Key Privacy for Identity Based Encryption *

Internet Security Research Lab Technical Report 2006-2

Jason E. Holt
Internet Security Research Lab
Brigham Young University

March 2006

Abstract

We define key privacy for IBE systems in terms of two properties, indistinguishability under chosen identity attack, and indistinguishability under chosen key generator attack. Further, we show that the BasicIdent system in the Boneh/Franklin IBE has these properties under chosen plaintext attack.

1 Introduction

In 2001, Bellare et al. [1] described the notion of key privacy for public key cryptosystems. Informally, key privacy means that an adversary cannot determine the recipient of a message given only its ciphertext. This property is obviously important for public key systems, and for IBE systems used as replacements for public key systems. However, IBE can also form the basis of a number of new digital credential systems [3, 5]. Hidden Credentials [4] in particular require that the underlying IBE possess key privacy in order to satisfy their requirement of credential indistinguishability. Their paper actually gives a somewhat informal proof under the FullIdent system from [2], but here we give more formal treatment to the problem, and address key privacy directly for the underlying BasicIdent system.

2 Definitions

2.1 Key Privacy

In [1], Bellare et al. define two properties for public key cryptosystems, IK-CPA and IK-CCA. IK-CPA stands for Indistinguishability of Keys under Chosen Plaintext Attack, while IK-CCA is the corresponding property for Chosen Ciphertext Attack.

In IBE systems, there are two properties used during encryption, rather than the single public key of the message recipient used in public key encryption. Thus, we define an ID-II-CPA game which can be used to prove whether the recipients identity string can be determined by an attacker, and an ID-IKG-CPA game where the attacker tries to distinguish between messages encrypted using public keys from differing Private Key Generators (PKGs). Using the random oracle model (ROM), we prove that the BasicIdent system in [2] has these properties. However, we leave the corresponding ID-II-CCA and ID-IKG-CCA properties of the CCA-secure FullIdent system for future work.

2.2 Identity Based Encryption

The four functions which implement an IBE system are Setup, Extract, Encrypt and Decrypt.

Setup creates a set of system parameters, including a public key for the PKG and its corresponding secret. Everything but the secret can be published to anyone who wishes to use (or attack) the system.

Extract takes a set of system parameters generated by Setup, an identity string and the PKGs secret, and outputs the private key corresponding to the identity.

Encrypt takes a set of system parameters, an identity and a message, and outputs a ciphertext which can only be decrypted with the private key for the identity.

Decrypt takes a set of system parameters and a private key, and returns the corresponding plaintext.

*This research was supported by funding from the National Science Foundation under grant no. CCR-0325951 and prime cooperative agreement no. IIS-0331707, and The Regents of the University of California.

2.3 ID-II-CPA

Identity-based Indistinguishability of Identity under Chosen Plaintext Attack is the first property we define to establish the key privacy characteristics of an IBE scheme. Informally, an ID-II-CPA attack is considered successful if an attacker can learn anything about the identity passed to Encrypt from the ciphertext it returns.

More precisely, we define an ID-II-CPA game in which an attacker Ida and a challenger Alice interact. This game is modelled after the IND-ID-CPA game defined in section 2 of the Boneh/Franklin paper. The game is defined as follows:

Setup: Alice runs Setup and publishes Params, which comprises all of the public system parameters for the PKG, including its public key.

Extraction queries: At any time, Ida may demand that Alice provide the private keys corresponding to any identity of her choosing except the identities she uses for the challenge. It follows that Ida can use such private keys to decrypt arbitrary ciphertexts encrypted against the corresponding identity.

Challenge: Ida chooses a message M and identities ID_0 and ID_1, and sends these to Alice. Alice chooses a random bit $b \in 0, 1$ and returns the output of $Encrypt(M, ID_b)$ to Ida.

Guess: Ida outputs a guess $b' \in 0, 1$. If $b = b'$, Ida wins the game. We define a function $Adv(k)$ which describes the Ida's probability, beyond random guessing, of winning the game. If, for any given polynomial f and sufficiently large k, $Adv(k) > f(k)$, then we say that the attack is successful. That is, it has a nonnegligible advantage in the ID-II-CPA game.

