Auditable Anonymity

Sonia Jahid[1] and Nikita Borisov
University of Illinois at Urbana-Champaign
{sjahid2,nikita}@illinois.edu

Cloud services offer an attractive platform for storing and sharing data, but they also bring with them new challenges in security and privacy. For sensitive data, such as electronic health records, or personal interactions over social networks, it is important to avoid revealing undue information to the cloud provider. Encrypting the data can solve part of the problem, as the data owner can control access by distributing cryptographic keys to appropriate users. However, certain information, such as access patterns, are still available to the cloud provider and it can tell, for example, what doctor or hospital looked up your medical record, or who you interact with socially. At the same time, cryptographic access control makes it difficult to implement important security protections, such as revoking access or maintaining an audit log of accesses.

In our work, we wish to allow a data owner to be able to receive a log of data accesses, while at the same time hiding this information from the cloud provider. The use of anonymizing tools, such as Tor [4], for reaching the cloud service can help with the latter. To address the auditability, we make use of anonymous credentials with revocable anonymity [2]. These credentials allow users to prove that they are authorized to access a service or a piece of data without revealing the contents of the credential and, importantly, their identity. Each access, however, generates an encrypted record that can be used by a special anonymity revocation service to learn who was behind such an access. In our architecture, this service is run by the data owner, who is able to decrypt the audit log records and get an access log.

We extend the revocable anonymous credential scheme to support credential chains that allow a credential owner can delegate some or all of its access rights to someone else by creating a signed credential. We adapt a technique [1] that is able to verify signatures on credentials using *zero-knowledge proofs*: a user who wishes to access data can prove that there exists a chain of credentials that authorizes access for the user without revealing any information about the credential chain itself. At the same time, the user generates an encrypted record containing the information in the chain and proves, once again in zero knowledge, that it is constructed correctly. This record is then kept by the cloud provider for later decryption and verification by the data owner.

Our ongoing work includes an implementation of this architecture, and a formal security analysis. In future work, we plan to integrate this approach with private information retrieval [3], a technique that allows one to hide which record in a database is being accessed.

Use Cases: A motivating use case for such a system is patient-centric electronic health records. The current provider-centric model for keeping health records makes it difficult to share medical information between providers and results in both increased costs and the possibility of errors or omissions in a patient's medical history. A number of efforts have been made to create patient-centric health record services, such as Microsoft's HealthValut, IndivoX, and Google Health. However, these services have run into a number of challenges, including significant privacy concerns, and Google Health has already been shut down. Our auditable anonymity techniques can mitigate the privacy risks associated with external storage of health records while allowing patients to control and monitor how their medical data is being used. Credential chains will allow patients to authorize an entire hospital to view their medical records while keeping an audit log of exactly which doctors or nurses actually accessed the data.

Another application can be found in social networks services: although there have been many previous proposals to encrypt social network data stored by third parties, they typically do not allow a user to track who is looking at their data, which can provide useful feedback in setting privacy policies. Credential chains can enable access by friends of friends, or people even further away in the social graph, while audit logs can act as a check on this level of activity.

References

[1] M. Backes, S. Lorenz, M. Maffei, and K. Pecina. Anonymous Webs of Trust. In *PETS*, 2010.

[2] J. Camenisch and A. Lysyanskaya. An efficient system for non-transferable anonymous credentials with optional anonymity revocation. In *EUROCRYPT*, 2001.

[3] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan. Private information retrieval. In *FOCS*, pages 41–50, Oct. 1995.

[4] R. Dingledine, N. Mathewson, and P. Syverson. Tor: The second-generation onion router. In *USENIX Security*, pages 303–320, 2004.