

A Scalable and Secure Cryptographic Service

Shouhuai Xu¹ and Ravi Sandhu²

¹ Department of Computer Science, University of Texas at San Antonio
shxu@cs.utsa.edu

² Institute for Cyber-Security Research, University of Texas at San Antonio
ravi.sandhu@utsa.edu

Abstract. In this paper we present the design of a scalable and secure cryptographic service that can be adopted to support large-scale networked systems, which may require strong authentication from a large population of users. Since the users may not be able to adequately protect their cryptographic credentials, our service leverages some better protected servers to help fulfill such authentication needs. Compared with previous proposals, our service has the following features: (1) it incorporates a 3-factor authentication mechanism, which facilitates compromise detection; (2) it supports immediate revocation of a cryptographic functionality in question; (3) the damage due to the compromise of a server is contained; (4) it is scalable and highly available.

Keywords: cryptographic service, scalability, security, compromise detection, compromise confinement, availability.

1 Introduction

Large-scale networked systems, such as peer-to-peer and grid systems, must be adequately protected; otherwise they may be abused or exploited to do more harm than good — Distributed Denial-of-Service (DDoS) attacks are just an example. An important aspect of secure large-scale networked systems is to enforce strong authentication, which would require a large population of users to utilize some cryptosystems such as digital signatures. Due to the very nature of cryptography, assurance offered by such authentications perhaps cannot be any better than security (or secrecy) of the corresponding cryptographic keys or functionalities. This is because compromise of a cryptographic key would allow the adversary to perfectly impersonate the victim user. The threat is amplified by the fact that average users often do not have the expertise or skill to secure their own computers, which may be justified by the fact that there have been many botnets that consist of many compromised computers.

Our contributions. We present the design of a scalable and secure cryptographic service that can be adopted to support large-scale networked systems, which require strong authentication from a large population of users. Specifically, our approach leverages some better protected servers to help protect the

users' private signing keys and functionalities. Compared with a standard two-party threshold digital signature system (i.e., a user's private key is split into two shares such as one is stored on the user's machine and the other is stored on a remote server) and previous proposals for a similar purpose, our service has the following features:

- * It incorporates a 3-factor authentication mechanism so that a signature may be produced in a certain way when a request: (1) presents a valid password (i.e., what you know); (2) is initiated from a party having access to the user's soft-token (i.e., what you have); (3) presents a valid fresh one-time secret (i.e., whether you *always* have access to the soft-token). As a result, the resulting service provides a *compromise detection* capability that may be of independent value.
- * It supports a convenient *key disabling* that can be done using a standard username/password authentication. This is useful, for instance, when a user's device is stolen on a business trip because the user can disable its private key without having access to a backup of the content on its stolen device.
- * The damage due to the compromise of a server is confined to a subset of the users subscribing to its service. Furthermore, the users associated with a compromised server do not have to re-initialize their private keys, unless they suspect that their own machines might have been compromised.
- * It is scalable due to its "decentralized" nature (i.e., each server may serve a subset of users). It is highly available since a single server is enough to help a user fulfill its task. While we do not explore the details of utilizing the state machine approach [28] to securely replicate a server, it should not be difficult to extend our solution to fulfill such replications. In particular, in a related prior scheme we proposed [29], threshold cryptosystems have been adopted to fulfill distributed password-based authentication and signing.

Related prior work. The simplest approach to securing cryptographic keys is to let each user utilize some tamper-resistant hardware device. The industry has started to provide machines equipped with Trusted Platform Module (TPM) as specified by the Trusted Computing Group (www.trustedcomputinggroup.com). However, there are many legacy computers that need be better protected, meaning that alternate solutions are still useful.

One alternate approach to protecting cryptographic keys is to encrypt a key with a password; this is indeed widely deployed in real-life systems. However, the resulting security assurance is quite weak because, once a computer is compromised, the adversary can obtain the cryptographic keys without even conducting an off-line dictionary attack. Another approach is to let a user store its cryptographic key in a remote server, and download the key from the server after a password-based authentication [24]. This approach still allows the adversary, who can compromise a user's computer, to obtain the private key in question. Moreover, the remote server has to be trusted.

Our scheme follows the paradigm of "cryptography as a service" [13]. (This paradigm is different from the server-aided protocols [23,2], which were

motivated to utilize a computationally powerful server to help a computationally poor device conduct expensive cryptographic computations.) In particular, we adopt as a starting point the proposal due to MacKenzie and Reiter [22], which follows [10,16]. The basic idea common to [22,10,16] as well as a similar result [9] is to split a private key into two shares such that one share (often called “soft-token”) is stored on the user’s device, and the other is stored on a remote server. These schemes have no single point of failure.

Finally, we should mention that there are some interesting cryptographic constructs that aim to protect private keys. Notable results include forward-security signatures [1,3,19], key-insulated signatures [15] and intrusion-resilient signatures [20]. The protections provided by these mechanisms are orthogonal to the protection provided by our approach. Nevertheless, they may be integrated together for a better protection.

Outline. In Section 2 we briefly review some necessary cryptographic preliminaries. In Section 3 we introduce the soft-token system model. In Section 4 we present a building block, which is utilized in Section 5 to construct the full-fledged service scheme. We conclude the paper in Section 6. Due to space limitation, the proofs of some theorems are deferred to the full version of the present paper [30].

2 Cryptographic Preliminaries

Let k be the primary security parameter (e.g., $k = 160$), and λ be a secondary security parameter for public keys (e.g., $\lambda = 1024$ means we use 1024-bit RSA moduli). A function $\epsilon : \mathbb{N} \rightarrow \mathbb{R}^+$ is negligible if for any c there exists k_c such that $\forall k > k_c$ we have $\epsilon(k) < 1/k^c$. Let H (with an additional subscript as needed) be a hash function that, unless otherwise stated, is assumed to behave like a random oracle [5] with range $\{0, 1\}^k$.