2.4 ID-IKG-CPA

Identity-based Indistinguishability of Key Generator under Chosen Plain-text Attack (ID-IKG-CPA) is the other property required for an IBE system to attain key privacy. Like chosen identity attacks, chosen key generator attacks reveal information about encrypted messages. The ID-IKG-CPA game is similar to the ID-II-CPA game:

Setup: None.

PKG Extraction queries: At any time, Ida may demand that Alice run Setup and publish the public parameters it returns. Ida may also demand the PKG secret for any of the public values she does not use in the challenge.

Challenge: Ida chooses a message M, an identity ID, and Setup public parameters P_0 and P_1. Ida sends these values to Alice. Alice chooses a random bit $b \in 0, 1$ and returns the output of $Encrypt(ID, M)$ using system parameters P_b.

Guess: Ida outputs a guess $b' \in 0, 1$. If $b = b'$, Ida wins the game. As in the previous proof, a successful attacker is one which has a nonnegligible advantage in guessing b.

2.5 ID-II-CCA and ID-IKG-CCA

The properties corresponding to ID-II-CPA and ID-IKG-CPA for adversaries who can obtain decryptions of chosen ciphertexts are ID-II-CCA and ID-IKG-CCA. The games for these two properties are identical to the corresponding CPA games, except that the adversary is also allowed decryption queries, just as in the traditional CCA game. Formally:

Decryption query: At any time, the attacker may demand the decryption of any ciphertext other than the challenge ciphertext, using the decryption key for any challenge identity or PKG.

Boneh and Franklin offer an augmentation of BasicIdent which offers CCA security. We expect this system, called FullIdent, to offer ID-II-CCA and ID-IKG-CCA security, but leave the proofs to future work.

3 Key Privacy and the Boneh/Franklin IBE

Here we show that the BasicIdent system from [2] has both ID-II-CPA and ID-IKG-CPA security.

3.1 The Bilinear Diffie Hellman Assumption

Boneh and Franklin show that any efficient Chosen Plaintext Attack (CPA) on BasicIdent implies an efficient solution to the Bilinear Diffie-Hellman (BDH) problem, which is assumed to be hard. The BDH problem is defined as follows:

System parameters: Given two groups G_1 and $G2$, and a so-called "admissible bilinear map" from G_1 to G_2, $\hat{e} : G_1 \times G_1 \rightarrow G_2$. $\hat{e}$ has the property that for integers $a, b \in \mathbb{Z}$ and a point $P \in G_1, \hat{e}(aP, bP) = \hat{e}(P, P)^{ab}$.

Challenge: Given P, aP, bP, cP, such that $P \in G_1$ and a, b, c are chosen at random from $\mathbb{Z}$, calculate $\hat{e}(P, P)^{abc}$.

3.2 BasicIdent

BasicIdent is an IBE system. Its simplified definition is as follows:

Setup defines two groups G_1 and G_2 according to the security parameter k, and two cryptographic hash functions H_1 and H_2 in the random oracle model. $H_1 : \{0, 1\}^* \rightarrow G_1$, while $H_2 : G_2 \rightarrow \{0, 1\}^*$. It also defines an admissible bilinear map $\hat{e} : G_1 \times G_1 \rightarrow G_2$. Setup then generates public parameters $PKG_{pub} = P, sP$ for a secret random integer s and a random element P of G_1.

Extract (ID) takes the string ID and returns its corresponding secret key, $sH_1(ID)$. Note that any "random" point Q returned by H_1 is equal to iP for some unknown (and hard to calculate) integer i.

Encrypt(M, ID) takes message M and identity string ID, chooses a random integer r and returns the ciphertext:

$$\begin{aligned} C &= < rP, M \oplus H_2(\hat{e}(H_1(ID), sP))^r > \\ &= < rP, M \oplus H_2(\hat{e}(iP, sP))^r > \\ &= < rP, M \oplus H_2(\hat{e}(P, P))^{isr} > \end{aligned}$$

where $\oplus$ denotes bitwise exclusive OR.

