Localization of Credential Information to Address Increasingly Inevitable Data Breaches*

Mohammad Mannan and P.C. van Oorschot
Carleton University, Canada

Abstract

Large-scale data breaches exposing sensitive personal information are becoming commonplace. For numerous reasons, conventional personal (identification) information leaks from databases that store online and/or on-site user transaction data. Collected ID numbers and supporting personal information enable malicious parties to commit large-scale identity fraud. Gates and Slonim (NSPW 2003) proposed the owner-controlled information paradigm to address privacy violations of personal information where users are expected to maintain all their information using a personal device. Rubin and Wright (FC 2001), Molloy et al. (FC 2007), and others explored the use of one-time numbers to address credit card fraud (mostly for online use). However, several other types of ID number are at least as sensitive as credit card numbers. Our fundamental assumption is that collected personal information will eventually be breached. To combat identity fraud under this new environmental attack paradigm, we introduce a more general approach involving localized or customized ID numbers for both card-present and card-not-present transactions. We also explore four variants of the general idea to spark more discussion and further research in this area.

1 Introduction and Motivation

Currently personal identity information is stored in a number of different places including small and large corporations, government agencies, educational institutes, hospitals, and financial data processing centers. Coupled with such data replication, insider abuse (e.g. [7]), negligence (e.g. [18, 8]), inadequacy of existing technology for protecting user data, and a computing environment arguably "under occupation" [24] (by e.g. malicious software and semantics attacks) have resulted in numerous large-scale data breaches. The U.S. Government Accountability Office (GAO) defines data breach as "an organization's unauthorized or unintentional exposure, disclosure, or loss of sensitive personal information, which can include personally identifiable information such as Social Security numbers (SSN) or financial information such as credit card numbers" [15]. Data breaches from organizations, small and large, considered to be highly secure or otherwise, make the news almost every day, and now seem to be the business norm. Beyond simple credit card numbers, leaked information now includes SSN, drivers' licenses, dates-of-birth, and bank account numbers. Aside from privacy exposure, these breaches facilitate *identity fraud*,[1] heavily exploited by underground criminal networks. For example, according to Symantec [43, p.23], an individual's *full identity* (which may include name, address, date of birth, SSN, driver's license number) can be bought for only $1-15. One primary reason for the enormous demand of compromised personal records is that most existing ID numbers are static, and thus reusable elsewhere (especially where the corresponding physical ID token is not required or the token can be easily forged).

In response to large-scale data breaches, security proponents have placed increased importance on data encryption, use of sophisticated intrusion detection technologies, etc. However, these conventional techniques are still not widely deployed, and also have been subject to a continual stream of innovative attacks, from *side-channel analysis* of cryptographic keys (e.g. timing/power analysis attacks), to the recent *cold boot attacks* [19]. Additionally, such technologies are of limited help in the case of organizational mismanagement.[2]

*Version: July 18, 2008. Contact author: mmannan@scs.carleton.ca.

[1] We define identity fraud as unauthorized exploitation of credential information through the use of false identity [29].

[2] According to one study [49], 87% of the breach cases analyzed could have been prevented with "reasonable security controls."

In the financial sector, as credit card number disclosures increased, some banks started to offer one-time use credit card numbers for online transactions around Sept. 2000. Several research proposals (e.g. [37, 25, 27]) have been made focusing mainly on enhancing such credit card number generation and user-friendliness.

We argue that improving data security mechanisms or new legislation for protecting consumer information is of limited use. Identity fraud originating from data breaches will grow as more and more identity information is collected and stored digitally.[3] In an environment where data breaches are evidently inevitable, it is our main thesis that the use of static/reusable ID numbers should be reduced, if not completely eliminated, to fight identity fraud. Building on existing ideas and experience regarding disposable credit card numbers, we propose a more general approach and technique to deploy an *ID number localization* approach to restrict and/or detect abuse of a wide variety of sensitive personal identification numbers. Our use of the term "localization" is primarily intended to mean customizing ID numbers to a specific relying party, which need not be tied to a particular geographic or physical location. We also outline four variants of our main proposal. Despite these proposals, our primary goal is to increase awareness of the new environmental attack paradigm by which ID numbers ultimately become compromised; different solutions are explored here to motivate further research in this area.

An individual may be required to provide their SSN or driver's license number to several parties (employers, banks, credit reporting and car rental agencies), all of whom store sensitive identification details for a long time. Confidentiality of such data may be breached by any of these parties. If a person has worked for five different companies in the past, her SSN may leak from any of those, and once disclosed may facilitate identity fraud. If a localized SSN scheme were in place, where each employer would get a "different (non-reusable) version" of the SSN of the given individual, then a disclosure of any such SSN would not be useful for identity fraud. We base example solutions on this idea and explore several variants. Again, our fundamental assumption is that user data will eventually be breached (cf. [45, 24]) primarily through relying parties; we focus on how to nonetheless mitigate identity fraud.

In a broader sense, one obvious reason for the severity of current identity fraud, spam, phishing, and many other Internet-related attacks is the leverage gained by using data compromised from one site at many others, repeated times. This "compromise once, reuse multiple times" feature provides significant advantages to attackers. Our approach is a defensive paradigm of (virtual) localization for the *use* of credential information on the Internet and in the physical world. Localized identification numbers as generated by our scheme are valid only for a particular relying party. This apparently reduces the value of compromised credential information to attackers, thereby reducing the threat and also the cost to defend adequately. Our approach attempts to undermine the asymmetric leverage attackers currently enjoy.

