ICNRG J. Seedorf
Internet-Draft NEC
Intended status: Informational March 5, 2015
Expires: September 6, 2015

Binding Self-certifying Names to Real-World Identities with a Web-of-Trust
draft-seedorf-icn-wot-selfcertifying-01

Abstract

Self-certifying names are one way of binding a given public key to a certain name in Information Centric Networking. However, an additional binding of a self-certifying name to a Real-World identity is needed in most cases, so that a recipient of some information cannot only verify that the publisher was in possession of the correct corressponding private key for the requested name, but that in addition the name itself is the intended one. This draft specifies how such a binding of Real-World identities with self-certifying ICN names can be done, taking existing IETF specifications into account.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction 2
2. High-Level Design 3
3. Standardisation Considerations 4
3.1. High-Level Considerations 4
3.2. Existing Information-Centric Naming Schemes in the IETF . 5
3.3. Existing Web-of-Trust Standards in the IETF 5
3.4. Hash Extension Techniques 5
4. Conclusion 5
5. References 5
5.1. Normative References 6
5.2. Informative References 6
Appendix A. Acknowledgment 7
Author’s Address 7

1. Introduction

Self-certifying names provide the useful property that any entity in a distributed system can verify the binding between a corresponding public key and the self-certifying name without relying on a trusted third party [Aura2003]. Self-certifying names thus provide a decentralized form of data origin authentication. This feature makes self-certifying names a prime candidate for addressing the security requirements in Information Centric Networking (ICN) (which are inherently different from IP networks): a source can digitally sign data associated with a self-certifying name, and any intermediate entity (e.g. ICN-router/Cache) or receiving entity (i.e. issuer of a request for the name) can verify the signature, without the need to verify the identity of the host that caches the object, nor relying on a trusted third party, or a Public Key Infrastructure (PKI). However, as noted in [Ghodsi2011] and elsewhere, self-certifying names lack a binding with a corresponding real-world identity (RWI): the concept enables to verify that whoever signed some data was in possession of the private key associated with the self-certifying name, but it does not provide any means to verify what real-world identity corresponds to the public key, i.e. who actually signed the data [Ghodsi2011] [Nom2014].

In principle, this binding between a public key and an RWI could be provided by a PKI, or alternatively by a Web-of-Trust (WoT)

[Ghodsi2011]. Several ICN approaches use a PKI [Survey] . However, until recently, there have not been concrete proposals for a WoT-based approach for binding a public key (or a self-certifying name) with an RWI in content-oriented architectures. A concrete approach on how this can be done has been proposed in [Nom2014]. This document has the objective of providing the corresponding necessary standards specification to enable this approach (or similar ones) in principle in an interoperable way.

2. High-Level Design

On a high level, binding of self-certifying names and a Web-of-Trust can be achieved in the following way (see [Nom2014] for a detailed example of such an approach): The WoT key-ID is equivalent to the self-certifying name part used in the naming scheme. This ties the self-certifying name with the ID of the corresponding public key in the WoT.

For instance, in the existing PGP Web-of-Trust, the V4 key ID is the lower 64 bits of the fingerprint of the public key, where the fingerprint is essentially the 160-bit SHA-1 hash of the public key [RFC2440]. So if a self-certifying name would be based on the same lower 64-bits of the fingerprint of a given public key, this public key would be tied to the self-certifying name and at the same time be tied to the real-world identity used in the WoT, e.g. an email-address or the real (i.e. non-self-certifying) name of a given ICN publisher.

Thus, if a user requests the content for a self-certifying name in a given ICN architecture, he/she would retrieve the content which contains a digital signature and the corresponding public key for the self-certifying name. The user can then verify that the content retrieved indeed belongs to the name by first hashing the public key and confirm that the hash (or part of it) matches the requested name, and second using the public key to verify the signature over the content. This is in principle the general way of using self-certifying names for data origin authentication in distributed systems. If, in addition, (part of) the self-certifying name is equivalent to a WoT key-ID, the user can use any WoT infrastructure (e.g. PGP keyservers) to retrieve certificates for the key ID that contain/confirm the binding between the corresponding (to the WoT key ID) public key with a real-world identity, such as an email address. This binding provides the requesting user with assurance that the self-certifying name indeed is owned by the intended publisher, i.e. is the correct, intended name from the requestor’s perpective.