Decrypt $(C, sH_1(ID))$ Decrypts C using private key $sH_1(ID)$, as follows:

$$\begin{aligned} C =< U, V >=< rP, M \oplus H_2(\hat{e}(H_1(ID), sP)^r) > \\ M = V \oplus H_2(\hat{e}(sH_1(ID), rP)) \\ = V \oplus H_2(\hat{e}(siP, rP)) \\ = V \oplus H_2(\hat{e}(P, P)^{isr}) \\ = (M \oplus H_2(\hat{e}(P, P)^{isr})) \oplus H_2(\hat{e}(P, P)^{isr}) \end{aligned}$$

3.3 Security Games

The Boneh/Franklin paper relies heavily on security games for its security proofs. In these games, a challenger gives cryptographic challenges to an attacker, who attempts to prove that she can solve the challenges. The type of game determines what facilities the challenger must provide the attacker. For example, in the CPA game, attacker Alice tries to prove to challenger Bob that she can distinguish the encryptions of two different plaintexts given the encryption of one of them and as many other encryptions of different messages as she wants. She provides the messages, Bob encrypts one and sends it back, and then Alice tries to guess which message it was, asking Bob for as many encryptions of other messages as she wants.

3.4 IND-ID-CPA and BasicIdent

Indistinguishability under Chosen Plaintext Attack is a property defined by Boneh and Franklin. Indistinguishability under CPA is a basic security property of many cryptosystems, and Boneh and Franklin extend it slightly for IBE systems (thus the -ID- in IND-ID-CPA). The IND-ID-CPA game is almost identical to the IND-ID-CIA game, except that instead of providing two identities during the Challenge phase, the attacker provides one identity and two equal length messages M1 and M2. The challenger then encrypts one of the messages under the identity, and its the attackers job to guess which one.

3.5 ID-II-CPA and BasicIdent

Here we show that BasicIdent has the ID-II-CPA property using the random oracle model. To aid our proof, we add the notion of H_2 queries, which allow Alice to serve as the random oracle for Ida during the game.

Theorem 3.1. *The Boneh-Franklin IBE scheme BasicIdent has the ID-II-CPA property.*

Proof. To prove this theorem we'll show that any chosen identity attack (CIA) against BasicIdent can be efficiently converted into a chosen plaintext attack (CPA) on BasicIdent. Since CPA is assumed to be hard, we will also be able to assume CIA is hard.

On page 13 of their paper, Boneh and Franklin assert that any IND-ID-CPA adversary with nonnegligible advantage over random guessing must obtain that advantage by calculating a proper input to H_2 in a challenge ciphertext:

$$C =< U, V >=< rP, M \oplus H_2(\hat{e}(H_1(ID), sP)^r >$$

and then calling H_2 on that input. We'll call that CPA adversary Alice. This claim is reasonable since we assume that H_2 is a random oracle. They then define an intermediary agent we'll call Bob who challenges Alice in the IND-ID-CPA game. Bob implements a random oracle which Alice uses for H_2. He uses the inputs which Alice gives to H_2 to solve a BDH challenge. That is, he uses Alice to compute $\hat{e}(P, P)^{isr}$ given P, iP, sP, rP. Consequently, Alice herself has an advantage in calculating BDH problems, which breaks our assumption.

Recall that in the IND-ID-CPA game, the attacker chooses two messages $M1, M2$, and receives one of them encrypted by the challenger. The attacker's job is to determine which message was encrypted.

Lemma 3.2. *Any IND-ID-CPA attacker on BasicIdent which can distinguish encryptions of two messages $M1, M2$ is equivalent to an attacker who can determine the input to H_2 given an encryption of an empty message.*

As we just pointed out, any successful attacker must be able to compute the input to H_2 for at least one of the ciphertexts. Once Alice finds the correct input to H_2, it-s quite trivial for her to answer Bob-s challenge of distinguishing which of her messages $M1, M2$ he encrypted. All she has to do is XOR the output of H_2 with V and see whether it turns up $M1$ or $M2$. Since XOR is such a trivial operation, Alice would have an equivalent task finding the input to H_2 if Bob just sent her:

$$C =< rP, 0 \oplus H_2(\hat{e}(H_1(ID), sP)^r) >$$

Lemma 3.3. *Any ID-II-CPA attacker on BasicIdent which can distinguish which of two identities $ID1$, $ID2$ was used in the encryption of a message M is equivalent to an attacker who can determine the input to H_2 given an encryption of an empty message.*

Likewise, our chosen identity attacker, Ida, has to calculate the input to H_2 in order to have any hope of distinguishing which of the identities she supplied to the challenger was used in encrypting her message M. QED.