In summary, our contributions and discussion points for NSPW include:

1. New Paradigm for Protecting Personal Identity Information. Legislative and technical efforts such as encryption alone to better secure personal identification data are evidently inadequate in today's untrusted computing environment. As one response, we propose the use of localized, restricted-use identification numbers instead of static, reusable ID numbers to limit large-scale identity fraud.

2. Breadth of Scope. We focus on protecting all types of identification numbers in general instead of solely credit card numbers. Where most previous solutions focus on card-not-present transactions, we address both card-present (ID-present) and card-not-present transactions. Furthermore, our approach addresses breaches resulting from real-world (offline) incidents such as lost/stolen disk drives, and backup tapes (i.e. independent of computers being compromised by malware).

3. Variations. To take into account deployment feasibility, and cost-benefit trade-offs, we explore several variants of our proposal (appropriate for varying scenarios).

Overview. We outline our main proposal along with threat model, notation and operational assumptions in Section 2. Four variants of the main proposal are introduced in Section 3. In Section 4 we briefly discuss related work and representative examples of recent data breach incidents. Section 5 concludes.

[3] "...the difference between such crimes [ID theft] today and in the future is the scale of the data involved" [51].

2 A Strawman Proposal

In this section, we outline our proposed ID localization scheme. Threat model, notation and operational assumptions are also discussed here.

Overview. A credential issuing party provides each user a smart card (chip-card) with a unique identification number for the user and a 'secret' key both stored on the chip and in print on the card itself.[4] For example, the credential issuing party for SSNs is the Social Security Administration (SSA), a user's SSN is a unique ID number (issued by SSA), and a secret key is a random string of digits or characters of sufficient length (e.g. 128 bits). As the secret key is stored on an ID card, the user does not have to memorize it. Using a software program (preferably on a trustworthy platform) or a chip card reader, a user generates a (virtual) localized identification number for a credential relying party from the issued identification number, the secret key, and the 'registered' identifier (e.g. a business name) of the credential relying party. For verification, the relying party forwards its registered identifier, the localized SSN, and the user's name and address to the issuing party. From name and address, the issuing party can uniquely index or identify the user, re-create a localized SSN, and verify whether the supplied SSN is valid (i.e. was created with the 'right' key). In essence, the proposal turns fixed (long-term) ID numbers into secrets that can be verified, but not reused across relying parties.

Threat model and notation. ID numbers are compromised in many different ways, including data breaches (through compromised merchants' databases or data-centers, and lost/stolen disks or backup tapes), phishing attacks, *dumpster diving*, corrupt insiders, workplace, theft of purses, wallets, or postal mails, social engineering, and existing/past relationship with victims. If a user's physical card is stolen or lost, valid localized numbers may be generated and used unless the user promptly reports the incident to card issuing parties (or the card is protected otherwise, e.g., through a traditional PIN). We focus on preventing identity fraud from large-scale data breaches, instead of attacks that are not much scalable (i.e. difficult to carry out in a comprehensive fashion). We primarily consider breaches of personal ID numbers that can be directly used to perpetrate identity fraud; breaches of other sensitive information, e.g., records of a person's health and education, business secrets (which are also commonly exposed), although important, are out of our scope. We assume that ID number issuing parties can be relied on to protect their customer credentials. User data is breached mostly from relying parties as ID numbers issued by one entity is generally used (and thus replicated) by many relying parties. Such replications increase the possibility of a breach. The following notation is used:

I, U, R	Issuer, user, and relying party respectively.
U_F	User's long time fixed ID number (issued by I).
U_R	User's localized ID number for R.
K_{IU}	Long-term secret key shared between I and U.
$f_{K_{IU}}(\cdot)$	A cryptographically secure MAC function f, keyed by K_{IU}.[5]
U_A	Name and address of U.

Detailed Steps. The steps required for ID number localization are as follows (see also Figure 1).

1. The credential issuer (I) provides a smart card to U with an ID number U_F (unique in I's domain), and a secret key K_{IU} upon verifying U's identity (e.g. through an in-person visit or equivalent). U_F is directly used only with I, and only I and U know U_F and K_{IU}. Additionally, I also keeps U's name and address U_A associated with U_F and K_{IU}.
2. In response to a relying party R's request, U generates a localized ID U_R for R.

$$U_R = f_{K_{IU}}(U_F, R) \tag{2.1}$$

   U sends U_R and U_A to R. The MAC output may require modifications to conform with the target ID format. For an on-site (card-present) transaction, U_R is generated using U's chip-card at R's chip-card

[4] The printed 'secret' key is used when a chip card or card reader is unavailable (variant 2 in Section 3); see item 1 under "Assumptions."

[5] To be more precise, $f(\cdot)$ should be a Pseudo-Random Function (PRF), as similarly used in "independent OTP" [36] and PwdHash [35].

reader (e.g. simply by 'swiping' the card). The reader provides the relying party's name[6] to the card for computing U_R; U does not input anything explicitly. For card-not-present (e.g. web) transactions, U may input the relying party's name to her chip-card reader. (See variant 2 in Section 3 for localized ID generation without a chip-card or chip-card reader.)

3. To verify the validity of U_R (i.e. whether U_R has been generated by using K_{IU} and U_F), R sends (U_R, U_A, R) to I.
4. Using name and address from U_A, I locates U_F and K_{IU}, and checks the validity of U_R; i.e., from U_F, K_{IU} and R, I generates U_R as in equation (2.1) and compares it with the received U_R. I then sends the verification result (accept/reject) to R. Appropriate integrity must of course be provided in this latter communication.

