Towards Building an Accountable and Trustworthy Cloud

Name: Ragib Hasan

Affiliation:
Assistant Research Scientist and NSF/CRA Computing Innovation Fellow,
Department of Computer Science, Johns Hopkins University.

Email: rhasan7@jhu.edu

Research interests:

- Cloud Security, Privacy, and Auditing
- Green Clouds

Summary of current research activities related to Cloud Computing

My research spans several disciplines related to trustworthy cloud computing, distributed systems and database security, secure electronic records. In previous projects, I addressed the problem of protecting integrity of file system and database transaction history in an untrusted distributed environment. I developed efficient and provably secure techniques for securing data provenance for databases and file systems. I also introduced novel techniques for securing databases (reducing audit times 100-fold compared to the previous best approach) for audit-based compliance with Sarbanes-Oxley Act and Gream-Leach-Bliley Act for financial records, and solving the trustworthy litigation hold problem.

My current work focuses on securing provenance of data while it resides inside a cloud, and designing accountable cloud computing architectures. I am also designing a structured threat model for evaluating cloud computing security, threats, and vulnerabilities.

My other work related to cloud computing focuses on design of novel cloud computing architectures that leverage portable mobile devices such as mobile phones and handheld devices in order to build a lightweight cloud. To this end, we have designed a mobile cloud architecture that is built on top of the existing mobile phone infrastructure, and can be cheaply and rapidly deployed for local computing needs.

At the Johns Hopkins University, I have solely designed a new graduate-level course on Cloud Computing security (http://www.cs.jhu.edu/~ragib/sp10/cs412/), which I have taught in Spring 2010 and am offering again in Spring 2011. This is among the very first courses on Cloud Security taught at the graduate-level in US universities. The goal of this class is two-fold – to explore the new landscape of cloud computing security in order to fathom the issues and challenges; and to encourage the next generation of researchers into learning and solving problems of cloud computing security.

ABSTRACT

With the advent of Web 2.0 and cloud-computing, the whole paradigm of computing is at a crossroads. I envision that in the next few years, a large part of computing will shift to the cloud-computing model. However, before that can happen, we have to solve some fundamental problems of cloud computing.

Why hasn't every organization moved to cloud computing rather than maintaining their own data processing systems? The answer is: lack of trust, security, and accountability. Today's data and compute clouds are essentially opaque systems, where the clients have no control over and limited information about what happens inside the cloud. This in turn leads to less trust on the computation done and data stored in clouds. Lack of accountability and concern over the integrity and confidentiality of data and computation limit widespread adoption of clouds in the mainstream of computing. Clients also do not have any trustworthy method of monitoring the data storage and progress of running jobs. These problems have caused many companies to resort to "private clouds" for their sensitive data and computations, which, unfortunately, is contrary to the main philosophy of cloud computing – to provide computing as a service in a flexible, on-demand manner.

How can we make clouds accountable? Accountability in clouds is needed for both customers and cloud providers. A large source of client concern is that, cloud users do not see what happens inside a cloud and how their data is handled. Clients have to fully trust the cloud providers to act honestly and not breach the confidentiality of data and computations. Cloud providers, on the other hand, do not want to disclose the cloud topology and operational details. We need to balance the opposing needs of the providers and clients.

In my future research, I want to make cloud computing more trustworthy and reliable, by bridging the accountability gap. I plan to approach this problem in two directions – by designing cryptographic constructs and mechanisms that will allow the cloud provider to prove the confidentiality and integrity of the data and computation, and by building distributed, efficient, and scalable systems for trustworthy monitoring of clouds without disclosure of sensitive cloud topology information.

One possible approach for overcoming this mutual distrust can be the application of secure data provenance in a cloud environment. My previous research on data provenance security lays the groundwork on which we can tackle the problem of data integrity in a cloud. If a client can receive a verifiable and non-forgeable provenance of the data item's journey through a cloud, then the client can determine the trustworthiness of the data item. At the same time, this would benefit cloud providers in a different way – by enabling them to determine the sources of misconfiguration, errors, and possible attacks on the cloud. An accountable cloud would be very attractive to customers worried about data and computational integrity. However, the problem is non-trivial, as the cloud providers control virtually every aspect of a cloud. I would like to extend the data provenance security schemes I have developed as part of my dissertation to take into account the asymmetric trust scenario in a cloud environment.

Monitoring of ongoing cloud computation is a more difficult problem. How can a client know that the cloud provider is accurately running the task? A possible solution is to make performance monitors an integral part of cloud computing systems. However, this is hard since the cloud provider virtually has complete control over the applications run in the cloud. A possibility is to use a combination of trusted hardware to perform remote attestation of the monitors. In my research, I want to design mechanisms to verify the authenticity of the performance monitor outputs.

In summary, the focus of my research in cloud computing would be to design new mechanisms and tools that will make clouds more transparent and accountable, while preserving the confidentiality of topology and other sensitive configuration information. I believe that this transparency will remove the barriers to widespread adoption of cloud computing in business and healthcare applications.