

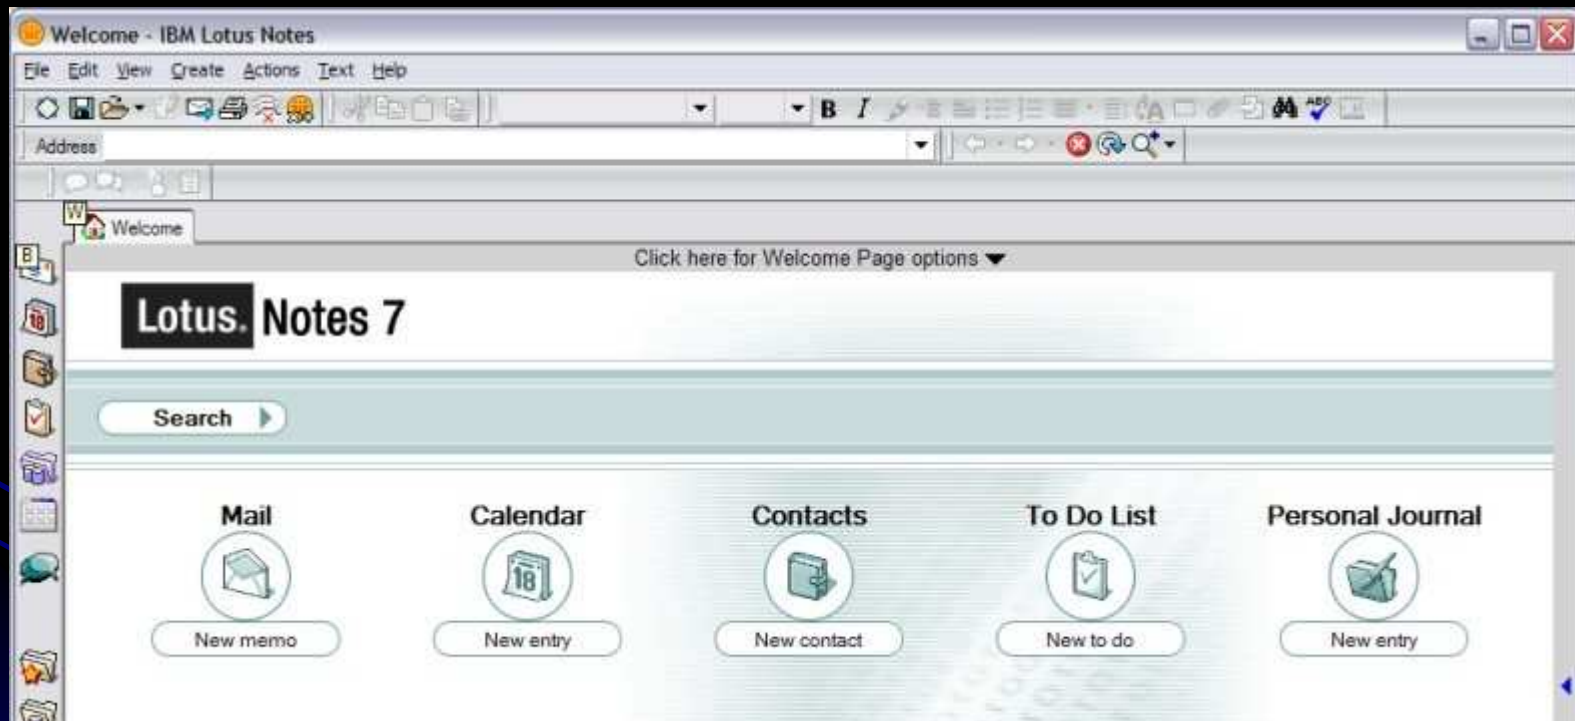
Mitigating Dictionary Attacks on Password- Protected Local Storage

Ran Canetti, **Shai Halevi**, Michael Steiner
IBM T. J. Watson Research Center

<http://eprint.iacr.org/2006/276>

Motivation

- 1 This is how we start our day at IBM:



- 1 What happens when I type in my password?

Encrypted Local File

- 1 Notes' startup does not rely on network
 - 1 Configuration/credentials stored on local disk
 - 1 Credentials are encrypted
- 1 Upon startup
 - 1 Derive key from password
 - 1 Use it to unlock the credentials
 - 1 Then use credentials for everything else

What if I lose my laptop?