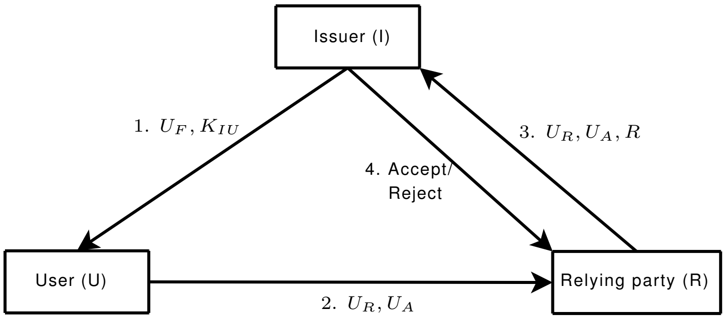

Figure 1: Steps in ID number localization

Assumptions. Operational assumptions in our main proposal and its variants (see Section 3) are as follows.

1. We assume a user does not reveal the printed long-term key on the card to any third parties (e.g. through phishing attacks). If chip-cards are used and users generate ID numbers only through an available chip-card reader, printing the secret key on the card can be avoided.
2. For a variant of our proposal (variant 2 in Section 3), we use a user's personal device (cellphone or PC). Such a device may expose the long-term user key if it contains malware. However, we focus on large scale data breaches, rather than individual information leaks (through malware, phishing, or shoulder-surfing).
3. In our main proposal, users must keep their name and address (U_A) updated with an ID issuing party. For variants 1, 3, and 4 (Section 3), this assumption may be relaxed. Arguably it is impractical to expect users to notify all their ID issuing parties of address changes. However, 'secondary' issuing parties may update U_A from 'primary' parties that are generally expected to have the most recent U_A information, e.g., banks, credit bureaus.
4. We assume that a localized ID number as generated in equation (2.1) is tied to a particular relying party, and can be reused at the same relying party but not anywhere else. This assumption allows *traceability*,[7] and apparently increases usability by requiring less user input (cf. [27]). However, our generalized proposal can be extended to generate more restricted ID numbers (even transaction-specific numbers), e.g., by including timestamp, validity period, transaction amount, etc. along with the name of a relying party (R) in equation (2.1). Such an extension may restrict insider abuse, and reuse of compromised IDs even at the same relying party from where the breach occurred.
5. In our localized ID scheme, an issuing party is directly able to keep track of the usage of a customer's ID. However, information aggregation by a centralized entity (e.g. credit reporting agencies, personal background check for law enforcement) from multiple sources is no longer straightforward under our proposal (due to unlinkability among different custom ID numbers). To achieve such aggregation,

[6] Additional relevant information may be provided as well; see item 4 under "Assumptions".

[7] Note however that, as per the current agreement between U.S. retailers and credit card companies (e.g. Visa and MasterCard), a merchant's identity may not be revealed even when the merchant is responsible for a data breach (see e.g. [28].)

we assume that ID issuing parties will collaborate when required/appropriate, for example, if compelled by law enforcement authorities. Note that for variants 1, 3, and 4 (Section 3), aggregation remains unaffected.

In contrast to many current uses of "identity information," in our proposal, verifying that identity information is "valid" involves the relying party carrying out a communication with the issuing party. This is part of the price we pay for the added security.

3 Variants

Here we discuss four variants of our main proposal. These variants are outlined to initiate further discussion, and for now we defer an in-depth analysis of implementation details, deployment strategy and associated costs, though critically important for rolling out any of these variants.

Variant 1: Localized authorization code. The localized ID scheme above uses U_R in place of U_F. This requires certain formatting of the MAC output. For example, in a regular credit card number, the first six digits identify the issuing bank and the last digit is the Luhn check digit. If U_R is used as a credit card number, it must conform to these restrictions (which may complicate U_R generation, depending on the ID number space of U_F). However, as such a number is identical to a real credit card number, it can be used in the existing infrastructure. An alternative approach is as follows: require the use of U_F along with U_R for a transaction, i.e., now U_R is used as a dynamic *authorization* code (cf. Card Verification Value 2 (CVV2) codes for credit cards [27]) accompanying the fixed ID number. By policy, U_F must not be accepted without a valid U_R. Now U_R need not conform to any strict formats. Existing implementations must still be changed to accommodate the extra authorization code check, but changes to many existing implementations would likely be significantly reduced, for example, databases which are indexed by U_F; this allows straightforward information aggregation from multiple sources. Theft of an U_F (or even an U_F, U_R pair) is no longer a concern, as for generating a new U_R attackers also require the key K_{IU}.

Variant 2: Without chip-card or card-reader. Some credential issuers may not adopt chip cards in the near future. Some ID cards do not even contain a magnetic stripe for storing extra or sensitive information. For example, Canadian Social Insurance Number (SIN) cards and (older) health cards contain only a user's name and ID number in a printed form. Our approach can be used in such cases if users are issued long-term secret keys (perhaps printed on the ID card itself). One-time ID numbers can be generated from a user's fixed ID number and the shared secret, using a personal computing device (e.g. a PC or cellphone containing an appropriate application). If such numbers are generated only infrequently, usability (e.g. having access to a computing device, providing user input) may not be affected much.

