independent software vendors. It should also know what data manipulation and access privileges third-party employees have and if audit trails are available.

Without transparency, SoftCom doesn't know if there's any mismatch between its enterprise security requirements and CloudX's security assurances. SoftCom's clients also need to know where their images are processed and stored and the security assurances of those sites. At the end of day, SoftCom is accountable to its own clients and thus must supply them with sufficient information for trusting CloudX.

Addressing These Issues

To fully trust CloudX, SoftCom needs the following assurances regarding its *control of the data*:

- CloudX will notify SoftCom when an entity accesses its images,
- CloudX and its other sites won't keep unauthorized copies of SoftCom images, and
- CloudX will destroy SoftCom's residue (temporary data, intermediate output, or data that's no longer needed) or outdated images at all the sites that it manages.

SoftCom also needs three additional assurances:

- the software (such as iFilter or iSearch) processing the SoftCom images must be reliable and trustworthy (*control of processes*),
- SoftCom must know where the persistent data storage resides and the processing occurs (*physical location*), and

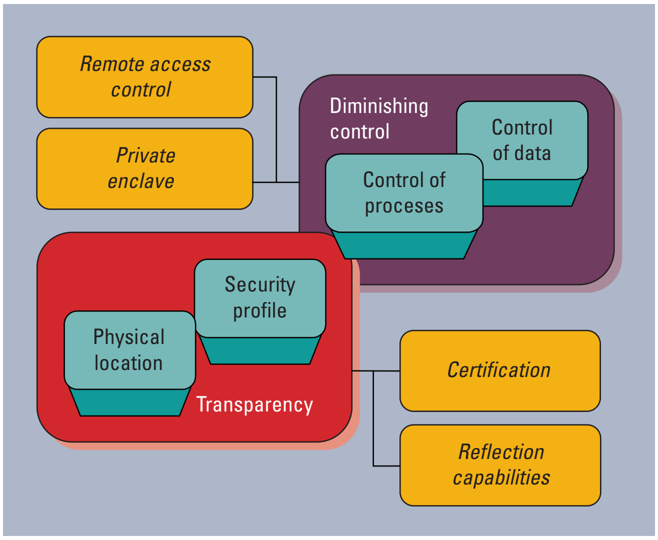

Figure 2. Trust in cloud computing. The figure shows issues related to diminishing control and transparency and the technologies that can address such issues.

- CloudX must make its service-level security properties transparent to SoftCom (*security profiles*).

Although it might seem daunting for cloud providers to offer such assurances, it is both necessary (if they hope to keep building a client base) and possible as the field of security and privacy evolves.

Emerging Technologies

Establishing trust in cloud computing will undoubtedly require identity and data privacy through encryption. It will also require data integrity, which security techniques such as digital signatures and access control can accomplish.

Additionally, advances in cryptography are addressing the issue of confidentiality. For example, although this is still in the research stage, cloud providers can now process encrypted data without decrypting it. They can also use partial encryption to prevent the cloud server from viewing or deciphering partially encrypted data.

Although these could make the system more secure, cultivating trust will require additional capabilities coupled with existing security practices (see Figure 2).

Remote Access Control

Cloud clients need remote access control capabilities, which can give them more jurisdictions over their data, regardless of the cloud provider's physical locations. Also, automatic tools with remote-tracking capabilities could let cloud consumers monitor how much access employees at the cloud service site have to their data. The consumer could then enable and disable data-manipulation commands at the remote sites.

This might change consumers' perception of the cloud as less trustworthy than in-house systems. A remote access tool would give consumers proactive control over their data at the remote location and the ability to better specify and enforce policies. When an employee at the cloud provider site logs out, the consumer could set the browser cache to automatically remove the contents.[11,12] Clients could also receive access logs and audit trails of all the cloud provider users and employees.[13]

Even when data is physically spread out and stored in various remote locations and processed by remote machines and software, the data owner could retain control of these activities using an approach similar to LongArm. The LongArm software package lets users control remote devices (see www.gdc4s.com/documents/LongArm.pdf).

Reflection

Transparency helps clients determine a priori whether a cloud is trustworthy based on profiles and security assurances associated with a service. The reflection mechanisms of a cloud provider's security profile inform consumers about the provider's strengths and weaknesses and reveal how their enterprise security policies would be addressed. Enterprises can then determine whether they need additional security to tackle any vulnerabilities they see in the cloud.

This capability, coupled with an automatic traceability facility, could help consumers determine the physical locations of various nodes in the cloud computing chain, so they'd know where their data is processed and stored.

Certification

In cloud computing's fluid and dynamic environment, ensuring security compliance can be difficult. The cloud is opaque, and the cloud providers can have differing security assurances.

To fully materialize a trusted cloud model, an independent security certification authority could certify cloud services in terms of their security properties and capabilities. The certificate would act as a quality stamp, guaranteeing secure services with a given degree of confidence. It could ensure that the implementation of the service security matched the published security profiles. The certificate could work as a trust model to boost consumers' confidence in cloud services.