Using these two facts, we now construct an ID-II-CPA game in which Alice challenges Ida and uses her responses in a modified IND-ID-CPA game with Bob. If Ida has a nonnegligible advantage in her game, Alice will have 1/2 that advantage in her own. That is, if Ida can win her CIA game, then Alice can win her CPA game. If Alice can win her game, then Bob can win his BDH game, which would break our assumption that BDH is hard.

Modification to IND-ID-CPA: Rather than requiring Alice to guess which of two messages was encrypted by Bob, Alice sends Bob an identity string which he uses to encrypt an empty message. Alice wins the game if she outputs the input to H_2, and loses if she outputs any other value. By the first lemma, this game is equivalently difficult to the original IND-ID-CPA game.

Setup: Bob runs Setup and publishes $G_1, G_2, H_1, H_2, \hat{e}, P, sP$, keeping s to himself. Ida and Alice will both use these system parameters.

Extraction queries: At any time, Ida may demand that Alice provide the private keys corresponding to any identity of her choosing except the identities she uses for the challenge. Alice, likewise, can demand any such keys from Bob, and does so on Ida's behalf.

H_2 **queries**: At any time, Ida may demand that Alice produce an output from H_2 for a given input. Ida may not use H_2 directly. Remember, this just lets Alice watch what inputs Ida gives to H_2.

Challenge: Ida chooses a message M and identities ID_0 and ID_1, and sends these to Alice. Alice chooses a random bit $b \in 0, 1$ and sends ID_b to Bob. Bob responds with:

$$C =< U, V >=< rP, H_2(\hat{e}(H_1(ID_b), sP)^r) >$$

Alice sends C to Ida.

Guess: Ida outputs a guess $b' \in 0, 1$. If $b = b'$, Ida wins the game. After Ida guesses, Alice checks her list of inputs to H_2. If any of the inputs produced V , Alice outputs that input. Otherwise, she outputs a random value.

Lemma 3.4. *Alice's advantage in winning the IND-ID-CPA game is at least 1/2 of Ida's advantage in the ID-II-CPA game.*

Ida has two ways to win her game. She can calculate $H_2(\hat{e}(sH_1(ID_b), rP))$ and XOR it with V to recover M. In this case, Alice learns the input to H_2 which she needs to win her game. Or, Ida can calculate $H_2(\hat{e}(sH_1(ID_{b'}), rP))$, where b' is the inverse of b. XORing that with V would yield something other than M, indicating that b is the correct answer. Because b is a random bit, Ida can win up to half the time without calculating the input to H_2 which Alice needs.

For example, assume Ida can win the ID-II-CPA game every time. She sends ID_0 and ID_1 to Alice, who randomly chooses ID_1 to send to Bob. Alice returns Bob's encryption under ID_1 to Ida. Ida picks ID_0 at random and calculates the $\hat{e}$ function accordingly, then sends it to Alice to learn the output of H_2. When it doesn't match the ciphertext Alice sent her earlier, Ida knows that the ciphertext must have used ID_1, and consequently wins the game with Alice. Alice is left without any useful information about the actual input to H_2 which Bob calculated when constructing the ciphertext, and so she guesses randomly. Random guessing has a very low probability of producing the right answer, so Alice almost certainly loses her game with Bob.

On the other hand, since Alice randomly chose an ID to send to Bob, Ida is just as likely to make H_2 queries on the ID Alice did choose as on the ID she didn't. If Ida had picked ID_1 instead of ID_0 when calculating the output of $\hat{e}$, or if Alice had picked ID_0, then Alice would have noticed that the response to Ida's H_2 query was identical to the ciphertext given her by Bob, and could then win her game with him. Since this case happens with probability 1/2 (or more, if Ida decides to calculate $\hat{e}$ for both ID_0 and ID_1), Alice wins her game half as often as Ida wins hers.

We have shown, then, that any ID-II-CPA attacker with nonnegligible advantage implies a modified IND-ID-CPA attacker with nonnegligible advantage who has the same advantage in a nonmodified IND-ID-CPA game. Boneh and Franklin showed that such an attacker would also have a nonnegligible advantage in solving BDH, which contradicts our assumption that BDH is hard.

□

3.6 ID-IKG-CPA and BasicIdent

Theorem 3.5. *BasicIdent also has the ID-IKG-CPA property.*

Proof. For this attack, we use the modified IND-ID-CPA attack from the previous theorem. We also require many IND-ID-CPA challengers that can operate independently and concurrently, and which reveal the PKG secret at the conclusion of the game.