Pseudorandom functions. A pseudorandom function (PRF) family $\{f_v\}$ parameterized by a secret value v has the following property [17]: It is computationally infeasible to distinguish f_v , where v is uniformly chosen at random, from a random function (with the same domain and range).

Message authentication codes. We assume the standard property of message authentication codes (MACs): If the key a is unknown, then given multiple pairs $\langle m_i, MAC_a(m_i) \rangle$ where the m_i ’s may be adaptively chosen, it is computationally infeasible to compute any pair $\langle m, MAC_a(m) \rangle$ where $m \neq m_i$.

Public key cryptosystems. A public key cryptosystem \mathcal{E} is a triple (GEN, E, D) of polynomial-time algorithms, where the first two are probabilistic. GEN , taking as input 1^λ , outputs a key pair (pk, sk) . E , taking as input a public key pk and a message m , outputs an encryption c for m . D , taking as input a ciphertext c and a private key sk , returns a message m when c is valid and

\perp otherwise. We assume an encryption scheme that is secure against adaptive chosen-ciphertext attack [26]. Basically, an attacker \mathcal{A} is given pk and allowed to query the decryption oracle. At some point \mathcal{A} generates two equal length strings X_0 and X_1 and sends them to a test oracle, which chooses $b \in_R \{0, 1\}$ and returns $Y = E_{pk}(X_b)$. Then \mathcal{A} continues querying the decryption oracle, with the restriction that it cannot query the decryption of Y . Finally, \mathcal{A} outputs b' . We say \mathcal{A} succeeds if $b = b'$. Practical schemes are available [6,12].

Digital signature schemes. A digital signature scheme \mathcal{S} is a triple (GEN, S, V) of polynomial-time algorithms, where the first two are probabilistic. GEN , taking as input 1^λ , outputs a key pair (pk, sk) . S , taking as input a message m and a private key sk , outputs a signature σ for m . V , taking as input a message m , a public key pk , and a candidate signature σ , returns $b = 1$ if σ is a valid signature for m and $b = 0$ otherwise. We assume a signature scheme that is existentially unforgeable under adaptive chosen-message attack [18]: a forger is given pk ; it is allowed to query a signature oracle on messages of its choice; it succeeds if it outputs a valid signature for m that is not one of the messages signed before.

3 Model and Goals

System model. There are a set of users and a set of semi-trusted servers. As in a standard Public Key Infrastructure (PKI), a server has a pair of public and private keys, and so does a user. All the users and servers are probabilistic polynomial-time algorithms. A user splits its private key into two shares after a cryptographic transformation such that one share is stored on the user side (perhaps being encrypted with a password) and the other is stored on a remote server. A soft-token is a data structure a user stores. A soft-token (containing a user-side key share) may be stateful, so we may denote by $token^{(i)}$ the soft-token after the i^{th} transaction in which it is utilized.

A server has two interfaces: one for producing signatures and the other for disabling private keys. The resulting signatures can be verified using the public keys of the users, and thus can be used for authenticating the users in higher-layer applications. In order for a user to produce a signature, the user conducts an interaction with a server, which collaborates with the claimed user only when the user successfully authenticates itself to the server (not to the higher-layer application). In order for a user to disable its private key, the user needs to succeed in a certain authentication operation. A server maintains a database for recording relevant information that would allow the service provider to take actions (e.g., gathering payment when the service is payment-based).

Adversary. We consider an adversary who may have control over the network. The adversary may compromise certain resources including a user's soft-token, a user's password, and a server's private key. The adversary may break into a user's device when the client software is *active*; this explains why we consider an adversary that is strictly more powerful than the adversary considered in

previous soft-token systems. We assume that the integrity of the server (e.g., the database) is guaranteed, even if an adversary can compromise the server's private key. This also models the situation where a semi-trusted server may be honest in performing the protocol, but curious about the users' private keys.

Goals. Recall that k is the primary security parameter. A cryptographic service should have the following properties:

- * **Abuse prevention.** Consider a fixed pair of $\langle user, server \rangle$, where the *user* possesses a soft-token *token* and a password *pwd*. Denote by $ADV(R)$ the type of adversary who succeeds in capturing the resource elements of $R \subseteq \{token, server, pwd\}$. When we say that an adversary ADV has access to *token* we mean that *token* is always available to ADV ; when we say that ADV has access to $\neg token^{(i)}$ we mean that ADV does not have access to $token^{(i)}$ but perhaps has access to $token^{(j)}$ for $0 \leq j < i$. Specifically,
 1. An adversary of type $ADV\{server, pwd\}$ can only forge signatures that are valid with respect to the user's public key with a negligible probability in k .
 2. An adversary of type $ADV\{token, server\}$ can forge signatures that are valid with respect to the user's public key only when the adversary succeeds in off-line dictionary guessing the user's password.
 3. An adversary of type $ADV\{token\}$ can forge signatures that are valid with respect to the user's public key with probability negligibly more than $q/|\mathbb{D}|$ after q invocations of the server, where \mathbb{D} is the dictionary from which the user's password is randomly drawn.
 4. An adversary of type $ADV\{token, pwd\}$ can output — with only a probability negligible in k — signatures that are valid with respect to the user's public key after the user's private key is disabled.
 5. An adversary of type $ADV\{\neg token^{(i)}, pwd\}$ can output — with only a probability negligible in k — signatures that are valid with respect to the user's public key after the user finishes the i^{th} transaction and before the user initiates the $(i + 1)^{th}$ transaction.
- * **Compromise detection.** The system itself can detect the compromise that an adversary has succeeded in impersonating the user for producing signatures.
- * **Immediate revocation.** A user can request a server to disable its private key by executing a standard username/password authentication.
- * **Compromise confinement.** The impact due to the compromise of a server is contained to a *subset* of the users subscribing to its service. Moreover, these users do not have to re-initialize their private keys.
- * **Scalability.** The system can serve a large population of users.
- * **High availability.** The system is highly available, even if some servers are under DDoS attacks.

