DOS-Resistant Authentication with Client Puzzles

Tuomas Aura[1], Pekka Nikander[1], and Jussipekka Leiwo[2]

[1] Helsinki University of Technology
P.O.Box 5400, FIN-02015 HUT, Finland
{Tuomas.Aura,Pekka.Nikander}@hut.fi
[2] Vrije Universiteit, Division of Sciences
De Boelelaan 1081A, 1081 HV Amsterdam, The Netherlands
leiwo@cs.vu.nl

Abstract. Denial of service by server resource exhaustion has become a major security threat in open communications networks. Public-key authentication does not completely protect against the attacks because the authentication protocols often leave ways for an unauthenticated client to consume a server's memory space and computational resources by initiating a large number of protocol runs and inducing the server to perform expensive cryptographic computations. We show how stateless authentication protocols and the *client puzzles* of Juels and Brainard can be used to prevent such attacks.

1 Introduction

Denial-of-service (DOS) attacks that exhaust the server's resources are a growing concern on the Internet and other open communications systems. For example, in the SYN attack, a client floods the server with the opening messages of the TCP protocol and fills the space reserved in the server for storing half-open connections.

A solution to such threats is to authenticate the client before the server commits any resources to it. The authentication, however, creates new opportunities for DOS attacks because authentication protocols usually require the server to store session-specific state data, such as nonces, and to compute expensive public-key operations. One solution is to begin with a weak but inexpensive authentication, and to apply stronger and costlier methods only after the less expensive ones have succeeded. An example of a weak authentication is the SYN-cookie protection against the SYN attack where the return address is verified not to be fictional by sending the client a nonce that it must return in its next message. This strategy is not entirely unproblematic because the gradually strengthening authentication results in longer protocol runs with more messages and the security of the weak authentication mechanisms may be difficult to analyze.

In this paper, we advocate the design principle that *the client should always commit its resources to the authentication protocol first and the server should*

B. Christianson et al. (Eds.): Security Protocols, LNCS 2133, pp. 170–177, 2001.

be able to verify the client commitment before allocating its own resources. The rule of thumb is that, at any point before reliable authentication, the cost of the protocol run to the the client should be greater than to the server. The client's costs can be artificially increased by asking it to compute solutions to puzzles that are easy to generate and verify but whose difficulty for the solver can be adjusted to any level. The server should remain stateless and refuse to perform expensive cryptographic operations until it has verified the client's solution to a puzzle.

2 Related Work

Classical models of denial of service by Gligor and Yu [6,17], Amoroso [1], and Millen [13] concentrate the specification and design of fair multi-user operating systems. They assume that all service requests are arbitrated by a trusted computing base (TCB) that enforces the policy set by a single security officer. Their ideas do not extend well to open distributed systems like the Internet where there is no central trusted administration and no global policy or means for enforcing one, and there are too many simultaneous users to theoretically guarantee the availability of any service.

Graph-theoretical models of network reliability by Cunningham [4] and Phillips [14] assess the vulnerability of a communications network to the destruction of nodes and links. These models are useful in the design of network topologies on the physical layer but their applicability does not easily extend to higher protocol layers.

The SYN attack against the TCP connection protocol on the Internet was reported e.g. in [3]. The attack and possible remedies were analyzed in detail by Schuba et al. [15]. Cookies have been previously used in the Photuris protocol by Karn and Simpson [11] and in the Internet Key Exchange (IKE) by Harkins and Carrel [7]. Criticism of the latter [16] shows that the gradually strengthening authentication is not straight-forward to design and a careful analysis of the server resource usage is needed.

Meadows [12] formalized the idea of gradually strengthening authentication. The design goals of a cryptographic protocol should specify how much resources the server may allocate at each level when its assurance of the client's identity and honest purposes step by step increases. This assurance is measured by the resources the client would need to mount a successful attack.

The advantages of statelessness in the beginning of an authentication protocol were recognized by Janson & al. [9] in the KryptoKnight protocol suite. Aura and Nikander [2] generalized the cookie approach to create stateless servers that maintain connections by passing the state data to the client. The paper also gives examples of authentication protocols where the server avoids saving a state until the authentication of the client is complete. Hirose and Matsuura [8] applied these ideas to a DOS-resistant version of their KAP protocol. In addition to remaining stateless, the server in their protocol postpones expensive exponentiation operations until it has verified that the client has performed sim-