For frequently used ID numbers such as credit card numbers, we assume the availability of chip cards with on-site card readers. For card-not-present transactions (e.g. e-commerce), it would be easier for a user if she has access to a chip-card reader (e.g. the user can avoid typing in the secret key). However, users can still use a personal device (with appropriate software on it) for generating localized ID numbers.

Variant 3: Database poisoning. Organizations storing a large number of personal records can create legitimate looking fake records, and insert those into their databases. An issuing party may share a unique secret key with each relying party, and the relying party creates fake records using the shared secret such that the fake records are indistinguishable to an attacker (without knowing the secret key), but the verification party can detect those as fabricated and linked to a particular relying party. The proportion of fake records can be configured to the sensitivity of stored information, storage/computation overhead, and/or company policy. For example, if a compromised database contains 1% fake records, on average, the breach is detected within 100 transaction attempts (i.e. the use of compromised records).

If this technique is implemented by all relying parties of a particular ID number, it will enable the ID issuer to distinguish the compromised relying party from a specific fake record during verification. However, the issuer must assign each fake ID number such that the number is attached to a specific relying party, and cannot be generated by a relying party without the assistance from the issuing party. Also, the issuer must require (through e.g. policy) that all relying parties insert fake records consistently. Satisfying these assumptions in practice could be difficult considering current compliance failures (e.g. [9]). However, this technique is apparently easy for ID issuing parties to implement (i.e. only requires self-compliance). Relying parties may also benefit from database poisoning by reducing their long-term liabilities due to breaches; according to one analysis [49], most organizations currently remain unaware of a compromise for months (63% of cases) and even years (2% of cases).

Responses to a fake record detection may vary depending on cost-benefit trade-offs, e.g., heightened scrutiny of incoming requests, activating additional verification processes, or deactivating the legitimate ID number temporarily or permanently (blocking new uses). To its advantage, this variant does not require any assistance from users (i.e. usability cost is non-existent), at the cost of increased background overheads.

Similar deceptive techniques have been in use for protecting postal mail addresses for a long time [47]. Inserting *honeytokens* [42] (bogus digital records) has been discussed for monitoring unauthorized access/use of different types of digital resources including databases with sensitive personal information. In the security literature, Kursawe and Katzenbeisser [24] discussed similar deceptive techniques to detect compromised personal records (e.g. a credit card number) from a user PC. For example, a user may store one valid credit card number along with 100 other legitimate-looking (but fake) card numbers. When these numbers are compromised, such an incident may be promptly detected by monitoring for the use of fake numbers. `HoneyIM` [52] uses 'decoy' Instant Messaging (IM) contacts for detecting IM worms in an enterprise environment. When a worm attempts to spread by sending its copy to every contact of a compromised account, the worm infected PC can be easily tracked by monitoring the decoy account.

Variant 4: User-centric authorization. In this variant, we propose to actively engage users in blocking critical misuses of breached data records, e.g., issuing of new credentials, transfer/sharing of existing credentials from one party to another, and high-value transactions. Assume that a user registers her personal device with each ID issuing party. When a relying party attempts to verify the user's ID with the corresponding issuer, the issuer notifies the user's personal device (through, e.g., phone call, SMS, email). The issuer may depend on the response from the user device to respond to the relying party, or simply keep the user device informed (i.e. in terms of 'log' messages). To counter automatic approval from a malware-infected device, *physical presence* mechanisms (e.g. a hardware switch, vertical/horizontal shaking) of Trusted Platform Module (TPM)-enabled devices [46] may be used.

Several techniques involving personal devices have been proposed in the recent past (possibly due to the increased proliferation of mobile phones, blackberries, PDAs). In contrast to `CROO` [29], this variant requires only critical transactions involving ID numbers to be verified through the personal device (not every transactions). Unlike "owner-controlled information" [14], the user is not expected to store and maintain all her privacy-sensitive information. The use of a personal device has also been proposed [48] as a "heartbeat locator" (securely detecting the location of the device) to counter identity fraud. A verification center continuously tracks the location of a registered device, and compares the location information with that of an attempted transaction before approving the transaction. A transaction may fail if the locations are not matched. This technique merely requires a user to keep her personal device with (or around) her, and the user does not need to interact with the device for approving a transaction. Issues related to device theft and cloning have also been discussed [48].

Note that our main proposal and variant 2 prevent ID number reuse across relying parties although a compromised ID may remain valid within the breached party's domain. Variant 1 is a prevention mechanism while variant 3 is detection-only. Variant 4 can prevent misuse if explicit user authorization is always required; otherwise, it becomes detection-only.

4 Related Work and Data Breach Incidents

Here we briefly discuss a few high-profile data breach incidents, and academic proposals related to our work.

4.1 Examples and Costs of Data Breaches

Examples of breaches are easily cited. Personal records of all 25 million child benefit recipients in the U.K., including their dates of birth, bank accounts, and national insurance numbers had been lost from a government agency when the agency mailed the records in discs [18]. Sensitive personal information on 26.5 million U.S. veterans had been reportedly stolen [8]. While the TJX data breach [30] is still fresh (affecting about 45 million users), millions of user records were stolen from the `Monster.com` job site [4]. A database admin reportedly [7] stole and sold 8.4 million customer records containing bank account and credit card information; another employee at the same company previously compromised 2.3 million records [39]. The theft of a computer with thousands of 'top-secret' mobile phone numbers, information regarding undercover terrorism and organized crime investigations was reported by a U.K. company [21]. A list of prominent data breaches in the U.S. from Jan. 2005 to July 16, 2008 reports [33] the exposure of more than 233 million records containing sensitive personal information.[8]

Erickson and Howard [11] analyzed news accounts of data breaches from 1980 to 2006, and identified *organizational mismanagement* as one prime reason for these breaches. Considering the incidents from 2005 and 2006, when most U.S. states legislated mandatory reporting, they found that in 68% of news stories concerning data theft, the theft could be attributed to organizational behaviour (e.g. administrative error, insider abuse). Apparently, even if we could 'remove' malicious outsiders (e.g. organized crime) as an element in data breaches (e.g. through security technology), data records with sensitive personal information will still be breached in large numbers.