4 Building Block: A Single Server Soft-Token Scheme

In this section we present a building block, which is extended from a scheme presented in [22] and will be incorporated into our full-fledged scheme in Section 5.

Suppose the server’s public data (e.g., public key) are available to the users, and consider a fixed pair of $\langle user, server \rangle$. The user runs the initialization process to generate a soft-token $token^{(0)}$ using its public and secret data as well as the server’s public data. In the i^{th} transaction ($i = 1, 2, \dots$), the user, who has access to $token^{(i-1)}$ that is generated in the $(i - 1)^{th}$ transaction or in the initialization process when $i = 1$, signs a message by interacting with the server. The server collaborates with the claimed user only when the user presents the password and the state information ϑ chosen by the server in the last transaction. At the end of the i^{th} transaction, the user obtains the signature, generates a new token $token^{(i)}$ using the cryptographic state information ϑ chosen by the server, and erases $token^{(i-1)}$. The server tracks the changes of the ϑ ’s in its database.

Denote a user’s public key by $pk_{user} = \langle e, N \rangle$ and private key by $sk_{user} = \langle d, N, \phi(N) \rangle$, where $ed = 1 \pmod{\phi(N)}$, N is the product of two large prime numbers, and ϕ is the Euler totient function. In the standard encode-then-sign paradigm, the signature sig on message m is $S_{\langle d, N, \phi(N) \rangle}(m) = \langle r, s \rangle$, where $r \in_R \{0, 1\}^{len_{pad}}$, $s = (\text{encode}(m, r))^d \pmod N$ for some encoding function encode . A signature $\langle r, s \rangle$ can be verified by checking if $s^e = \text{encode}(m, r) \pmod N$. The function encode could be either deterministic (e.g., $len_{pad} = 0$ in the case of hash-and-sign [14]) or probabilistic (e.g., PSS [7]), these types of signatures were proven secure against adaptive chosen-message attacks in the random oracle model. The basic idea for splitting the private key d is to let $d_1 + d_2 = d \pmod{\phi(N)}$. The scheme has the following components.

Token initialization. Suppose $H_1 : \{0, 1\}^* \rightarrow \{0, 1\}^k$ and $f : \{0, 1\}^* \rightarrow \{0, 1\}^{\lambda+k}$. The inputs are the server’s public encryption key pk_{server} , the user’s password pwd , the user’s public key $pk_{user} = \langle e, N \rangle$ and private key $sk_{user} = \langle d, N, \phi(N) \rangle$. The initialization proceeds as follows.

$$\begin{aligned}
 uid &= \text{username} \\
 v &\in_R \{0, 1\}^k \\
 a &\in_R \{0, 1\}^k \\
 b &= H_1(pwd) \\
 d_1 &= f(v, pwd) \\
 d_2 &= d - d_1 \pmod{\phi(N)} \\
 \tau &= E_{pk_{server}}(\langle a, b, uid, d_2, N \rangle) \\
 ct &= 0 \\
 st &\in_R \{0, 1\}^k \\
 token &= (ct, st, v, a, \tau, e, N, pk_{server})
 \end{aligned}$$

The user chooses its own uid , a unique and memorizable string such as its email address. The soft-token is $token = (ct, st, v, a, \tau, e, N, pk_{server})$, where ct is an incremental counter indicating the serial number of a transaction (which is used for simplifying the description), st is the state information that will be chosen by the server (for the time being of $ct = 0$, it is just a placeholder). All the other values, including $b, d, d_1, d_2, \phi(N)$, and pwd , are erased.

Server database. The server maintains a database of $\mathcal{Y} = (\tau, uid, count, \vartheta, m, r)$, where uid is obtained after receiving the first service request, and $count$ is an incremental counter (with initialized value zero). Each record of the database represents a transaction corresponding to $\tau = E_{pk_{server}}(\langle a, b, uid, d_2, N \rangle)$. The database has two operations: $append(\tau, uid, count, \vartheta, m, r)$ for appending a record $(\tau, uid, count, \vartheta, m, r)$ to the database, and $last(\tau, uid, count, \vartheta, m, r)$ for returning either the record $(\tau, uid, count, \vartheta, m, r)$ corresponding to the *last* transaction corresponding to τ or NULL (meaning that the token corresponding to τ has never been used before).

Signing protocol. The client software prompts the user to enter the password pwd , to get the to-be-signed message m as well as the soft-token $token = (ct, st, v, a, \tau, e, N, pk_{server})$. The protocol is depicted in Fig. 1.

USER	SERVER
$token = (ct, st, v, a, \tau, e, N, pk_{server})$ $\beta = H_1(pwd), \rho_1 \in_R \{0, 1\}^k$ $\rho_2 \in_R \{0, 1\}^\lambda, \rho_3 \in_R \{0, 1\}^k$ $r \in_R \{0, 1\}^{len_{pad}}$ $\gamma = E_{pk_{server}}(\langle m, r, \beta, st, \rho_1, \rho_2, \rho_3 \rangle)$ $\delta = MAC_a(\langle \gamma, \tau \rangle)$	
	(γ, τ, δ) \rightarrow
	abort IF τ has been disabled $\langle a, b, uid, d_2, N \rangle = D_{sk_{server}}(\tau)$ abort IF $MAC_a(\langle \gamma, \tau \rangle) \neq \delta$ $\mathcal{Y} = last(\tau, uid, count, \vartheta', m', r')$ $\langle m, r, \beta, st, \rho_1, \rho_2, \rho_3 \rangle = D_{sk_{server}}(\gamma)$ abort IF $\beta \neq b$ abort IF $\mathcal{Y} \neq NULL \wedge st \neq \vartheta'$ $\sigma = (encode(m, r))^{d_2} \bmod N$ $\vartheta \in_R \{0, 1\}^k, \theta = \vartheta \oplus \rho_1$ $\eta = \sigma \oplus \rho_2, \varphi = MAC_{\rho_3}(\langle \theta, \eta \rangle)$ $count = count + 1$ $append(\tau, uid, count, \vartheta, m, r)$
	(θ, η, φ) \leftarrow
abort IF $MAC_{\rho_3}(\langle \theta, \eta \rangle) \neq \varphi$ and $\sigma = \eta \oplus \rho_2$ and $d_1 = f(v, pwd)$ $s = \sigma \cdot (encode(m, r))^{d_1} \bmod N$ abort IF $s^e \neq encode(m, r)$ $st = \theta \oplus \rho_1, ct = ct + 1$ $token' = (ct, st, v, a, \tau, e, N, pk_{server})$ erase $\beta, d_1, \rho_1, \rho_2, \rho_3, \theta, \eta, \varphi, \vartheta, token$	

