EXHIBIT 21

Authentication Confidences

Gregory R. Ganger
ganger@ece.cmu.edu

April 28, 2001

April 2001
CMU-CS-01-123

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

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.*

We thank the members and companies of the Parallel Data Consortium (at the time of this writing: EMC Corporation, Hewlett-Packard Labs, Hitachi, IBM Corporation, Intel Corporation, LSI Logic, Lucent Technologies, Network Appliances, Panasas, Inc., Platys Communications, Seagate Technology, Snap Appliances, Sun Microsystems and Veritas Software Corporation) for their insights and support.

Keywords: security, authentication, biometric authentication, system access.

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.

This white paper promotes an alternative approach in which the system remembers 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. The remainder of this section discusses various causes of identity uncertainty and existing mechanisms for dealing with it. The following section discusses how authentication confidences might be added to systems.

1.1. Human identification and confidence

In current computer systems, authentication of a user's digital identity relies on one or more mechanisms from three categories:

∞ **Something one knows.** The concept here is that if the user knows a pre-determined secret, it must be the right person. The common type of secret is a password, though other schemes like images [5] and patterns are being explored. The conventional wisdom is that since it is a secret, 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 — the likelihood of an imposter providing my password from my office is much lower than the likelihood of them providing it over the network (especially from the Internet or the dormitories). 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.

∞ **Something one has.** The concept here is that if a user has a pre-configured item, it must be the right person. The common item is some kind of smart card or ID badge. The conventional wisdom is that anyone who has the token should have full access and that no other information is needed. Again, we disagree. As with the password example, location of token and time since session use can both affect the confidence that a system should have in the corresponding authentication. More radical out-of-band information, such as the owner's expected location based on scheduled appointments, could also provide insight.

- **Something one is.** The concept here is that the system compares measured features of the user to pre-recorded values, allowing access if there is a match [1]. 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. More drastic changes result when switching between eyeglasses and contact lenses or upon breaking one's nose. 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. When setting the cut-off value, an administrator makes a trade-off between the likelihood of false positives (allowing the wrong person in) and false negatives (denying access to the right person). Figure 1 illustrates this process and the corresponding trade-off. Note that we are not suggesting elimination of the cut-off. Instead, we are suggesting that the amount by which the observed value exceeds this cut-off be remembered as part of confidence.

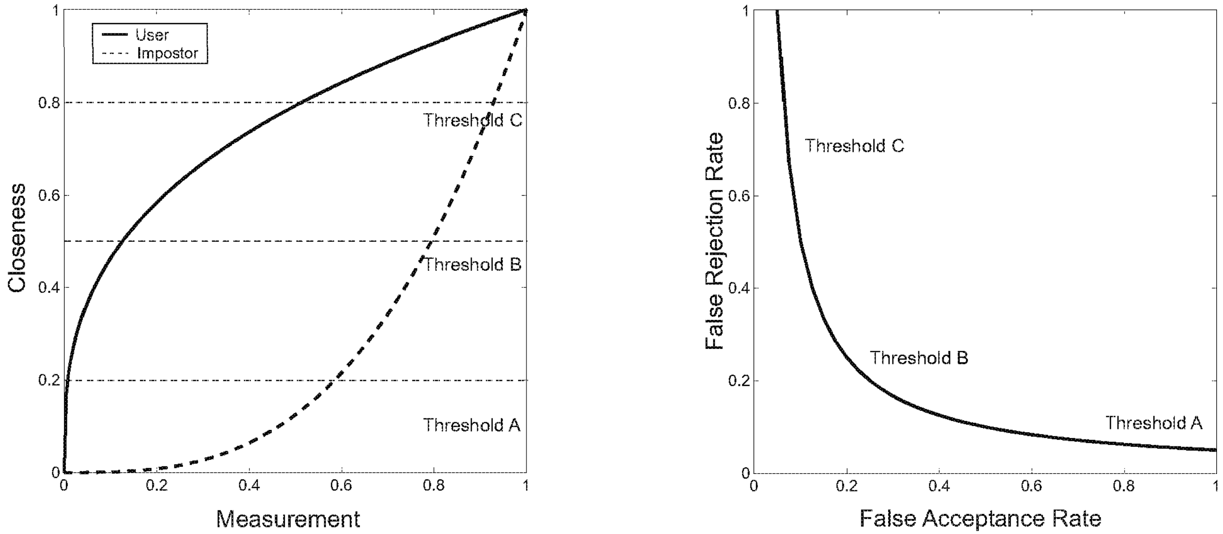

(a) closeness and cut-off points

(b) false accept/reject trade-off

Figure 1: Illustrative example of closeness-of-match thresholds for biometric-based authentication and the corresponding trade-off between false acceptance rate and false rejection rate. On the left, (a) shows possible distributions of closeness values for a user and an imposter. Notice that each cut-off threshold will sometimes reject the real user and sometimes accept the imposter. Specifically, at a given cut-off threshold, false accepts are to the right of the dashed line and false rejects are to the left of the solid line. As biometric accuracy improves, the area beneath the user's distribution will increase and that beneath the imposter's curve will decrease. On the right, (b) illustrates the trade-off between false acceptance rate and false rejection rate more directly with the common "Receiver Operating Characteristic" curve. Better biometric accuracy would reduce the space beneath this curve.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.

WHAT WILL YOU BUILD? | sales@docketalarm.com | 1-866-77-FASTCASE

fastcase®
Smarter legal research.