Private Enclaves

Cloud computing providers could form a security enclave for their consumers, as is widely practiced in the defense industry. An enclave is a set of computing environments connected by one or more networks that a single authority controls using a common security policy.

Enclaves could provide a set of standard capabilities, such as incident detection and response, boundary defense, and monitoring. They could be specific to an enterprise or to a set of similar services that various enterprises consume.

At the same time, providers could also compartmentalize users' data so that it's not mixed up with other users' data. This would solve the problem of cross-VM side-channel attacks. Cloud providers should also prevent attackers from creating a cloud cartography[14] of the enclave by refusing to disclose the mapping of the physical topology of the cloud computing for a service or user. In an enclave, it's easier to enforce the enterprise's security policy because you're only dealing with the part of the cloud related to the client data or processes, rather than the entire cloud.

Of course, there's no blanket solution to convince consumers that a cloud is fully trustworthy. The importance of trust varies from organization to organization, depending on the data's value. Furthermore, the less trust an enterprise has in the cloud provider, the more it wants to control its data—even the technology. However, it's crucial that consumers and providers change their mindsets.

Trusting cloud computing might differ from trusting other systems, but the goal remains the same—improve business and remain competitive by exploiting the benefits of a new technology. Any new technology must gradually build its reputation for good performance and security, earning users' trust over time. The security problems of some cloud providers—such as Google's widespread service outages in May 2009—have prompted consumers to become more security-aware than ever before. To regain consumers' trust, cloud providers must offer better transparency and more consumer control of data and processes. IT

Any new technology must gradually build its reputation for good performance and security, earning users' trust over time.

References

1. C. Costa and K. Bijlsma-Frankema, "Trust and Control Interrelations," *Group and Organization Management*, vol. 32, no. 4, 2007, pp. 392–406.
2. M. Lund and B. Solhaug, "Evolution in Relation to Risk and Trust Management," *Computer*, May 2010, pp. 49–55.
3. D. Gambetta, "Can We Trust Trust?" *Trust: Making and Breaking Cooperative Relations*, Basil Blackwell, 1988, pp. 213–237.
4. S. Perez, "In Cloud We Trust?" *ReadWriteWeb*, Jan. 2009; www.readwriteweb.com/enterprise/2009/01/in-cloud-we-trust.php.
5. B. Michael, "In Cloud Shall We Trust?" *IEEE Security & Privacy*, Sept./Oct. 2009, p. 3.
6. B. Dournaee, "Taking Control of the Cloud for Your Enterprise," white paper, Intel SOA Expressway, June 2010.
7. N. Riter, "VMware Unveils Security API," *Search Security*, Apr. 2009; http://searchsecurity.techtarget.com.au/articles/31679-VMware-unveils-security-API.
8. G. Peterson, "Thinking Person's Guide to the Cloud, Part 3b," blog, 1 Nov. 2009; http://1raindrop.typepad.com/1_raindrop/2009/10/thinking-persons-guide-to-the-cloud-part-3b.html.
9. D. Blum, "Cloud Computing: Who Is in Control?" Burton Group Blogs' Security and Risk Management Blog, 25 June 2009; http://srmsblog.burtongroup.com/2009/06/cloud-computing-who-is-in-control.html.
10. L. Kaufmann, "Can a Trusted Environment Provide Security?" *IEEE Security & Privacy*, Jan./Feb. 2010, pp. 50–52.

11. J. Boles, “Security and the Cloudy Cloud: A Revolution for the Infrastructure,” *Computerworld*, Oct. 2008; http://blogs.computerworld.com/security_in_the_cloud.
12. K. Khan, “Addressing Cloud Computing in Enterprise Architecture,” *Cutter IT Journal*, vol. 22, no. 11, 2009, pp. 27–33.
13. R. Chakraborty et al., “The Information Assurance Practices of Cloud Computing Vendors,” *IT Professional*, July/Aug. 2010, pp. 29–37.
14. T. Ristenpart et al., “Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds,” *Proc. 16th ACM Conf. Computer and Comm. Security*, ACM Press, 2009, pp. 199–212.

***Khaled M. Khan** is an assistant professor and the graduate program coordinator in the Computer Science and Engineering Department at Qatar University. He also holds an honorary adjunct fellow position in the School of Computing and Mathematics at the University of Western Sydney, Australia. His research interests include secure software engineering, cloud computing, measuring security, and health informatics. Khan received his PhD in computing from Monash University. He's the Editor in Chief of the* International Journal of Secure Software Engineering *and is a member of the IEEE Computer Society. Contact him at k.khan@qu.edu.qa.*

***Qutaibah Malluhi** is a professor and head of the Computer Science and Engineering Department at Qatar University. His research interests include storage networking, high performance computing, cloud computing, and computer networks. Malluhi received his PhD in computer science from the University of Louisiana, Lafayette. He's a member of the IEEE Computer Society. Contact him at qmalluhi@qu.edu.qa.*

Top articles, podcasts, and more.
cn
computingnow.computer.org