Fig. 1. Building block: a single-server scheme with stateful soft-tokens

Let us briefly explain the functions of the protocol elements:

- * β is a value showing that the user knows the password pwd .
- * ρ_1 and ρ_2 are two one-time pads chosen by the user, and will be used by the server to encrypt the state information ϑ and the partial signature σ (produced using the partial private key d_2), respectively.
- * ρ_3 is a one-time message authentication key that allows the user to detect compromise of $token$ because the adversary could tamper with θ while keeping η intact.
- * r is a len_{pad} -bit random string used in the encoding function.
- * γ is the encryption of $m, r, \beta, st, \rho_1, \rho_2$, and ρ_3 .
- * δ and φ are message authentication codes computed using a and ρ_3 , respectively. (Note that both δ and φ are not for preventing abuses, but for detecting attacks.)

Key disabling protocol. In order to disable its private key, the user authenticates itself to the server by conducting a standard username/password authentication protocol corresponding to uid/pwd . The server will query its database to get $b = H_1(pwd)$ from $\langle a, b, uid, d_2, N \rangle = D_{sk_{server}}(\tau)$, where τ corresponds to uid . Any secure password protocol (e.g., [8,4,11,21]) can be used for this purpose.

4.1 Discussions

On 3-factor authentication. Previous key-split schemes (such as [22]) employ a 2-factor authentication mechanism based on a password and a soft-token. Whereas, we employ a 3-factor authentication mechanism so that a signature is produced when a request (i) presents a valid password (i.e., what you know), (ii) is initiated from a party having access to the user's soft-token (i.e., what you have), and (iii) presents a valid fresh one-time secret (i.e., whether you *always* have access to the soft-token). The last factor helps achieve the newly introduced compromise detection and the enhanced abuse prevention (see Section 4.2).

On atomicity of the transactions. We assumed that atomicity of the transactions is ensured. This may be problematic when, for example, the servers are under a DDoS attack. This issue is addressed via another layer of assurance for synchronization in Section 5.

On light-weight key disabling. Allowing a user to disable its private key via a standard username/password authentication has the advantage that a user does not have to resort to its soft-token, which may not be available (e.g., when the user's device is stolen). Although this convenience seemingly gives an adversary the chance to impose denial-of-service attack (i.e., the adversary can request the server to disable the user's private key), we argue that there are no sever consequences. First, suppose an adversary does not know a user's token or password. Then, the adversary can conduct an on-line dictionary attack against the user's

password. This is severe if the on-line dictionary attack can be launched automatically by a software program. Fortunately, there exist some effective methods (e.g., [25]) to force the adversary to conduct a *manual* on-line dictionary attack, which may be unlikely. Second, suppose an adversary knows a user's token but not password. Then, it is seemingly more attractive for the adversary to manage to get the user's password so that it can produce signatures (rather than disable the user's private key). Moreover, the user might also have to disable its private key once the user realized that its soft-token has been compromised. Third, suppose an adversary knows a user's password but not token. Then, the adversary can always disable the user's private key. This may not be seen as a drawback because the user has to disable its private key once the user realized that its password has been compromised. Fourth, suppose an adversary knows a user's token and password. Then, the adversary is already able to produce signatures by contacting the server, which is perhaps more attractive than to conduct a denial-of-service attack by disabling the victim user's private key.

4.2 Analysis

Our scheme does not incur any significant extra complexity, when compared with the starting-point scheme in [22]. Specifically, a soft-token keeps some state information (e.g., 160 bits), and a server keeps some state information linear to the number of users (which can be easily mitigated by letting the server use a pseudorandom function). Moreover, no extra exponentiations are imposed on a user or a server. Below we analyze the security properties.

Proposition 1. *The single server scheme implements some of the requirements specified in Section 3 (the others will be fulfilled in the full-fledged scheme via another layer of protection).*

- * **Abuse prevention.** *This is analyzed in Theorems 1-5.*
- * **Compromise detection.** *Suppose atomicity of the transactions and integrity of the server are guaranteed. Lack of synchronization means that either an adversary had succeeded in impersonating the user, or the token had been tampered with.*
- * **Immediate revocation.** *This is true since the request for disabling a private key is authenticated by pwd , whereas uid is also remembered by the user.*

In order to prove the **abuse prevention**, we introduce the following formal security model. Denote $D\text{-RSA}[\mathcal{E}, \mathbb{D}]$ the real-world single-server signing system based on an encryption scheme \mathcal{E} for the server and dictionary \mathbb{D} . An adversary is given $\langle e, N \rangle$ where $(\langle e, N \rangle, \langle d, N, \phi(N) \rangle) \leftarrow GEN_{RSA}(1^\lambda)$, the public data generated in the initialization procedure, and certain secret data of the user and/or server (depending on the type of the adversary). The goal of the adversary is to forge RSA signatures with respect to $\langle e, N \rangle$. The adversary is allowed to have the following types of oracle queries:

1. **start(m)** – This results in a user to initiate the protocol. The oracle may execute according to the protocol, maintain state as appropriate (i.e., there is an implicit notion of sessions), and return (γ, τ, δ) .

2. $\text{serve}(\gamma, \tau, \delta)$ – This represents the receipt of a message ostensibly from the user. The oracle may execute according to the protocol to return (θ, η, φ) .
3. $\text{finish}(\theta, \eta, \varphi)$ – This represents the receipt of a response ostensibly from the server. The oracle may execute according to the protocol to return a valid signature.