- 1 No worries, my credentials are protected
 - 1 My top secret password (sh@1) to the rescue
- 1 What about users with weak passwords?
 - 1 Attacker can mount off-line dictionary attack:
 - 1 For each password in the dictionary
 - 1 Derive a key from the password
 - 1 Check if it decrypts the file
- 1 Is it possible to protect against it?
 - 1 **Without connectivity or secure hardware?**

Other solutions (out of scope)

- 1 Relying on “secure hardware”:

- 1 Store secrets in secure hardware, use password to unlock hardware

- 1 Restrict password-guessing attacks

- 1 Relying on the network:

- 1 Store secrets on server, use password to authenticate to server, then get secrets

- 1 Mitigate on-line attacks

- 1 **This work: what if you cannot do either?**


Key-derivation from passwords

- 1 Common practice: use salt and a deliberately-slow key derivation function
 - 1 E.g., $\text{key} = \text{SHA1}^{65536}(\text{salt} \mid \text{password})$
 - 1 Different salt values for different users, salt is stored on disk.
- 1 Linear slow-down for the attacker
 - 1 But 65536 must grow as computers get faster
- 1 Can we do better?

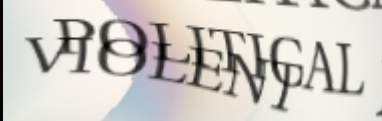
A Different Approach

- 1 A key-derivation protocol (user \leftrightarrow laptop)
 - 1 User does more than just providing password
 - 1 Using “human-only solvable puzzles”
 - 1 People can solve these puzzles
 - 1 Computers cannot
- 1 User enter password, solves puzzles
- 1 Key is derived from both password and the solutions to these puzzles

Puzzles



CAPTCHAs [Na96, vABHL03]

- 1 Example: What's written here? 
- 1 Can generate automatically with solution
 - 1 Without any secret information
- 1 People can usually solve them
- 1 Solving them automatically is beyond the state of the art

CAPTCHAs & pwds [PS02]

- 1 Limit on-line attacks in client-server setting
 - 1 Server generates CAPTCHAs + answers, sends to user
 - 1 User sends back solutions, server checks
 - 1 Then run pwd-based key-exchange protocol
- 1 Not applicable in our setting
 - 1 Where to store the solutions?

Inkblot authentication [SS04]

- 1 A different type of puzzles: one user's solution is unpredictable by other users



What do you see in this picture?

- 1 Solutions used to generate a strong pwd

Is Inkblots the answer?

- 1 Nice work if you can get it...
 - 1 Unpredictability from other people a tall order
 - 1 Need “many bits” in each puzzle
 - 1 People should remember their answers
- 1 How do we know that they are really unpredictable?
 - 1 Maybe with some demographic information they become predictable...

This Work

- 1 A more robust protocol
- 1 Same protection when puzzles are inkblots
- 1 But strong protection even if solutions are predictable by other people
 - 1 Can use CAPTCHAs (or in between)

Our protocol

- 1 Many puzzles are stored on disk ($z_1 \dots z_n$)
- 1 User's password is used to select a few
 - 1 $\langle i_1, \dots, i_\ell \rangle \leftarrow \mathbf{Expand}(\text{salt}, \text{password})$
- 1 User is asked to solve these few puzzles
 - 1 Solutions are s_1, \dots, s_ℓ
- 1 Key is derived from password+solutions
 - 1 $\text{Key} \leftarrow \mathbf{Extract}(\text{salt}, s_1, \dots, s_\ell, \text{password})$
- 1 **Goal: attacker must solve MANY puzzles to find the key**

An example

- 1 Store 2^{20} CAPTCHAs (fit on one DVD-R)
- 1 User needs to solve eight CAPTCHAs
 - 1 $\langle i_1, \dots, i_8 \rangle \leftarrow \text{HMAC-SHA1}_{\text{salt1}}(\text{pwd})$
 - 1 Each index is 20-bit long
- 1 $\text{key} \leftarrow \text{HMAC-SHA1}_{\text{salt2}}(s_1, \dots, s_8, \text{pwd})$
- 1 An attacker that solves 10,000 CAPTCHAs has $< 1\%$ chance of hitting four of the eight CAPTCHAs that the honest user uses

Properties of puzzles

- 1 Automatically-generated problems
 - 1 $z \leftarrow G(\text{aux})$, G is randomized
 - 1 aux can be user-supplied input (family pics?)
- 1 Consistently solved by each human user
 - 1 $s \leftarrow H(z)$, consistent across time*
 - 1 Different users need not agree on an answer
 - 1 But answers need not be unpredictable
- 1 Hard to solve for a machine

* Can we use fuzzy extractors to correct a few errors?

