

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 2, 2018

C. Cremers
L. Garratt
University of Oxford
S. Smyshlyaev
CryptoPro
N. Sullivan
Cloudflare
C. Wood
Apple Inc.
March 01, 2018

Randomness Improvements for Security Protocols
draft-cremers-cfrg-randomness-improvements-00

Abstract

Randomness is a crucial ingredient for TLS and related security protocols. Weak or predictable "cryptographically-strong" pseudorandom number generators (CSPRNGs) can be abused or exploited for malicious purposes. The Dual EC random number backdoor and Debian bugs are relevant examples of this problem. This document describes a way for security protocol participants to mix their long-term private key into the entropy pool(s) from which random values are derived. This augments and improves randomness from broken or otherwise subverted CSPRNGs.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 2, 2018.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Randomness Wrapper	3
3. Tag Generation	4
4. Application to TLS	4
5. IANA Considerations	4
6. Security Considerations	4
7. Normative References	5
Authors' Addresses	6

1. Introduction

Randomness is a crucial ingredient for TLS and related transport security protocols. TLS in particular uses Random Number Generators (RNGs) to generate several values: session IDs, ephemeral key shares, and ClientHello and ServerHello random values. RNG failures such as the Debian bug described in [DebianBug] can lead to insecure TLS connections. RNGs may also be intentionally weakened to cause harm [DualEC]. In such cases where RNGs are poorly implemented or insecure, an adversary may be able to predict its output and recover secret Diffie-Hellman key shares that protect the connection.

This document proposes an improvement to randomness generation in security protocols inspired by the "NAXOS trick" [NAXOS]. Specifically, instead of using raw entropy where needed, e.g., in generating ephemeral key shares, a party's long-term private key is mixed into the entropy pool. In the NAXOS key exchange protocol, raw entropy output x is replaced by $H(x, sk)$, where sk is the sender's private key. Unfortunately, as private keys are often isolated in HSMs, direct access to compute $H(x, sk)$ is impossible. An alternate yet functionally equivalent construction is needed.

The approach described herein replaces the NAXOS hash with a keyed hash, or pseudorandom function (PRF), where the key is derived from raw entropy output and a private key signature. Implementations SHOULD apply this technique when indirect access to a private key is available and CSPRNG randomness guarantees are dubious, or to provide stronger guarantees about possible future issues with the randomness.

2. Randomness Wrapper

Let x be the raw entropy output of a CSPRNG. When properly instantiated, x should be indistinguishable from a random string of length $|x|$. However, as previously discussed, this is not always true. To mitigate this problem, we propose an approach for wrapping the CSPRNG output with a construction that artificially injects randomness into a value that may be lacking entropy.

Let $\text{PRF}(k, m)$ be a cryptographic pseudorandom function, e.g., HMAC [RFC2104], that takes as input a key k of length L and message m and produces an output of length M . For example, when using HMAC with SHA256, L and M are 256 bits. Let $\text{Sig}(sk, m)$ be a function that computes a signature of message m given private key sk . Let G be an algorithm that generates random numbers from raw entropy, i.e., the output of a CSPRNG. Let tag be a fixed, context-dependent string. Let KDF be a key derivation function, e.g., HKDF-Extract [RFC5869] (with first argument set to nil), that extracts a key of length L suitable for cryptographic use. Lastly, let H be a cryptographic hash function that produces output of length M .

The construction works as follows: instead of using x when randomness is needed, use:

```
PRF(KDF(G(x) || H(Sig(sk, tag1))), tag2)
```

Functionally, this computes the PRF of a string (tag2) with a key derived from the CSPRNG output and signature over a fixed string (tag1). See Section 3 for details about how " tag1 " and " tag2 " should be generated. The PRF behaves in a manner that is indistinguishable from a truly random function from $\{0, 1\}^L$ to $\{0, 1\}^M$ assuming the key is selected at random. Thus, the security of this construction depends upon the secrecy of $H(\text{Sig}(sk, \text{tag1}))$ and $G(x)$. If the signature is leaked, then security reduces to the scenario wherein the PRF provides only a wrapper to $G(x)$.

In systems where signature computations are not cheap, these values may be precomputed in anticipation of future randomness requests. This is possible since the construction depends solely upon the CSPRNG output and private key.