An adversary of type $\mathcal{ADV}\{\text{server}, \text{pwd}\}$, $\mathcal{ADV}\{\text{token}, \text{server}\}$, or $\mathcal{ADV}\{\text{token}\}$ succeeds in breaking the scheme if it can output a valid signature $\langle r, s \rangle$ on message m and there was no $\text{start}(m)$ query. An adversary of type $\mathcal{ADV}\{\text{token}, \text{pwd}\}$ succeeds in breaking the scheme if it can output a valid signature $\langle r, s \rangle$ on message m and there was no $\text{serve}(\gamma, \tau, \delta)$ query, where $\langle m, *, *, *, *, *, *, * \rangle = D_{sk_{\text{server}}}(\gamma)$. An adversary of type $\mathcal{ADV}\{\neg\text{token}^{(i)}, \text{pwd}\}$ succeeds in breaking the scheme if it can output a valid signature $\langle r, s \rangle$ on message m after the user finishes the i^{th} transaction and before the user initiates the $(i + 1)^{\text{th}}$ transaction, and there was no $\text{serve}(\gamma, \tau, \delta)$ query such that $\langle m, *, *, *, *, *, *, * \rangle = D_{sk_{\text{server}}}(\gamma)$.

Denote by q_{user} the number of $\text{start}(\cdot)$ queries to the user, q_{server} the number of $\text{serve}(\cdot, \cdot, \cdot)$ queries to the server. Let q_h and q_f be the number of queries to the random oracles h and f , respectively. Let q_o be the number of other oracle queries not counted above. Let $\bar{q} = (q_{\text{user}}, q_{\text{server}}, q_h, q_f, q_o)$. We also denote by $|\bar{q}| = q_{\text{user}} + q_{\text{server}} + q_h + q_f + q_o$ as the total number of oracle queries.

We say an adversary (\bar{q}, ε) -breaks D-RSA if it makes \bar{q} oracle queries (of the respective types and to the respective oracles) and succeeds with probability at least ε . In the following, “ \approx ” means equality within negligible factors.

We say an attacker \mathcal{A} (q, ε) -breaks \mathcal{E} if the attacker makes q queries to the decryption oracle and $2 \cdot \Pr[\mathcal{A} \text{ succeeds}] - 1 \geq \varepsilon$, which implies $\Pr[\mathcal{A} \text{ outputs } 0|b = 0] - \Pr[\mathcal{A} \text{ outputs } 0|b = 1] \geq \varepsilon$. Note that if \mathcal{E} uses random oracles, the oracles may be queried by the attacker along with the encryption oracle.

We say a forger (q, ε) -breaks a signature scheme if it makes q queries and succeeds with probability at least ε . Note that if \mathcal{S} uses random oracles, the oracles may be queried by the forger along with the signature oracle.

Now we present the theorems for the abuse prevention property of the single-server signing protocol. Theorem 1-4 are extended from [22] and their proofs are deferred to the full version of the present paper [30] (due to space limitation).

Theorem 1. *Suppose $\{f_v\}$ is a pseudorandom function family. If an adversary \mathcal{F} of type $\mathcal{ADV}\{\text{server}, \text{pwd}\}$ can (\bar{q}, ε) -break D-RSA $[\mathcal{E}, \mathbb{D}]$ system, then there exists a forger \mathcal{F}^* able to $(q_{\text{user}}, \varepsilon')$ -break the underlying RSA signature scheme, where $\varepsilon' \approx \varepsilon$.*

Theorem 2. *Let H_1 and f be random oracles. If \mathcal{F} of type $\mathcal{ADV}\{\text{token}, \text{server}\}$ can (\bar{q}, ε) -break D-RSA $[\mathcal{E}, \mathbb{D}]$ system, there exists a forger \mathcal{F}^* able to $(q_{\text{user}}, \varepsilon')$ -break the underlying RSA signature scheme, where $\varepsilon' \approx \varepsilon - \frac{q_h + q_f}{|\mathbb{D}|}$.*

Theorem 3. *Suppose H_1 has a negligible probability of collision over \mathbb{D} . If an adversary \mathcal{A} of type $\mathcal{ADV}\{\text{token}\}$ can (\bar{q}, ε) -break the D-RSA $[\mathcal{E}, \mathbb{D}]$ system for $\varepsilon = \frac{q_{\text{server}}}{|\mathbb{D}|} + \psi$, then there exists either a forger \mathcal{F}^* able to $(q_{\text{user}}, \varepsilon')$ -break*

the underlying RSA signature scheme for $\varepsilon' \approx \frac{\psi}{2}$, or an attacker \mathcal{A}^* able to $(2q_{server}, \varepsilon'')$ -break \mathcal{E} for $\varepsilon'' \approx \frac{\psi}{2(1+q_{user})}$.

Theorem 4. *Suppose the underlying RSA signature scheme is deterministic. If an adversary \mathcal{A} of type $\text{ADV}\{\text{token}, \text{pwd}\}$ can (\bar{q}, ε) -break the $D\text{-RSA}[\mathcal{E}, \mathbb{D}]$ system, then there exists either a forger \mathcal{F}^* able to $(q_{server}, \varepsilon')$ -break the underlying RSA signature scheme for $\varepsilon' \approx \frac{\varepsilon}{2}$, or an attacker \mathcal{A}^* able to $(2q_{server}, \varepsilon'')$ -break \mathcal{E} for $\varepsilon'' \approx \frac{\varepsilon}{2(1+q_{user})}$.*

