Appears in HotOS-VIII (IEEE Workshop on Hot Topics in Operating Systems), May 2001.

Position Summary: Authentication Confidences

Gregory R. Ganger
Carnegie Mellon University
ganger@ece.cmu.edu

Abstract

"Over the Internet, no one knows you're a dog," goes the joke. Yet, in most systems, a password submitted over the Internet gives one the same access rights as one typed at the physical console. We promote an alternate approach to authentication, in which a system fuses observations about a user into a probability (an **authentication confidence***) that the user is who they claim to be. Relevant observations include password correctness, physical location, activity patterns, and biometric readings. Authentication confidences refine current yes-or-no authentication decisions, allowing systems to cleanly provide partial access rights to authenticated users whose identities are suspect.*

1 The Case for Authentication Confidences

Access control decisions consist of two main steps: authentication of a principal's digital identity and authorization of the principal's right to perform the desired action. Well-established mechanisms exist for both. Unfortunately, authentication in current computer systems results in a binary yes-or-no decision, building on the faulty assumption that an absolute verification of a principal's identity can be made. In reality, no perfect (and acceptable) mechanism is known for digital verification of a user's identity, and the problem is even more difficult over a network. Despite this, authorization mechanisms accept the yes-or-no decision fully, regardless of how borderline the corresponding authentication. The result is imperfect access control.

Using authentication confidences, the system can remember its confidence in each authenticated principal's identity. Authorization decisions can then explicitly consider both the "authenticated" identity and the system's confidence in that authentication. Explicit use of authentication confidences allows case-by-case decisions to be made for a given principal's access to a set of objects. So, for example, a system administrator might be able to check e-mail when logged in across the network, but not be able to modify sensitive system configurations. This position paper discusses identity indicators, and our full white paper [1] completes the case.

2 Human identification and confidence

Identity verification in most systems accepts any user presenting a predetermined secret (e.g., password) or token (e.g., ID card). The conventional wisdom is that, since they are private, no additional information about the likelihood of true identity is necessary or available. We disagree. For example, a system's confidence in the provided password could certainly depend upon the location of its source. As well, a gap of idle time between when the password was provided and a session's use might indicate that the real user has left their workstation and an intruder has taken the opportunity to gain access.

A controversial emerging authentication mechanism compares measured features of the user to pre-recorded values, allowing access if there is a match. Commonly, physical features (e.g., face shape or fingerprint) are the focus of such schemes, though researchers continue to look for identifying patterns in user activity. Identifying features are boiled down to numerical values called "biometrics" for comparison purposes. Biometric values are inherently varied, both because of changes in the feature itself and because of changes in the measurement environment. For example, facial biometrics can vary during a day due to acne appearance, facial hair growth, facial expressions, and ambient light variations. Similar sets of issues exist for other physical features. Therefore, the decision approach used is to define a "closeness of match" metric and to set some cut-off value — above the cut-off value, the system accepts the identity, and below it, not.

Confidence in identity can be enhanced by combining multiple mechanisms. The simplest approach is to apply the mechanisms independently and then combine their resulting confidences, but more powerful fusing is also possible. For example, merged lip reading and speech processing can be better than either alone. Note that if the outcomes conflict, this will reduce confidence, but will do so appropriately.

References

[1] Gregory R. Ganger. *Authentication Confidences*. CMU-CS-01-123. Technical Report, Carnegie Mellon Univeristy School of Computer Science, April 2001.