Lemma 3.6. *An IND-ID-CPA game challenger can reveal the PKG secret at the conclusion of a game without giving the attacker any advantage in other IND-ID-CPA games.*

Any IND-ID-CPA attacker can also implement an IND-ID-CPA challenger, since the procedures are public and efficient. Any IND-ID-CPA challenger knows his PKG secret, which is simply a random integer. So any useful advantage an attacker could gain from an external challenger could also be gained, efficiently, by the attacker working alone and could be used as part of her attack without external help.

Lemma 3.7. *IND-ID-CPA games may be performed in parallel.*

Again, since the attacker can efficiently play games in parallel alone, playing games in parallel with external attackers gives her no useful advantage. QED.

As in the previous proof, Ida will play the attacker in an ID-IKG-CPA game against Alice. Alice will play the attacker in multiple modified parallel IND-ID-CPA games against Bob. These games are the same as in the last theorem, except that Bob outputs the PKG secret for that game at its conclusion. Alice initiates a new game with Bob each time Ida makes an extraction query for a new set of PKG parameters. If Ida makes n queries and has advantage A in winning the game, Alice will have advantage $\geq A/n$ in her game.

Setup: None.

PKG Extraction queries: Whenever Ida demands that Alice return a set of PKG parameters, Alice initiates a new IND-ID-CPA game with Bob and forwards the parameters Bob creates to Ida. Any time Ida demands the PKG secret for a set of parameters, Alice forfeits the game with Bob (by outputting a random challenge and guess, which will be wrong with overwhelming probability), and passes the revealed PKG secret along to Ida. Ida may not demand the PKG secret for a set of parameters she uses in the challenge.

H_2 **queries**: At any time, Ida may demand that Alice produce an output from H_2 for a given input. Ida may not use H_2 directly.

Challenge: Ida chooses a message M, an identity ID, and Setup parameters P_0 and P_1 from the set of parameter sets returned in PKG extraction queries. Ida sends these values to Alice. Alice sends M, ID as her challenge in the games corresponding to P_0, P_1, receiving ciphertexts C_0, C_1. Alice chooses a random bit $b \in 0, 1$ and returns C_b to Ida.

Guess: Ida outputs a guess $b' \in 0, 1$. If $b = b'$, Ida wins the game. After Ida guesses, Alice checks her list of inputs to H_2. As in the last game, if any of the inputs hashes to one of the ciphertexts, Alice outputs that input for that game and wins. Then she makes random guesses for all remaining games.

Lemma 3.8. *If Ida makes n PKG extraction queries, Alice's advantage in winning the IND-ID-CPA game is at least* $1/n$ *of Ida's advantage in the ID-IKG-CPA game.*

As in the last game, Ida's advantage comes from calculating the input to H_2 for at least one of the two ciphertexts. Alice creates n games, and has Ida's advantage in winning at least one of them. As in the last proof, this shows that any attack on ID-IKG-CPA with nonnegligible advantage implies an efficient solution to BDH. □

4 Conclusion

Since IBE uses both an identity string and a PKG public key when encrypting messages, we define key privacy separately for these two parameters. We gave security games for our two properties for chosen plaintext attackers, and then extended them for chosen ciphertext attacks. Furthermore, we proved that the BasicIdent system from [2] has these properties under the CPA assumption. We leave proofs for key privacy under the CCA assumption and the CCA-secure FullIdent as an avenue for future work.

References

[1] Mihir Bellare, Alexandra Boldyreva, Anand Desai, and David Pointcheval. Key-privacy in public-key encryption. In ASIACRYPT, pages 566582, 2001.

[2] D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. In Proceedings of Crypto 2001, volume 2139 of Lecture Notes in Computer Science, pages 213229. Springer, 2001.

[3] R. Bradshaw, J. Holt, and K. E. Seamons. Concealing complex policies with hidden credentials. In Eleventh ACM Conference on Computer and Communications Security, Washington D.C., Oct 2004. ACM Press.

[4] J. Holt, R. Bradshaw, K. E. Seamons, and H. Orman. Hidden credentials. In 2nd ACM Workshop on Privacy in the Electronic Society, pages 18, Washington, DC, October 2003. ACM Press.

[5] N. Li, W. Du, and D. Boneh. Oblivious signature-based envelope. In Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC 2003), pages 182189, Boston, Massachusette, July 2003. ACM Press.