A Ponemon Institute benchmark study [32] investigates the costs of a data breach using data from 35 U.S. organizations for the year 2007. On average, it costs an organization \$197 per record compromised, an increase of 8% since 2006 (for financial services firms, the cost is \$239 per record). The cost of lost business due to a breach (from the loss of existing customers, and diminished new customers) is estimated on average \$128 per record (a 30% increase from 2006). Acquisti et al. [1] provides a comprehensive analysis (using data from 1999 to 2006) of the impacts of a privacy breach incident on a company's stock market value; these effects are generally negative and statistically significant in the short term, but not so visible in the long run. Costs to consumers affected by a data breach is even more difficult to estimate. According to one estimation,[9] in 2007, the average fraud amount per ID fraud victim in the U.S. is \$5720 (about 9% decrease from 2006). However, in most cases, it is difficult to clearly establish a link between breached data and fraud [15]. This fact is also often exploited by breaching parties to understate their legal responsibilities.

General observations. Large-scale data breaches occur frequently, and current legal and technical measures are failing to slow down this trend. Costs of these breaches are significant for both consumers and corporations. Establishing a concrete link between data breaches and identity fraud is often difficult because misuse may occur long after a breach, and misused information cannot directly be attributed to a particular breach incident. However, the growing underground (criminal) market for stolen personal information strongly suggests that breached ID numbers can be easily sold and exploited [43].

4.2 Related Work and Comparison

In NSPW 2007, Beaumont-Gay et al. [5] proposed a policy-based solution called Data Tethers where enforced policies are dependent on the operating environment; i.e. access control policies for stored data on a computing device differ depending on whether the device is inside a *secure* environment or otherwise. Data Tethers may encrypt or remove sensitive data when an insecure environment (e.g. stolen laptop) is detected. This technique assumes the existence of an "actually secure" computing environment, and that Data Tethers policies will always be flawlessly enforced. While such techniques can substantially improve data security in certain environments, in general, we believe the most prudent assumption is that data will be compromised, irrespective of protection mechanisms deployed to prevent such leaks.[10] Also, in many cases it is only realistic to acknowledge that the prevention of compromise is beyond the control of end-users, and of relying parties who hold such information.

Gates and Slonim [14] introduced the owner-controlled information paradigm to address the issues of "privacy, consistency and mobility" in regard to personal information. Users are expected to maintain all of their personal information, identification information, as well as medical history and financial information using a personal device. Organizations must contact a user directly to collect and use personal information. Although this technique provides greater control over a user's sensitive data, it apparently comes with several unique challenges and high usability costs (some of which have been discussed in the paper, e.g.,

[8] `Attrition.org` and Identity Theft Resource Center [23] also maintain similar but independent lists of data breaches. Verizon [49] provides a comprehensive analysis (including breach sources, attack types and paths, time span of breach events) of 500 such data breach cases from 2004 to 2007; see also [20].

[9] Javelin Strategy and Research Survey (Feb. 2007); for excerpts see `http://www.privacyrights.org/ar/idtheftsurveys.htm`.

[10] For example, several breach incidents have been reported from the U.S Department of Defense [17], presumably one of the most security-aware organizations.

lost/stolen device, unauthorized access, backup and recovery). Ashley et al. [2] propose a framework to addresses "privacy management" (e.g. publishing concrete privacy promises, user consent management, privacy enforcement, auditing) for collected customer data in an enterprise environment. This framework may provide higher level of privacy assurance, although it may incur significant costs (in addition to requiring an enterprise to develop a comprehensive privacy policy, and to enforce that policy honestly and consistently).

To reduce customers' fear of using credit cards online (i.e. for card-not-present transactions), several banks enable users to generate limited-use (e.g. one-time) card numbers through their websites. These dynamically generated card numbers are tied to a user's fixed credit card, and can be used for online purchases instead of the fixed card itself. Examples from real-world deployments of such schemes include American Express' Private Payments,[11] Discover card's Secure Online Account Numbers,[12] and SecureClick [40].

Rubin and Wright [37] proposed an *offline* scheme for generating limited-use credit card numbers (i.e. without requiring direct interaction between a user and card issuer for generating new numbers). A user (U) and a card issuer (I) share a long-term secret key (K). U possess a computing device, and stores K on it. To generate a new credit card number number, U selects a monetary restriction (e.g. \$100 limit), expense category (e.g. food), limited validity period, merchant name, timestamp etc. and encrypts these restrictions using K (in an arbitrary finite domain encryption scheme [6]). U then transmits the newly generated limited-use number and her identifying information (e.g. name and address) to the card issuer via the merchant. From the identifying information, I selects K and verifies the limited-use number (e.g. checks the restrictions).

