Stronger Security of Authenticated Key Exchange

Brian LaMacchia[1], Kristin Lauter[2], and Anton Mityagin[3]

[1] Microsoft Corporation, 1 Microsoft Way, Redmond, WA
bal@microsoft.com
[2] Microsoft Research, 1 Microsoft Way, Redmond, WA
klauter@microsoft.com
[3] Microsoft Live Labs, 1 Microsoft Way, Redmond, WA
mityagin@microsoft.com

Abstract. Recent work by Krawczyk [12] and Menezes [16] has highlighted the importance of understanding well the guarantees and limitations of formal security models when using them to prove the security of protocols. In this paper we focus on security models for authenticated key exchange (AKE) protocols. We observe that there are several classes of attacks on AKE protocols that lie outside the scope of the Canetti-Krawczyk model. Some of these additional attacks have already been considered by Krawczyk [12]. In an attempt to bring these attacks within the scope of the security model we extend the Canetti-Krawczyk model for AKE security by providing significantly greater powers to the adversary. Our contribution is a more compact, integrated, and comprehensive formulation of the security model. We then introduce a new AKE protocol called NAXOS and prove that it is secure against these stronger adversaries.

1 Introduction

In this paper we extend the Canetti-Krawczyk [11,12] security model for authenticated key exchange (AKE) to capture attacks resulting from leakage of ephemeral and long-term secret keys. Our security model for authenticated key exchange is defined in the spirit of Bellare and Rogaway [3] and Canetti and Krawczyk [11] by an experiment in which the adversary is given many corruption powers for various key exchange sessions and must solve a challenge on a test session. We extend adversarial capabilities to the following extent: the only corruption powers we do not give an adversary in the experiment are those that would trivially break an AKE protocol. We also define a new AKE protocol which is secure in our new model.

More specifically, in an authenticated key exchange protocol, two parties exchange information and compute a secret key as a function of at least four pieces of secret information: their own long-term (static) and ephemeral secret keys and the other party's long-term and ephemeral secret keys. Of the four pieces of

W. Susilo, J.K. Liu, and Y. Mu. (Eds.): ProvSec 2007, LNCS 4784, pp. 1–16, 2007.

information, we allow an adversary to *reveal*[1] any subset of the four which does not contain both the long-term and ephemeral secrets of one of the parties. To explain this more precisely, we divide AKE test sessions (sessions which are subject to attack by an adversary) into two types. In sessions of the first type ("passive" sessions), the adversary does not cancel or modify communications between the two parties. In sessions of the second type ("active" sessions), the adversary may forge the communication of the second party. Another way to phrase the distinction, as done by Krawczyk in the analysis of the HMQV protocol [12], is whether the adversary actively intervenes in the key exchange session or is a passive eavesdropper.

In addition to distinguishing between passive and active sessions, we identify which pieces of secret information the adversary can reveal without being able to trivially break the AKE protocol (compute the session key for any AKE protocol). In both types of sessions, if an adversary can reveal the long-term and the ephemeral secret keys of one of the parties in the session, then the adversary can trivially compute a session key as it has all the secret information of one of the legitimate parties in the session.

For passive sessions, an adversary may reveal both ephemeral secret keys, both long-term secret keys, or one of each from the two different parties without trivially breaking the protocol. Thus security in our model implies weak Perfect Forward Secrecy, defined by Krawczyk to be security against revelation of long-term secret keys after the session is completed (without active adversarial intervention in the session establishment).

For active sessions, the adversary may forge communications from one of the parties. Thus, if the adversary can also reveal the long-term secret key of that same party, then the adversary can trivially compute the session key. The same argument was used by Krawczyk to show that no 2-round AKE protocol can achieve full perfect forward secrecy (PFS). Still, an adversary can reveal a long-term secret key or ephemeral secret key of the other party without trivially breaking the session. So for another example, our extension to the Canetti-Krawczyk model also implies security against Key Compromise Impersonation (KCI) attacks, where the adversary first reveals a long-term secret of a party and then impersonates others to this party.

Considering attacks involving both types of sessions, it is natural to define a single security model which captures all of them. In our model, in passive test sessions we allow the adversary to reveal any subset of the four pieces of secret information which does not contain both the long-term and ephemeral secrets of one of the parties. In active test sessions, we allow the adversary to reveal only the long-term secret or the ephemeral secret key of the party which is executing the test session. In our security experiment, a test session is still considered *clean* even if the adversary has revealed any of the allowable combinations of secret keys of the two parties.

[1] We say that an adversary "reveals" a piece of secret information when that adversary chooses to learn the value of that information by performing the corresponding key reveal query as defined in Section 3.2.