Travel report from Secure Cloud Services and Storage Workshop 2017

Martin Strand, martin.strand@ntnu.no

The day before ESORICS 2017, the project Cryptographic Tools for Cloud Security hosted the 'Secure Cloud Services and Storage Workshop 2017', with an impressive lineup of speakers.

Let us start by considering the problem(s). The *client* would like to outsource storage and/or computations in order to save money on capacity variations, hardware and maintenance, but without compromising on secrecy, speed, reliability and integrity. On the other hand, the *service provider* (or just the server) wants to make the most of its resources.

Here are some of the challenges.

1. It is a waste of bandwidth and storage space to upload and store the exact same file several times. The simple solution to this *deduplication* problem is to hash the file locally, and check if the same hash already exists on the server. If yes, the file is not uploaded, and the user simply gets access to the file already on the server. However, the client can use this information to search for the existence of some file, so it is a breach of secrecy for another customer.

2. Secret data must be encrypted to protect against an untrusted provider, but the trivial solution means that the customer must download and decrypt the complete dataset in order to search it, reducing the advantages of outsourcing the storage in the first place. We return to possible solutions to this problem.

3. Large amounts of backup data will probably never be used, so the provider has a strong economic incentive to discard data that the user will never again access. This can be catastrophic for the user, but downloading the complete data set to verify it is not feasible.

4. When outsourcing a computation, there is an apparent deadlock: the customer does not want to disclose the data, the provider tries to keep the algorithm secret (otherwise, he would be out of business). There exist solutions to the problem, but efficiency is still an issue. Also, the customer would perhaps like a ensure that the server is actually performing the right computation, and not something much cheaper, and then outputting a less precise answer.

At the workshop, presenters communicated the status on several of these issues. The interested reader can find the slides at `http://scs.iik.ntnu.no/program.php`.

The first topic in the discussion was *multiparty searchable encryption*; that multiple writers and multiple readers can access the system, but possibly with different read privileges. Next, the first speaker discussed *proof of* *retrievability*, which is a way to handle the third issue listed above. The first idea one might come up with is to store hashes of some random data blocks, and query those once in a while. After a number of queries, the selection is exhausted, and the dishonest and resourceful adversary will then respond to his problem by delete all data except those blocks, or even just the hashes of them. The speaker then introduced StealthGuard, which can support arbitrarily many queries.

The next speaker was Moti Yung representing Snap. He showed how the company has made it possible to create an encrypted section only available for the eyes of the user. The trouble is that a password is too weak to function as a key to the encrypted material, so Snap has solved the problem by what is essentially a secret sharing scheme: the key material can only be created if several servers work together, and then combine the material with a PIN and the user's password. While sharing theory may deem it as insecure since all servers could be considered to be under the control of the adversary – Snap itself – Yung argued that this is mitigated using system surveillance, and that it is sufficient for the real world.

David Pointcheval discussed the more theoretical solutions to the 4th challenge, including *Fully homomorphic encryption* (which is still considered highly impractical) and *Function encryption* (which in turn is based on FHE).

I refer to the slides listed on the workshop homepage for details of the other talks.