Assuming the availability of chip-cards and chip-card readers, Li and Zhang [25] (see also [26]) proposed a one-time credit card scheme with limited involvement of a user (i.e. no transaction specific user inputs) for card-present (on-site) and card-not-present (web, and phone/fax/email) payment scenarios. A user generates one-time use numbers simply by inserting her credit card into a chip-card reader. In this scheme, a credit card stores a secret value (S) and an initial one-time credit card transaction number (CCT). Assuming T_{cur} is the current CCT number, the next CCT T_{new} is generated by hashing (T_{cur}, S). At the end of the current transaction, T_{new} replaces T_{cur} on the card.

Using a personal device such as a cellphone, Molloy et al. [27] proposed an offline scheme for generating virtual credit card numbers similar to Rubin and Wright [37]. Instead of using finite domain encryption, Molloy et al. used a MAC to avoid several limitations of such encryption, e.g., encoding merchant names in a compact format. Also, the long-term shared key between a user and card issuer is a user-memorable password P (generally 'weak'). The MAC key is generated from the hashed output of P and the user's real credit card number as assigned by the issuer. A *transaction string* (including U_A, expiration date, R, transaction amount) is MACed to generate the virtual credit card number. The authors claim the *forgery resistant* property, i.e., an attacker cannot (easily) forge credit card numbers even if he knows the user's real credit card number and some virtual credit card transactions. However, this property relies on the assumption that a user-chosen password is *strong* (i.e. has high entropy), whereas in practice user-chosen passwords are often weak, and may be discovered easily through a dictionary attack.

The main appeal of using disposable credit card numbers is apparently to alleviate the inconvenience of customers contacting their bank (and replacing a compromised card), as users are typically liable for at most \$50 in case of fraudulent use of their credit card. Generally, credit card numbers alone cannot be used for identity fraud. On the other hand, efforts to reduce misuse of more sensitive information such as SSN are apparently scarce (see e.g. limiting the use of SSNs as an identifier [10], and the FTC workshop [13]; see also [16]).

Similar to our proposal, `CROO` [29] attempts to address the generic identity fraud problem albeit by a different use of one-time passwords; for example, CROO is more complex and seeks to secure individual transactions, whereas we focus on securing ID numbers.

From a legal perspective, Solove [41] identifies several inadequacies in the traditional model for addressing privacy violations using ID theft as an example. ID theft is a "consequence of an architecture" [41] exploited by ID thieves (e.g. the use of SSNs for indexing a large array of sensitive personal information held by government agencies and private businesses). A new architecture has been proposed based on the Fair Information Practices (originating from a 1973 report by the U.S. Department of Housing, Education, and

[11] Introduced in Oct. 2000, discontinued since Oct. 2004.

[12] http://www2.discovercard.com/deskshop. Orbiscom's (`orbiscom.com`) Controlled Payment Numbers technology enables Discover and several other one-time disposable credit card providers.

Welfare). The Fair Information Practices focus on increasing an individual's involvement (e.g. participation in the collection, storage and use) in personal information systems. As an example mechanism, Solove [41] proposes that user-chosen passwords or account numbers be used for accessing credit reports instead of using SSNs or other sensitive personal information.

Partly driven by increasing public demand, most U.S. states (44 out of 50, as of July 2008) have legislated data breach notification laws, requiring organizations to report breach incidents to a state agency. While the question of whether these laws will reduce data theft in the long run is yet to be answered, it has been reported that so far their effect appears to be statistically insignificant [34]. However, another study [38] reported that notification laws are increasing "awareness of the importance of information security" among organizations surveyed. Payton [31] provides a review of current U.S. state and federal laws regarding data breaches, and possible legal remedies available to fraud victims. Costs and benefits of a national data breach notification requirement have also been analyzed [15].

Some businesses attempt to prevent identity theft by providing a service which places *fraud alerts* on a customer's credit bureau profiles. However, in one incident [50], the identity of the CEO of such a company was exploited to obtain a $500 loan (using the CEO's SSN which is displayed publicly on the company website and TV commercials).

Advantages of our proposal. Advantages of our proposal relative to existing ones include the following.

1. A localized ID number as generated in our proposal is bound uniquely to the relying party. While this does not offer the advanced restrictions of Rubin and Wright [37], their additional restrictions require additional user input; thus we expect that our simpler proposal may enjoy better usability.
2. While Li and Zhang [25] assume the availability of user-level chip-card readers, our proposal (variant 2) can work when the user has access to a wide variety of computing devices (e.g. cellphone, PC).
3. Our proposal is computationally immune to offline dictionary attacks as it does not rely on user-chosen passwords (in contrast to Molloy et al. [27]).
4. Our proposal may also limit *synthetic ID theft* [22] where imposters use real identifiers (e.g. SSN) along with other fake attributes, e.g., name, address.

5 Concluding Remarks

Once personal ID numbers are collected by third parties, we believe that the most prudent assumption in today's Internet environment is that they will be breached at some point in time, despite best efforts (if any) of the collecting parties. In addition to lost personal privacy (e.g. medical history, purchase habits, online and real-world activities under surveillance), these breaches enable large-scale identity fraud. Some of these fraudulent activities remain undetected by their victims for years [12, 49]. While direct monetary losses for consumers from such fraud are recoverable to some extent, nonmonetary damages (productivity/time lost to resolve identity theft [44], denied credit or other financial services, harassment by debt collection agencies, criminal investigation or arrest [3]) are not; see e.g. the FTC 2003 report [12]. One of the main problems is that agencies/corporations responsible for these breaches of customer records are not generally held accountable for the breaches, and presently there is no significant financial penalty. We expect that if it was corporate data that was being compromised, corporations would pursue legal remedies; but since it is primarily the personal information of individuals, and the perceived dollar amount likely to be gained through legal remedy is small compared to the cost of litigation, individuals generally do not pursue legal remedies.