Theorem 5. *Suppose the underlying RSA signature scheme is deterministic. If an adversary \mathcal{A} of type $\text{ADV}\{\neg\text{token}^{(i)}, \text{pwd}\}$ can (\bar{q}, ε) -break the $D\text{-RSA}[\mathcal{E}, \mathbb{D}]$ system, then there exists either a forger \mathcal{F}^* able to $(q_{server}, \varepsilon')$ -break the underlying RSA signature scheme for $\varepsilon' \approx \frac{\varepsilon}{2q_{user}}$, or an attacker \mathcal{A}^* able to $(2q_{server}, \varepsilon'')$ -break \mathcal{E} for $\varepsilon'' \approx \frac{\varepsilon}{4q_{user}}$.*

Proof. (sketch) Suppose there exists $i, 1 \leq i \leq q_{user}$, such that \mathcal{A} outputs a valid signature for a new message with probability at least ε after the user finishes the i^{th} transaction and before the user initiates the $(i + 1)^{\text{th}}$ transaction. Consider an algorithm \mathcal{F}^{**} that is the same as \mathcal{F}^* in Theorem 4, except that:

- * Let τ be the encryption of 0-string of appropriate length, the first $i - 1$ γ 's be the encryptions of normal messages, and the i^{th} γ be the encryption of 0-string of appropriate length.
- * A token is sent to \mathcal{A} only when it issues a `get()` query, which is not allowed for the i^{th} token generated by the user in the i^{th} transaction. Also, \mathcal{A} can maintain the consistency between the user side and the server side by issuing a `coordinate(token)` query.

Given the simulation generated by \mathcal{F}^{**} , if $\text{Pr}[\mathcal{A} \text{ succeeds}] \geq \frac{\varepsilon}{2}$, then it is clear that there exists \mathcal{F}^* able to $(q_{server}, \varepsilon')$ -break the underlying RSA signature scheme, where $\varepsilon' \approx \frac{\varepsilon}{2q_{user}}$; otherwise, there exists \mathcal{A}^* able to $(2q_{server}, \varepsilon'')$ -break \mathcal{E} , where $\varepsilon'' \approx \frac{\varepsilon}{4q_{user}}$.

Consider an algorithm \mathcal{A}^{**} that is given a public key pk' . Now, \mathcal{A}^{**} can perfectly simulate the real-world system as follows. It generates the pair of public and private keys on behalf of the user. It need have access to the decryption oracle, but at most $2q_{server}$ times. Note that \mathcal{A} is given a newly generated token only when it issues a `get()` query, and that no `get()` query is allowed after the user finishes the i^{th} transaction and before the user initiates the $(i + 1)^{\text{th}}$ transaction. Therefore, \mathcal{A} succeeds in forging with probability at least ε .

Now consider the same simulation, except that τ is the encryption of 0-string of appropriate length, the first $i - 1$ γ 's are the encryptions of normal messages, and the i^{th} γ is the encryption of 0-string of appropriate length. This simulation is equivalent to the simulation of \mathcal{F}^{**} . Therefore, $\text{Pr}[\mathcal{A} \text{ succeeds}] < \frac{\varepsilon}{2}$.

A standard hybrid argument shows that \mathcal{A}^{**} can $(2q_{server}, \frac{\varepsilon}{4})$ -break \mathcal{E} , which means that there exists \mathcal{A}^* able to $(2q_{server}, \frac{\varepsilon}{4q_{user}})$ -break \mathcal{E} . \square

5 Full-Fledged Scheme

Now we present the full-fledged scheme, which is built on top of the building-block discussed in the previous section. The basic idea underlying the scheme is the following. In order to achieve **compromise confinement**, we augment $b = H_1(pwd)$ to $b = H_1(c, pwd)$, where c is a cryptographic secret used to ensure that $b = H_1(c, pwd)$ itself does not leak any significant information of a password pwd . A drawback of this approach is that a user cannot disable its private key using a standard username/password authentication mechanism. Nevertheless, this can be resolved by letting a user choose two passwords, one for generating signatures and the other for disabling its key. Specifically, this can be done by further augmenting $\tau = E_{pk_{server}}(\langle a, b, uid, d_2, N \rangle)$ to $\tau = E_{pk_{server}}(\langle a, b, b^*, uid, d_2, N \rangle)$ such that $b = H_1(c, pwd)$ and $b^* = H_2(\pi)$, where password pwd is for generating signatures, and password π is for disabling the key.

Special care is also taken to address atomicity and availability issues. Suppose a server $server_i$ cannot finish a transaction within a certain time interval. Then, it is reasonable to allow the user to contact another server. This flexibility complicates the facilitation of transaction atomicity. We resolve this issue by incorporating a simple “commit/rollback” mechanism. Moreover, a *commit* or *rollback* request should be authenticated. This can be done by augmenting $\gamma = E_{pk_{server}}(\langle m, r, \beta, st, \rho_1, \rho_2, \rho_3 \rangle)$ to $\gamma = E_{pk_{server}}(\langle m, r, \beta, st, \rho_1, \rho_2, \rho_3, \rho_4, \rho_5 \rangle)$, where $\rho_4 = H_3(\rho_4^*)$ is used for committing a transaction and $\rho_5 = H_3(\rho_5^*)$ is used for rolling back a transaction. The scheme has the following components.

Token initialization. Suppose $H_1 : \{0, 1\}^* \rightarrow \{0, 1\}^k$, $H_2, H_3 : \{0, 1\}^k \rightarrow \{0, 1\}^k$, and $f : \{0, 1\}^* \rightarrow \{0, 1\}^{\lambda+k}$ are appropriate hash and pseudorandom functions, respectively. A user chooses two passwords – pwd for generating signatures and π for disabling the private key, and a pair of public and private keys ($pk_{user} = \langle e, N \rangle$; $sk_{user} = \langle d, N, \phi(N) \rangle$). A user chooses n servers as its service providers, and each server has a public key $pk_{server,i}$, where $1 \leq i \leq n$. The initialization process proceeds as follows.

$$\begin{aligned}
uid &= username \\
v_i &\in_R \{0, 1\}^k, 1 \leq i \leq n \\
a_i &\in_R \{0, 1\}^k, 1 \leq i \leq n \\
c_i &\in_R \{0, 1\}^k, 1 \leq i \leq n \\
b_i &= H_1(c_i, pwd), 1 \leq i \leq n \\
b^* &= H_2(\pi) \\
d_{1,i} &= f(v_i, pwd), 1 \leq i \leq n \\
d_{2,i} &= d - d_{1,i} \bmod \phi(N), 1 \leq i \leq n \\
\tau_i &= E_{pk_{server,i}}(\langle a_i, b_i, b^*, uid, d_{2,i}, N \rangle), 1 \leq i \leq n \\
ct_i &= 0, 1 \leq i \leq n \\
st_i &\in_R \{0, 1\}^k, 1 \leq i \leq n \\
token_i &= (ct_i, st_i, v_i, a_i, c_i, \tau_i, e, N, pk_{server,i}), 1 \leq i \leq n
\end{aligned}$$

The user keeps n soft-tokens on its device and erases all the other values.