“Human-only solvable puzzles”

- 1 Fairly weak requirements
- 1 Need not be CAPTCHAs
 - 1 Don't need to generate puzzle+solution
 - 1 Not necessarily one right solution
- 1 Need not be Inkblot
 - 1 One user's solution not necessarily unpredictable by other users
- 1 Can be many things in between

Toy Examples

1 Rank these people by coolness



Andrew Yao



Um Kulthum



Helen Keller



Tom Cruise

1 Which of these pictures doesn't belong?



Hardness of puzzles

- 1 What we need: hard to distinguish the “real solution” from a “random solution”
- 1 (G,H) is μ -hard if there exists distribution R with μ bits of min-entropy such that
 - 1 $z \leftarrow G()$ is a random puzzle
 - 1 $s \leftarrow H(z)$ is the right solution
 - 1 $s' \xleftarrow{\$} R(z)$ is a random solution
 - 1 Attacker (PPT) cannot tell (z,s) from (z,s')

Challenge: design good puzzles

- 1 Need “many bits of hardness” for a construction to be useful
 - 1 Else user is bothered with many puzzles
- 1 Aside: must we store puzzles on disk?
 - 1 Use $r \leftarrow f(\text{salt}, \text{pwd})$ as randomness to generate the puzzles?
 - 1 Say, f is a random oracle
 - 1 Puzzles must be hard even if attacker knows the randomness

Security Analysis

Adversarial model for protocol

- 1 Attacker: not just PPT TM
 - 1 Can also get help from people
- 1 Protocol has access to human help, why not the attacker?
- 1 This is a realistic attack
 - 1 Used against deployed CAPTCHA systems
 - 1 Attackers ship CAPTCHAs to their own web-sites, ask their visitors to solve them

Modeling “human attackers”

- 1 People can do many things
 - 1 Outside the model: invite target to dinner, get her to disclose her password
- 1 We assume: attacker only uses humans as puzzle-solvers
 - 1 Attacker has oracle access to H
 - 1 Or a “noisy version” of it (?)
 - 1 Makes analysis possible
 - 1 Keep in mind that it is not entirely realistic

Formal adversarial model

- 1 Attacker: efficient automated program (PPT TM) with puzzle-solving oracle
- 1 Resources: time, number of queries
 - 1 E.g., polynomial-time, sub-linear # of queries
- 1 Goal: distinguish key from random

Notions of security (1)

- 1 Indistinguishability ([BR93]-style)
 - 1 Attacker gets key, puzzles, salt, needs to decide if key is real or random
 - 1 Will focus on this notion in this talk
- 1 Bound attacker's advantage in terms of:
 - 1 Parameters (n puzzles on disk, user solves ℓ)
 - 1 Size of password dictionary ($|D|=d$)
 - 1 Number of oracle queries ($q \ll n$ queries)
 - 1 Hardness of puzzles vs. key-length

Notions of security (2)

- 1 UC: define “ideal functionality” that only allows a limited number of password guesses:
 - 1 Parameters: **D: dictionary**, p : #-of-pwd-guesses
 - 1 Init(pw,aux) from user U
 - 1 store pw, generate and store random key
 - 1 Send aux to adversary
 - 1 **Check that $pw \in D$, else give it to adversary**
 - 1 Recover(pw') from anyone
 - 1 If $pw' = pw$ then return key to U
 - 1 Password(pw') from adversary
 - 1 If already made p such queries then ignore
 - 1 Else if $pw' = pw$ return key to adversary

A “generic” attack

- 1 An attack with complexity $\sim |D| \cdot 2^\mu$
 - 1 Attacker does not solve puzzles
 - 1 Works even in the random-oracle model
- 1 Assume: given a puzzle z , attacker can generate a list of 2^μ potential solutions
 - 1 The right solution is in the list
 - 1 But the attacker does not know where
- 1 This is consistent with a μ -hard puzzle
 - 1 And realistic attackers often have this ability

The attack setup

- 1 Attacker gets (z_1, z_2, \dots, z_n) , salt (if any), and an alleged m -bit key k^*
 - 1 Can make $\leq q$ queries to puzzle-solving oracle
- 1 Needs to distinguish between:
 - 1 “Random”: k^* is random,
 - 1 “Real”: $k^* \leftarrow \mathbf{Extract}(\text{salt}, s_{i_1}, \dots, s_{i_\ell}, \text{pwd})$, where $\langle i_1, \dots, i_\ell \rangle \leftarrow \mathbf{Expand}(\text{salt}, \text{pwd})$ and s_i is the right solution for z_i