We outline an ID number localization approach and schemes to reduce identity fraud due to large-scale data breaches that expose reusable fixed ID numbers. Rather than focusing on analysis of a particular solution, our proposed variants are intended to initiate further discussion on how to better address the current problem of identity fraud resulting from breached databases of personal information on millions of customers. There are certainly deployment challenges with several of our proposals; consequently, we raise the question, "Are there better proposals that can address the same problem?" We believe that fundamentally new approaches are required to address this problem, which clearly is not addressed by existing solutions.

Acknowledgements. We thank anonymous NSPW 2008 reviewers for their comments and suggestions, and our NSPW shepherd Brian Snow. The first author is supported in part by an NSERC CGS. The second author is Canada Research Chair in Network and Software Security, and is supported in part by an NSERC Discovery Grant, and the Canada Research Chairs Program.

References

[1] A. Acquisti, A. Friedman, and R. Telang. Is there a cost to privacy breaches? An event study. In *International Conference of Information Systems (ICIS)*, Dec. 2006.

[2] P. Ashley, C. Powers, and M. Schunter. From privacy promises to privacy management: A new approach for enforcing privacy throughout an enterprise. In *New Security Paradigms Workshop (NSPW)*, Virginia Beach, VA, USA, Sept. 2002.

[3] BBC News. 'I was falsely branded a paedophile'. News article (Apr. 3, 2008). `http://news.bbc.co.uk/1/hi/magazine/7326736.stm`.

[4] BBC News. Monster attack steals user data. News article (Aug. 21, 2007). `http://news.bbc.co.uk/2/hi/technology/6956349.stm`.

[5] M. Beaumont-Gay, K. Eustice, and P. Reiher. Information protection via environmental data tethers. In *New Security Paradigms Workshop (NSPW)*, New Hampshire, USA, Sept. 2007.

[6] J. Black and P. Rogaway. Ciphers with arbitrary finite domains. In *RSA Security (CT-RSA)*, Feb. 2002.

[7] ChannelRegister.co.uk. IT pro admits stealing 8.4M consumer records. News article (Dec. 4, 2007). `http://www.channelregister.co.uk/2007/12/04/admin_steals_consumer_records/`.

[8] CNN.com. FBI seeks stolen personal data on 26 million vets. News article (May 23, 2006). `http://www.cnn.com/2006/US/05/22/vets.data/index.html`.

[9] ComputerWorld.com. TJX violated nine of 12 PCI controls at time of breach, court filings say. News article (Oct. 26, 2007).

[10] Electronic Privacy Information Center (EPIC). Protecting the privacy of the Social Security Number from identity theft. EPIC testimony before the Committee on Ways and Means in the U.S. House of Representitives (June 21, 2007). `http://www.epic.org/privacy/ssn/idtheft_test_062107.pdf`.

[11] K. Erickson and P. N. Howard. A case of mistaken identity? News accounts of hacker, consumer, and organizational responsibility for compromised digital records. *Journal of Computer-Mediated Communication*, 12(4), July 2007.

[12] Federal Trade Commission. Identity theft survey report, Sept. 2003. `http://www.ftc.gov/os/2003/09/synovatereport.pdf`.

[13] Federal Trade Commission. Security in numbers: SSNs and ID theft workshop, Dec. 2007. `http://ftc.gov/bcp/workshops/ssn/index.shtml`.

[14] C. Gates and J. Slonim. Owner-controlled information. In *New Security Paradigms Workshop (NSPW)*, Ascona, Switzerland, Aug. 2003.

[15] Government Accountability Office (GAO). Personal information: Data breaches are frequent, but evidence of resulting identity theft is limited; however, the full extent is unknown, June 2007. Report to Congressional Requesters, GAO-07-737, `http://www.gao.gov/new.items/d07935t.pdf`.

[16] Government Accountability Office (GAO). Social Security Numbers: Use is widespread and protection could be improved, June 2007. Testimony Before the Subcommittee on Social Security, Committee on Ways and Means, House of Representatives, GAO-07-1023T, `http://www.gao.gov/new.items/d071023t.pdf`.

[17] Government Reform Committee. Agency data breaches since January 1, 2003. Staff Report, U.S. House of Representatives (Oct. 13, 2006). `http://oversight.house.gov/documents/20061013145352-82231.pdf`.

[18] Guardian.co.uk. Lost in the post - 25 million at risk after data discs go missing. News article (Nov. 21, 2007). `http://www.guardian.co.uk/politics/2007/nov/21/immigrationpolicy.economy3`.

[19] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten. Lest we remember: Cold boot attacks on encryption keys. In *USENIX Security*, San Jose, CA, USA, 2008.

[20] R. Hasan and W. Yurcik. A statistical analysis of disclosed storage security breaches. In *ACM Workshop on Storage Security and Survivability (StorageSS)*, 2006.

[21] Help Net Security (HNS). Server with top-secret data stolen from Forensic Telecommunications Services. News article (Aug. 13, 2007). `http://www.net-security.org/secworld.php?id=5418`.

[22] C. J. Hoofnagle. How SSNs are used to commit ID theft: Synthetic identity theft. In *Security in Numbers: SSNs and ID Theft Workshop*, Dec. 2007. Panel presentation at the workshop, hosted by the FTC.

