NSF PI Meeting: The Science of Cloud Computing

Anna Cinzia Squicciarini
College of Information Sciences and Technology
Pennsylvania State University
University Park (PA)
Email: acs20@psu.edu
Phone: 814-863-7614

Topics:
6. Cloud Security, Privacy, and Auditing,
3. Data Portability, Consistency, and Management

Current Research Activities. In collaboration with faculty from Missouri Science and Technology (Dan Lin), the PI has published several collaborative works on data privacy and protection issues in the cloud. In particular, we analyzed existing access control techniques that may apply and summarized desired properties that an efficient and effective data protection approach should have. Such properties are, for example, that users actual data file should be protected according to the specified access control policies whenever and wherever the file is located in the cloud; users should have certain degree of ability to control other parties usage of the data files; techniques should allow users to obtain different levels of privacy preferences over their data files. To achieve these goals, we have proposed a comprehensive data protection model which enables users to select service providers compatible to their privacy preferences, to reach service agreement and to enforce these agreements throughout the entire life cycle of the user data in the cloud. To facilitate the selection of protection models and techniques according to specic scenarios, we also defined cost functions based on factors like privacy preservation levels, communication cost and processing cost. Moreover, we have also studied an interesting privacy problem about information leakage caused by indexing in the cloud, for which we developed a portable data binding technique to ensure strong enforcement of users privacy requirements at server side.

Future Research Directions

Cloud computing enables highly scalable services to be easily consumed over the Internet on an a need basis. While cloud computing is expanding rapidly and used by many individuals and organizations internationally, data protection issues in the cloud have not been carefully addressed at current stage. Users' fear of confidential data (particularly financial and health data) leakage and loss of privacy in the cloud may become a significant barrier to the wide adoption of cloud services. For example, the software as a service paradigm has increased the frequency of collaborations across enterprise boundaries, raising concerns from organizations that sensitive information may be leaked to outsiders due to mistakes or inappropriate le sharing decisions by their employees.

The current scenario therefore demands systematic and usable approaches to assist users to effortlessly monitor and control the access to their data as it travels across the cloud. Access control techniques appear to be suitable for preventing data leakage and facilitate control of users data managed on the cloud. However, our research team has identified many unique challenges regarding access control in the cloud which have not been addressed by any existing works. In particular, one of the most compelling issues is to *ensure that data usage agreements are actually honored*. This is especially challenging in the cloud since users' data may be distributed to multiple parties that the users do not know, and these parties may satisfy different security requirements using various mechanisms, interfaces and semantics. The expected protection technique should be able to reassure individuals that their personal data is being used for the use they consented and being delivered to only the parties they allowed. Moreover, the technique should not be intrusive to the data recipients' system in order to be widely adopted.

To address this important challenge *our research group is currently exploring innovative policy enforcement techniques that support self-protection policies, automatic accountability engine and strong enforcement mechanism*. Our envisioned approach is unique in that it is flexible and highly distributed to be adopted in the cloud; it goes beyond traditional access control by monitoring and ensuring the actual data usage whenever and wherever the data is located or transferred in the cloud. Specifically, we are investigating how to leverage mobile code techniques to *cage user's content* and protect it by automatically executable access control policies. By means of our caging technique, we aim ensure that the content (i.e., data) is mobile without requiring any specialized software or reliance on third parties; and also ensure that any access to the content can be audited and the corresponding policies are strongly enforced.