Generic attack, phase 1

- 1 For each $p \in D$
 - 1 Compute $\langle i_1, \dots, i_\ell \rangle \leftarrow \mathbf{Expand}(\text{salt}, p)$
 - 1 Generate 2^μ solutions for each z_{ij}
 - 1 For a total of $2^{\mu\ell}$ solution-vectors
 - 1 Keep only those solution-vectors with $\mathbf{Extract}(\text{salt}, s_1, \dots, s_\ell, p) = k^*$
 - 1 These are the “consistent vectors”
- 1 So far, didn't make any oracle queries
 - 1 If $m \gg \mu\ell$, can already distinguish

Generic attack, phase 2

- 1 Query the puzzle-solver upto q times
- 1 Purge vectors (s_1, \dots, s_ℓ, p) for which any of the solutions is not the right one
- 1 Choose queries to maximize mutual-info between the answer and your decision
 - 1 Or use some greedy strategy


Generic attack, phase 3

- 1 Count remaining solution-vectors that are consistent with each password
- 1 Maximum-likelihood: are these numbers more likely for “real” or “random”?
- 1 The point: once we ask on “too many puzzles” for a password, we expect to have zero remaining consistent solutions in the “random” case

The moral

- 1 Attacker must query its oracle on many puzzles that are mapped to right pwd
- 1 We prove: not just a feature of this attack
 - 1 As long as remaining un-queried puzzles have more pseudo-entropy than key-length, key is secure
- 1 Many many details/open problems
 - 1 Some examples next, more on ePrint

The function **Expand**



- 1 Goal: attacker's queries “almost cover” only small fraction of the actual dictionary

The function **Expand** (2)

- 1 \forall large enough D , $\forall |Q|=q$ ($q \ll n$),
 $a\text{-cover}(Q)$ is a very small fraction of D
 - 1 If we were talking about $\text{cover}(Q)$: the neighbor-set of any large subset of D contains more than $|Q|$ neighbors (expansion)
 - 1 Since we want $a\text{-cover}(Q)$: same holds even when dropping many edges
 - 1 As long as the degree remains $> m/\mu$
 - 1 Similar to fault-tolerant expansion

Constructing Expand

- 1 Huge pwd-universe ($\{0,1\}^{160}$)
 - => no deterministic construction
 - 1 Deterministic construction for small D? (open?)
 - 1 Randomized construction?
- 1 Expand as a truly random function
 - 1 Standard analysis using Chernoff
 - 1 Ugly bound, but useful in specific cases
- 1 Expand as n -wise independent
 - 1 Use n -th moment inequality

Constructing Expand (2)

- 1 Can we do better?
- 1 Speculation: ℓ independent random linear maps over $\text{GF}(2)$ work well
 - 1 Old result of Alon et al. (“linear hashing yields small buckets”): a linear map over $\text{GF}(2)$ works well for $q=\ell=1$, and $|D|=n$

The function **Extract**

- 1 Extracts m -bit key from puzzle solutions
- 1 Key is pseudo-random as long as $>m$ bits of pseudo-min-entropy are left in un-queried puzzles
- 1 Strong randomness extractor is sufficient
 - 1 From $m^* > m$ bits of min-entropy, extracts a key that is δ away from uniform m -bit string
 - 1 ℓ^* - number of puzzles needed to get m^* bits of pseudo-min-entropy

Security Statement

- 1 Assuming puzzles are μ -hard (and fixing parameters $n, \ell, m, D, \ell^*, \delta$), an attacker making q queries has advantage at most

$$\text{a-cover}(q, D) + \delta \cdot \binom{\ell}{\ell^*} + \text{negligible}$$

Caveat Emptor

Non-malleability of puzzles

- 1 Current analysis allows the attacker to query its humans only on puzzles that are stored on the disk
- 1 To remove this restriction, puzzles need to be non-malleable
- 1 Not clear how to define/achieve non-mal
 - 1 Even if we don't care about human-solvable, e.g., non-malleable OWFs, PRGs, ...

Some open problems

- 1 Design good puzzle systems
- 1 Design of **Expand** function
 - 1 Should be good (fault-tolerant) expander
- 1 Better protocols (< storage, > security, etc.)
- 1 (Non)-malleability of puzzles
- 1 Better modeling of the attacker
- 1 Better UC analysis

Thank you