Sig(sk, tag1) MUST NOT be used or exposed beyond its role in this computation. Moreover, Sig MUST be a deterministic signature function, e.g., deterministic ECDSA [RFC6979].

3. Tag Generation

Both tags SHOULD be generated such that they never collide with another accessor or owner of the private key. This can happen if, for example, one HSM with a private key is used from several servers, or if virtual machines are cloned.

To mitigate collisions, tag strings SHOULD be constructed as follows:

- o tag1: Constant string bound to a specific device and protocol in use. This allows caching of Sig(sk, tag1). Device specific information may include, for example, a MAC address. See Section 4 for example protocol information that can be used in the context of TLS 1.3.
- o tag2: Non-constant string that includes a timestamp or counter. This ensures change over time even if randomness were to repeat.

4. Application to TLS

The PRF randomness wrapper can be applied to any protocol wherein a party has a long-term private key and also generates randomness. This is true of most TLS servers. Thus, to apply this construction to TLS, one simply replaces the "private" PRNG, i.e., the PRNG that generates private values, such as key shares, with:

```
HMAC(HKDF-Extract(nil, G(x) || Sig(sk, tag1)), tag2)
```

Moreover, we fix tag1 to protocol-specific information such as "TLS 1.3 Additional Entropy" for TLS 1.3. Older variants use similarly constructed strings.

5. IANA Considerations

This document makes no request to IANA.

6. Security Considerations

A security analysis was performed by two authors of this document. Generally speaking, security depends on keeping the private key secret. If this secret is compromised, the scheme reduces to the scenario wherein the PRF provides only an outer wrapper on usual CSPRNG generation.

The main reason one might expect the signature to be exposed is via a side-channel attack. It is therefore prudent when implementing this construction to take into consideration the extra long-term key operation if equipment is used in a hostile environment when such considerations are necessary.

The signature in the construction as well as in the protocol itself MUST be deterministic: if the signatures are probabilistic, then with weak entropy, our construction does not help and the signatures are still vulnerable due to repeat randomness attacks. In such an attack, the adversary might be able to recover the long-term key used in the signature.

Under these conditions, applying this construction should never yield worse security guarantees than not applying it assuming that applying the PRF does not reduce entropy. We believe there is always merit in analysing protocols specifically. However, this construction is generic so the analyses of many protocols will still hold even if this proposed construction is incorporated.

7. Normative References

[DebianBug]

Yilek, Scott, et al, ., "When private keys are public - Results from the 2008 Debian OpenSSL vulnerability", n.d., <<https://pdfs.semanticscholar.org/fcf9/fe0946c20e936b507c023bbf89160cc995b9.pdf>>.

[DualEC]

Bernstein, Daniel et al, ., "Dual EC - A standardized back door", n.d., <<https://projectbullrun.org/dual-ec/documents/dual-ec-20150731.pdf>>.

[NAXOS]

LaMacchia, Brian et al, ., "Stronger Security of Authenticated Key Exchange", n.d., <<https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/strongake-submitted.pdf>>.

[RFC2104]

Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997, <<https://www.rfc-editor.org/info/rfc2104>>.

[RFC5869]

Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", RFC 5869, DOI 10.17487/RFC5869, May 2010, <<https://www.rfc-editor.org/info/rfc5869>>.

- [RFC6979] Pornin, T., "Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)", RFC 6979, DOI 10.17487/RFC6979, August 2013, <<https://www.rfc-editor.org/info/rfc6979>>.
- [X9.62] American National Standards Institute, ., "Public Key Cryptography for the Financial Services Industry -- The Elliptic Curve Digital Signature Algorithm (ECDSA). ANSI X9.62-2005, November 2005.", n.d..

Authors' Addresses

Cas Cremers
University of Oxford
Wolfson Building, Parks Road
Oxford
England

Email: cas.cremers@cs.ox.ac.uk

Luke Garratt
University of Oxford
Wolfson Building, Parks Road
Oxford
England

Email: luke.garratt@cs.ox.ac.uk

Stanislav Smyshlyaev
CryptoPro
18, Sushevsky val
Moscow
Russian Federation

Email: svs@cryptopro.ru

Nick Sullivan
Cloudflare
101 Townsend St
San Francisco
United States of America

Email: nick@cloudflare.com

Christopher A. Wood
Apple Inc.
One Apple Park Way
Cupertino, California 95014
United States of America

Email: cawood@apple.com