GENI Attribution Scenarios

Matt Bishop[1], Mina Doroud[1], Jeffrey Hunker[2], Carrie Gates[3]

Scenario I: Illicit downloading of music

The RIAA has detected that pirated music has been downloaded over GENI. They serve an order to determine the culprit on the GENI Projects Office. Under the current structure, based on the IP address, the GPO must determine slice which were used, PI who created the slice and the institution which did the downloading. They must then forward the order to that institution (or inform the RIAA which institution to serve with the order). That institution must then find the downloader, probably by examining the systems involved in running the experiment on the particular slice. At that point, the owner of the system can be identified. Since most of the time, the owner will not be the user who downloaded the music, further investigation is also needed. Based on the information provided by RIAA (IP address, time, etc.) the ultimate user can (hopefully) be determined by the owner.

Here, the type of attribution required is any of the following:

- Perfect attribution, because then all actors and systems are known to everyone;
- Perfect selective attribution, meaning that all actors and systems are known to a select group (in this case, the GPO, institution, and RIAA need to know only some of the components; for example, the RIAA only needs to know the actor who did the downloading or the responsible party)

These can be combined, so some components will be known to all and others only to the PI, GPO, institution, or RIAA.

Scenario II: DDoS attack on GENI

A group of attackers launch a distributed denial of service (DDOS) attack on a GENI slice, disrupting an experiment. The experimenters report the attack to the GENI Projects Office. The GPO needs to trace the attack to its (ultimate) starting point or points. It needs information about the DDOS-related packets so it can attribute the attack to its originators and respond appropriately.

Here, *for the GPO (or the target of the attack)* the type of attribution required is any of the following:

- Perfect attribution, because then the GPO can find the systems that the packets originated from; if the attribution can be ascribed to a particular

[1] Dept. of Computer Science, University of California at Davis, Davis CA 95616-8562 USA
[2] Jeffrey Hunker Associates, Pittsburgh, PA 15232 USA
[3] CA Labs, New York, NY 10022 USA

software component or an individual, then the "real world" source of the attack can also be identified;
- Perfect selective attribution, so that the GPO can obtain the information described above (but others cannot); this may be appropriate if the DDoS attack was accidental due to a faulty software component.

However, if the attack is intentional, the *originators* of the attack will undoubtedly want other types of attribution:
- Perfect Non-Attribution, because they do not want to be identified as the origin (or originators) of the DDoS packets;
- False Attribution, for the same reason as perfect non-attribution, with the added advantage (in the attackers' eyes) of misdirecting the GPO.

This raises an interesting question: if the originators of the attack are within the GENI network, can they utilize aspects of an attribution choice framework to hide their attack, or determine what the investigators are using to track them?

Scenario III: Collecting and sharing data

Institutions share documents with each other in online environments. Documents can belong to different projects and institutions. And many different people within those projects and institutions can work on them. For legal purposes (such as patent or commercialization rights), the institutions involved need to know who modified these documents.

A good instance of this need (also one requiring high levels of assurance) are the needs of lawyers and legal firms to send and share documents with attribution as to origin and changes made.

Here, the type of attribution required is any of the following:
- Perfect attribution, because then the institutions can identify those who modified the documents;
- Perfect selective attribution, for reasons as before, except that only the institutions know.

Some people who modify the documents may want to be anonymous. In that case:
- Perfect non-attribution, so the modifiers cannot be identified

Scenario IV: Elections[4]

Elections are the foundation of democratic and republican societies. Recently, many jurisdictions began exploring people voting over the Internet. Each of the 50 states

[4] Scenarios IV and V are from the paper: Matt Bishop, Mina Doroud, Carrie Gates, Jeffrey Hunker "Using an Attribution Framework: The Second Summer of the Sisterhood", 11th European Conference on Information Warfare and Security ECIW-2012, Laval, France, July 2012.

is responsible for holding its own elections; most delegate this responsibility to the counties, with the state having the ultimate authority to certify the results. In general three transactions are involved in elections. First, the voter registers to vote. Second, the voter receives the correct ballot from the ballot generator. Third, the voter marks his ballot and then transmits the marked ballot to Election Central, where the votes are counted.

Consider each transaction separately. The first transaction speaks to authorization of the voter. By the voter supplying attributes that uniquely identify her, the registration authorities can determine which ballot type she should receive, and set the "authorization to vote" attribute to a value that will enable her to obtain the correct ballot. The voter requires that the registration authority have the attribute "authorized to register voters" with value "true".

The second transaction occurs when the voter acquires the ballot. The voter first verifies that the ballot generation system is authorized to generate the ballots by checking the value of the attribute "authorized to generate ballots". The voter presents the value of the "authorization to vote" attribute, which the ballot server validates as having an acceptable level of assurance. That authorization is used to determine the ballot type that the voter requires. A ballot of that type, with the attribute "issued to authorized voter" and value a nonce is generated and sent to the user.

The third transaction is the casting of the vote. The important artifact here is the ballot; it must come from the ballot server, and be voted by an authorized voter. The voter first contacts Election Central, and checks that the attribute "Election Central" is "true". The voter then transmits her ballot to Election Central, which checks that the attribute "issued to authorized voter" has a nonce that is unused so far. It then processes the ballot and tallies the votes.

Thus, the attributes of interest is any of the following:

- Perfect selective attribution: At each transaction step, before the voting, only few attributions is required by a specific party. For example at the first transaction, the registration authority specifies the unique attributes it requires and the assurance evidence that the value of that attribute (in practice, the address) is correct (in practice, evidence that the prospective voter lives there);
- Perfect non-attribution: After the voting take place, nobody is able to bond the individuals to the vote they made.

Scenario IV: Telephone-to-Tweet Service, Arab Spring

Social activists have learned to exploit Internet and social networking services to communicate across geographic and political boundaries. The "Arab Spring", a term

for the uprisings in the Middle East, is a good example of this. The role that the social networks played was critical not only within the countries, but also in communicating events with the rest of the world. Most of the online services need some sort of attribution provided by the user.

The type of attribution required by the service is any of the following:

- Perfect attribution, then all users are known to everyone;
- Perfect selective attribution, meaning that users are known to a select group (the service itself or to a private network)

During the Egyptian Arab Spring uprisings, the government attempted to block the peoples' access to the Internet. Not only was access to the online services infeasible but it was also dangerous. Many of the on-line services made available attributes that could identify the users even if they tried to remain anonymous—and government agencies could, and did, make use of this information to pursue those involved in anti-government activities. In response, Google and Twitter provided a "telephone-to-tweet" service. Google established 3 telephone numbers that people could call and record a message. Google then posted it to Twitter. Associated with each message was a hash tag indicating the geographic origin of the message, when that could be determined; otherwise, the message was posted without a hash tag.

The only attribute the "telephone-to-tweet" service requested was the country from which the message originated. In most cases, the telephone system would supply this information to an acceptable level of assurance. If the information were not available, no associated hash tag would be generated. In other words, the server would request the "country of origin" attribute, but if the client could not supply it, the policy negotiation mechanism rolled back to accepting the (unattributed) message.

The "telephone-to-tweet" service is an example of

- Perfect non-attribution: none of Google, Twitter, nor any listener is to be able to identify the speaker from any metadata.

These particular scenarios demonstrate that whatever GENI-specific attribution framework is developed, it must be expandable to a larger, more universal network. These "outside GENI" scenarios are examples of this.