The Science of Cloud Computing
(application to attend PI meeting)

Name: Reza Curtmola
Affiliation: Department of Computer Science, New Jersey Institute of Technology
Email: crix@njit.edu

Topics of interest: Cloud Security, Privacy, and Auditing; Cloud Architectures and Systems

Current research activities: Security and Privacy of Cloud Storage Systems.

In cloud storage, the storage and management of data is outsourced to third parties called *cloud storage providers (CSPs).* Cloud storage allows thin clients to leverage the massive storage infrastructure available at CSPs in order to store, manage, and distribute large amounts of data at low costs. When data is outsourced, security is a primary concern because data owners lose control over the faith of their data. CSPs must be trusted unconditionally, because existing cloud storage platforms are opaque and outside auditors are not allowed to inspect claims about the data redundancy and protection levels. Coupled with numerous data loss incidents, this prevents organizations from assessing the risk posed by outsourcing data storage to untrusted clouds, making cloud storage unsuitable for applications that need to store sensitive information or that require strong long-term security and reliability guarantees.

We have pursued the following research questions:

- *How can we build long-term data preservation systems on untrusted systems?* Towards this goal, we have developed secure auditing protocols (called *Provable Data Possession* - PDP [1]) that allow for public-verifiability of archival data on untrusted cloud stores. Auditing schemes are tightly coupled with erasure correcting codes and we have demonstrated how to make PDP interoperate with replication [5], MDS codes [4], and network codes [2].
- *How can we ensure privacy of sensitive data without sacrificing useful functionality?* We have developed efficient *searchable encryption schemes* [3] that allow cloud-stored data to be searched, while providing strong confidentiality guarantees against the cloud storage provider.

References

[1] G. Ateniese, R. Burns, R. Curtmola, J. Herring, L. Kissner, Z. Peterson, and D. Song. Provable data possession at untrusted stores. In *Proc. of ACM CCS '07*, 2007.

[2] B. Chen, R. Curtmola, G. Ateniese, and R. Burns. Remote data checking for network coding-based distributed storage systems. In *Proc. of ACM Cloud Computing Security Workshop (CCSW '10)*, 2010.

[3] R. Curtmola, J. Garay, S. Kamara, and R. Ostrovsky. Searchable symmetric encryption: Improved definitions and efficient constructions. In *Proc. of the ACM CCS '06*, 2006.

[4] R. Curtmola, O. Khan, and R. Burns. Robust remote data checking. In *Proc. of ACM StorageSS '08.*

[5] R. Curtmola, O. Khan, R. Burns, and G. Ateniese. MR-PDP: Multiple-replica provable data possession. In *Proc. of ICDCS '08*, 2008.

Future research problems: We seek to address several security and privacy research problems that arise in the context of cloud storage.

How can we build long-term data preservation systems on untrusted systems? Archival storage systems hold large bodies of information that are rarely or never accessed. Preservation of archival data is valuable for future generations. Even is such data sets are rarely/never accessed, they require constant maintenance. These include integrity checks to ensure that the data is intact and security checks to make sure the data has not been lost or altered maliciously. The storage community is now increasingly aware of *silent errors* that damage data, but are not self-evident. These differ from disk failures that raise alerts when they occur. Moreover, cloud storage providers (CSPs) continue to be plagued with with periodic outages and losses of customer data due to power outages, media failures, software bugs, malware, negligence and poor data management practices. This is problematic because data owners can outsource the storage of the data, but cannot outsource the liability for data loss! We plan to seek answers to the following research questions:

- How can we restore control over cloud-stored data to data owners? How to improve the transparency of cloud storage platforms? There is a need to depart from the current model of unconditional trust in the CSPs. If we can better quantify the risks of storing data in the cloud, this will facilitate the migration to the cloud for applications that require long-term security and reliability.

- How can we incorporate strong security and reliability guarantees into untrusted cloud storage while preserving the efficiency advantages that are present in benign settings? How to provide efficient auditing guarantees, while preserving the advantages of the outsourcing paradigm?

- How to *prevent* data loss by designing efficient auditing mechanisms that minimize the workload of the CSP? When components of a storage system fail, how to design *repair* mechanisms that meet various constraints of cloud storage systems? For example, some systems may require to minimize the network bandwidth during repair, whereas other systems may require to minimize the amount of I/O operations in order to reduce power consumption.

A software layer for security and privacy. We aim to design, prototype and integrate into existing cloud storage platforms a software layer that provides support for *searchable encryption* and *remote data integrity checking* services. This layer will allow users to maintain control over their cloud-stored data by enhancing existing cloud storage functionality to include support for data privacy and long-term reliability. The prototyping and integration of this layer will help answer several important questions: How to bridge the gap between existing theoretical constructs and existing cloud storage infrastructure (*e.g.*, Windows Azure, Amazon Web Services)? How to avoid infrastructural changes and accommodate different APIs, structures, user needs and architectural attributes specific to different cloud storage platforms? There is also a need to better understand the performance impact of strong privacy and reliability guarantees on existing cloud inftrastructures.

Improving the security stance of cloud services. Traditional security services (confidentiality, integrity, authentication, access control) must be re-evaluated in the new context of a cloud environment, in order to preserve benefits such as lower resource costs and improved scalability, but also to address new challenges such as multi-tenancy and the need for energy-efficient storage and computation.