Server databases. Each server, $server_i$ for $1 \leq i \leq n$, maintains a database of $\mathcal{Y}_i = (\tau, uid, count, \vartheta, m, r)$, where uid is obtained after receiving the first

service request, *count* is an incremental counter (with initialized value zero) maintained by the server. Each record of the database represents a transaction with respect to $\tau = E_{pk_{server}}(\langle a, b, b^*, uid, d_2, N \rangle)$, There are two types of operations regarding the database. First, $\mathbf{append}(\tau, uid, count, \vartheta, m, r)$ appends $(\tau, uid, count, \vartheta, m, r)$ into the database. Second, $\mathbf{last}(\tau, uid, count, \vartheta, m, r)$ returns either $(\tau, uid, count, \vartheta, m, r)$ corresponding to the *last* transaction and τ , or NULL (meaning that the token corresponding to τ has never been used before).

Signing protocol. The protocol is depicted in Figure 2. The user first sends a request to a (randomly chosen) server. If the transaction can not be finished within certain time, the user contacts another server. Note that such a switching process can be made transparent to the user.




Fig. 2. The full-fledged scheme

Key disabling protocol. In order to disable its private key, the user authenticates itself to each of the n servers by proving that it knows the password π corresponding to *uid*. This process is the same as in the underlying single server protocol, and can be made automatic via an appropriate software design (for ease of deployment).

5.1 Analysis and Discussion

Since transaction atomicity plays an important role in our system, we must show that it has no significant security consequence if atomicity is violated.

Proposition 2. *Suppose there is no system crash resulting in the loss of system state information. Suppose a server keeps a database of state information $(\tau, uid, count, \vartheta, m, r)$. Then there is no denial-of-service attack because of an out-of-synchronization between a user and a server.*

Proof. Recall that a server classifies transactions into two categories: “committed” and “pending”. Note that a successfully rolled-back transaction is treated as a “committed” one, but corresponding to the last successfully committed transaction. This is so because that a rolled-back transaction can be viewed as one where no actions have been taken whatsoever. So, when a user sends a request with certain state information ϑ , there are three cases.

- * ϑ matches the state information corresponding to a “committed” transaction \mathcal{T} . There are further two cases.
 1. \mathcal{T} is the last “committed” transaction. This is the normal case and the protocol proceeds as specified.
 2. \mathcal{T} is not the last “committed” transaction. The server simply rejects this request. In this case, the server could inform the user, perhaps via an out-of-band channel, that the user side might have been compromised. This is so because in our design a user updates its state information before sending a commitment request.
- * ϑ matches the state information corresponding to a “pending” transaction \mathcal{T} . There are further two cases.
 1. \mathcal{T} is the last “pending” transaction. This means that the server did not commit the transaction yet. Then the server can simply accept the request and change the “pending” transaction into a “committed” one.
 2. \mathcal{T} is not the last “committed” transaction. This is impossible, because our design ensures that whenever a server accepts a request (and sends new state information to a user), it always treat the *last* transaction as committed.
- * ϑ matches no state information in the server database at all. The server simply rejects the request. In this case, the server could inform the user, perhaps via an out-of-band channel, that the user side might have been compromised. This is so because in our design we let a server choose the state information.

□

The above proposition implies that we can treat all transactions as atomic. Now we are ready to look at the properties of the full-fledged scheme. We claim that the full-fledged scheme implements all of the goals specified in Section 3. Informally, we observe the following:

- * **Abuse prevention.** This can be formally analyzed by extending the theorems in the underlying single server scheme.

- * **Compromise detection.** Suppose atomicity of the transactions and integrity of the server are guaranteed. Lack of synchronization means either an adversary had successfully impersonated the user, or the token had been tampered with. In either case, the user needs to disable its private key.
- * **Immediate revocation.** This is true since the request for disabling a private key is authenticated by π , whereas *uid* is also remembered by the user.
- * **Compromise confinement.** Once it is known that a server has been compromised, only the users who suspects that their tokens have been compromised need to re-initialize their keys, whereas the others just need to erase their tokens corresponding to the compromised server (i.e., no need to re-initialize their private keys). We notice that when a user knows that its password for disabling its private key, namely π , may have been compromised (e.g., due to the compromise of a server), the user can, if desired, execute an appropriate password-change protocol (e.g., the one associated with the adopted password authentication scheme) to update π to some π' . Such a process would be much more light-weight than a process for updating private keys. Note, however, that the password update process is not necessary from the perspective of the security of the user's signature scheme.
- * **Scalability.** The system is scalable because the servers are decentralized, namely that the servers are operated by possibly many service providers.
- * **High availability.** The system is highly available since a single server is sufficient for the users to generate signatures.

6 Conclusion and Future Work

We presented a scalable and secure cryptographic service, which has the following features: (1) it incorporates a 3-factor authentication mechanism; (2) it supports immediate revocation of a cryptographic key or functionality in question; (3) the damage due to the compromise of a server is contained; (4) it is scalable and highly available.

Acknowledgement. We thank Philip MacKenzie for valuable feedbacks, and the anonymous reviewers for their comments.