[23] Identity Theft Resource Center (ITRC). ITRC 2008 breach list. Security breaches from 2005 to 2008 (Apr. 8, 2008). `http://www.idtheftcenter.org/artman2/publish/lib_survey/ITRC_2008_Breach_List.shtml`.

[24] K. Kursawe and S. Katzenbeisser. Computing under occupation. In *New Security Paradigms Workshop (NSPW)*, New Hampshire, USA, Sept. 2007.

[25] Y. Li and X. Zhang. A security-enhanced one-time payment scheme for credit card. In *IEEE Workshop on*

Research Issues on Data Engineering: Web Services for E-Commerce and E-Government Applications (RIDE), Boston, MA, USA, Mar. 2004.

[26] Y. Li and X. Zhang. Securing credit card transactions with one-time payment scheme. *Journal of Electronic Commerce Research and Applications (ECRA)*, 4(4), 2005.

[27] I. Molloy, J. Li, and N. Li. Dynamic virtual credit card numbers. In *Financial Cryptography and Data Security (FC)*, Scarborough, Trinidad and Tobago, Feb. 2007.

[28] MSNBC. Breach exposes 4.2 million credit, debit cards. News article (Mar. 17, 2008). http://www.msnbc.msn.com/id/23678909/.

[29] D. Nali and P. C. van Oorschot. CROO: A universal infrastructure and protocol to detect identity fraud. In *European Symposium on Research in Computer Security (ESORICS)*, Malaga, Spain, Oct. 2008.

[30] News.com. TJX says 45.7 million customer records were compromised. News article (Mar. 29, 2007).

[31] A. M. Payton. Data security breach: Seeking a prescription for adequate remedy. In *ACM Conference on Information Security Curriculum Development (InfoSecCD)*, Kennesaw, Georgia, USA, Sept. 2006.

[32] Ponemon Institute. 2007 annual study: U.S. cost of a data breach, Nov. 2007. Research report sponsored by PGP and Symantec, http://www.pgp.com/downloads/research_reports/.

[33] Privacy Rights Clearing House. A chronology of data breaches. http://www.privacyrights.org/ar/ChronDataBreaches.htm.

[34] S. Romanosky, R. Telang, and A. Acquisti. Do data breach disclosure laws reduce identity theft? In *Workshop on the Economics of Information Security (WEIS)*, June 2008.

[35] B. Ross, C. Jackson, N. Miyake, D. Boneh, and J. C. Mitchell. Stronger password authentication using browser extensions. In *USENIX Security*, 2005.

[36] A. Rubin. Independent one-time passwords. *USENIX Journal of Computing Systems*, 9(1), Feb. 1996.

[37] A. Rubin and R. Wright. Off-line generation of limited-use credit card numbers. In *Financial Cryptography (FC)*, Grand Cayman, British West Indies, Feb. 2001.

[38] Samuelson Law, Technology & Public Policy Clinic, University of California-Berkeley School of Law. Security breach notification laws: Views from chief security officers, Dec. 2007. http://groups.ischool.berkeley.edu/samuelsonclinic/files/cso_study.pdf.

[39] SecurityFocus.com. Employee steals 2.3m records from data firm. News article (July 5, 2007). http://www.securityfocus.com/brief/541.

[40] A. Shamir. SecureClick: A web payment system with disposable credit card numbers. In *Financial Cryptography (FC)*, Grand Cayman, British West Indies, Feb. 2001.

[41] D. J. Solove. Identity theft, privacy, and the architecture of vulnerability. *Hastings Law Journal*, 54, 2003. Available at SSRN: http://ssrn.com/abstract=416740.

[42] L. Spitzner. Honeytokens: The other honeypot. SecurityFocus Infocus technical article (July 17, 2003). http://www.securityfocus.com/infocus/1713.

[43] Symantec.com. Symantec global Internet security threat report: Trends for July - December 07. Published in Apr. 2008 (Volume XIII).

[44] The Arizona Republic. Man's ID theft nightmare takes 2 years to iron out. News article (Feb. 13, 2008). http://www.azcentral.com/community/mesa/articles/0213mr-idtheft0214.html.

[45] W. Treese. Once collected, data isn't private. *netWorker*, 9(3), 2005.

[46] Trusted Computing Group. TCG physical presence interface specification, Apr. 2007. Version 1.00, final revision 1.00, for TPM family 1.2; level 2.

[47] US Monitor. Mail monitoring and list protection service. http://www.usmonitor.com.

[48] P. van Oorschot and S. Stubblebine. Countering identity theft through digital uniqueness, location cross-checking, and funneling. In *Financial Cryptography and Data Security (FC)*, Roseau, Commonwealth of Dominica, 2005.

[49] Verizon Business Risk Team. 2008 data breach investigations report. Analysis of 500 security breach and data compromise engagements between 2004 – 2007. http://www.verizonbusiness.com/resources/security/databreachreport.pdf.

[50] Wired.com. LifeLock founder resigns amid controversy. Wired blog article (June 11, 2007). http://blog.wired.com/27bstroke6/2007/06/lifelock_founde_1.html.

[51] D. Wright, M. Friedewald, W. Schreurs, M. Verlinden, S. Gutwirth, Y. Punie, I. Maghiros, E. Vildjiounaite, and P. Alahuhta. The illusion of security. *Communications of the ACM*, 51(3), Mar. 2008.

[52] M. Xie, Z. Wu, and H. Wang. HoneyIM: Fast detection and suppression of instant messaging malware in enterprise-like networks. In *Annual Computer Security Applications Conference (ACSAC)*, Dec. 2007.