This work was supported in part by NSF and UTSA.

References

1. Anderson, R.: Invited Talk at ACM CCS'97.
2. Asokan, N., Tsudik, G., Waidner, M.: Server-Supported Signatures. *Journal of Computer Security* 5(1) (1997)
3. Bellare, M., Miner, S.: A forward-secure digital signature scheme. In: Wiener, M.J. (ed.) *CRYPTO 1999*. LNCS, vol. 1666, Springer, Heidelberg (1999)
4. Bellare, M., Pointcheval, D., Rogaway, P.: Authenticated Key Exchange Secure against Dictionary Attacks. In: Preneel, B. (ed.) *EUROCRYPT 2000*. LNCS, vol. 1807, Springer, Heidelberg (2000)

5. Bellare, M., Rogaway, P.: Random Oracles Are Practical: A Paradigm for Designing Efficient Protocols. In: Proc. ACM CCS'93, pp. 62–73.
6. Bellare, M., Rogaway, P.: Optimal asymmetric encryption – How to encrypt with RSA. In: De Santis, A. (ed.) EUROCRYPT 1994. LNCS, vol. 950, Springer, Heidelberg (1995)
7. Bellare, M., Rogaway, P.: The Exact Security of Digital Signatures - How to Sign with RSA and Rabin. In: Maurer, U.M. (ed.) EUROCRYPT 1996. LNCS, vol. 1070, Springer, Heidelberg (1996)
8. Bellare, S., Merritt, M.: Encrypted Key Exchange: Password-based Protocols Secure against Dictionary Attack. In: Proc. IEEE Security and Privacy, IEEE Computer Society Press, Los Alamitos (1992)
9. Boneh, D., Ding, X., Tsudik, G., Wong, C., Method, A.: for Fast Revocation of Public Key Certificates and Security Capabilities. In: Proc. Usenix Security Symposium (2001)
10. Boyd, C.: Digital Multisignatures. In: Beker, H.J., Piper, F.C. (eds.) Cryptography and Coding, pp. 241–246. Clarendon Press (1989)
11. Boyko, V., MacKenzie, P., Patel, S.: Provably Secure Password Authentication and Key Exchange Using Diffie-Hellman. In: Preneel, B. (ed.) EUROCRYPT 2000. LNCS, vol. 1807, Springer, Heidelberg (2000)
12. Cramer, R., Shoup, V.: A Practical Public Key Cryptosystem Provably Secure Against Adaptive Chosen Ciphertext Attack. In: Krawczyk, H. (ed.) CRYPTO 1998. LNCS, vol. 1462, Springer, Heidelberg (1998)
13. Dean, D., Berson, T., Franklin, M., Smetters, D., Spreitzer, M.: Cryptography as a Network Service. In: Proc. NDSS'01 (2001)
14. Denning, D.E.: Digital Signature with RSA and other Public-Key Cryptosystems. C. ACM 27(4), 388–392 (1984)
15. Dodis, Y., Katz, J., Xu, S., Yung, M.: Strong Key-Insulated Signature Schemes. In: Desmedt, Y.G. (ed.) PKC 2003. LNCS, vol. 2567, Springer, Heidelberg (2002)
16. Ganesan, R., Yaksha: Augmenting Kerberos with Public Key Cryptography. In: Proc. NDSS'95 (1995)
17. Goldreich, O., Goldwasser, S., Micali, S.: How to Construct Random Functions. J. ACM 33(4), 210–217 (1986)
18. Goldwasser, S., Micali, S., Rivest, R.: A Digital Signature Scheme Secure against Adaptive Chosen-Message Attacks. SIAM J. Computing 17(2), 281–308 (1988)
19. Itkis, G., Reyzin, L.: Forward-Secure Signatures with Optimal Signing and Verifying. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, Springer, Heidelberg (2001)
20. Itkis, G., Reyzin, L.: SiBIR: Signer-Base Intrusion-Resilient Signatures. In: Yung, M. (ed.) CRYPTO 2002. LNCS, vol. 2442, Springer, Heidelberg (2002)
21. Katz, J., Ostrovsky, R., Yung, M.: Efficient Password-Authenticated Key Exchange Using Human-Memorizable Passwords. In: Pfitzmann, B. (ed.) EUROCRYPT 2001. LNCS, vol. 2045, Springer, Heidelberg (2001)
22. MacKenzie, P., Reiter, M.: Networked Cryptographic Devices Resilient to Capture. In: Proc. IEEE Security and Privacy, IEEE Computer Society Press, Los Alamitos (2001)
23. Matsumoto, T., Kato, K., Imai, H.: Speeding Up Secret Computations with Insecure Auxiliary Devices. In: Goldwasser, S. (ed.) CRYPTO 1988. LNCS, vol. 403, Springer, Heidelberg (1990)
24. Perlman, R., Kaufman, C.: Secure Password-based Protocol for Downloading a Private Key. In: Proc. NDSS'99 (1999)

25. Pinkas, B., Sander, T.: Securing Passwords Against Dictionary Attacks. In: Proc. ACM CCS'02 (2002)
26. Rackoff, C., Simon, D.: Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In: Feigenbaum, J. (ed.) CRYPTO 1991. LNCS, vol. 576, Springer, Heidelberg (1992)
27. Rivest, R.A., Shamir, A., Adleman, L., Method, A.: A Method for Obtaining Digital Signatures and Public-Key Cryptosystems. C. ACM 21(2), 120–126 (1978)
28. Schneider, F.: Implementing Fault-Tolerant Services Using the State Machine Approach: A Tutorial. ACM Comput. Surv. 22(4), 299–319 (1990)
29. Xu, S., Sandhu, R.: Two Efficient and Provably Secure Schemes for Server-Assisted Threshold Signatures. In: Proc. RSA Con. – Cryptographer's Track (2003)
30. Xu, S., Sandhu, R.: A Scalable Secure Cryptographic Service. Full version of the present paper, available at www.cs.utsa.edu/~shxu