The current PGP specification [RFC2440] considers only a bitlength of 64-bit for forming the key-ID, which is not very collision-resistant (collision-resistance among different key-IDs was not a design goal for PGP [RFC2440]). For securely binding a self-certifying name to a WoT key-ID, collision-resistance is a design goal, because otherwise attckaers could potentially forge a binding of their public key with a given self-certifying name. Thus, either a longer bitlength of the hash of the public key (or its fingerprint) must be used, or hash extension techniques [Aura] must be used, which effectively make collision attacks harder for constant bitlengths at the price of the time needed to create a public/private key pair. Future versions of this document will take these design considerations into account.

3. Standardisation Considerations

Future versions of this document will outline a concrete protocol specification for binding self-certifying names to a Web-of-Trust as outlined on a high level in the previous Section. Below some initial standardisation considerations are highlighted, as well as an assessment of existing IETF standards that could be used as building blocks. Also, future versions of this document will look in more detail into existing IETF specifications, e.g. regarding ICN naming ([RFC6920]) and Web-of-Trust ([RFC2440]), and inspect to what extend such existing specifications can be used directly or in a modified form.

3.1. High-Level Considerations

An initial list of details that need to be specified is the following:

- (List of) Asymmetric cryptography algorithm(s) and corresponding bit-length(s)
- (List of) Hash algorithm(s) and corresponding bit-length(s)
- Rules that define what part of the hash is used for forming the self-certifying part of the name, i.e. the Web-of-Trust Key-ID
- Rules for forming a self-certifying name based on a public key
- Semantics of a signature in the Web-of-Trust
- Defintion of how many bits are used in case of hash extension techniques [Aura][RFC3972]

3.2. Existing Information-Centric Naming Schemes in the IETF

RFC 6920 'Naming Things with Hashes' defines a standard for correctly identifying data 'using the output from a hash function' [RFC6920]. In particular, it specifies a '(ni) URI Format' (see [RFC6920], Section 3) and a 'Named Information Hash Algorithm Registry' (see [RFC6920], Section 9.4). These building blocks allow to specify a format for self-certifying names as hashes of WoT public keys, as outlined above,. In particular, truncated hash formats are clearly defined which can be used to form a self-certifying name from a Web-of-Trust public key by defining what part of the hash is used for forming the WoT key-ID self-certifying part of the name (e.g. 'sha-256-64' for a truncated SHA-256 hash to 64 bits).

3.3. Existing Web-of-Trust Standards in the IETF

RFC 2440 asymmetric cryptography algorithms and corresponding bit-length for usage in a Web-of-Trust [RFC2440]. Thus, there is an existing IETF specification that provides this building block needed for binding Self-certifying Names to Real-World Identities with a Web-of-Trust.

3.4. Hash Extension Techniques

RFC 3972 discusses hash extension techniques, i.e. approaches that 'increase the cost of both address generation and brute-force attacks by the same parameterized factor while keeping the cost of address use and verification constant' [RFC3972]. This can be a building block for using hash extension techniques for binding Self-certifying Names to Real-World Identities with a Web-of-Trust.

4. Conclusion

One option for binding self-certifying names to real-world identities is using a Web-of-Trust. This document aims at a concrete specification for providing such a binding, taking existing IETF specification into account. An inspection of existing Web-of-Trust and Naming Scheme standards in the IETF reveal that the basic building blocks for the intended specification for binding Self-certifying Names to Real-World Identities with a Web-of-Trust are already available as IETF standards. Future versions of this document will provide a more detailed specification.

5. References

5.1. Normative References

[RFC2440] Callas, J., Donnerhacke, L., Finney, H., and R. Thayer, "OpenPGP Message Format", RFC 2440, November 1998.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.

[RFC6920] Farrell, S., Kutscher, D., Dannewitz, C., Ohlman, B., Keranen, A., and P. Hallam-Baker, "Naming Things with Hashes", RFC 6920, April 2013.

5.2. Informative References

[Aura] Aura, T. and M. Roe, "Strengthening Short Hash Values", http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.145.7681, .

[Aura2003] Aura, T., "Cryptographically Generated Addresses (CGA)", 6th International Conference on Information Security (ISC), 2003, .

[Ghodsi2011] Ghodsi, A., Koponen, T., Rajahalme, J., Sarolahti, P., and S. Shenker, "Naming in Content-oriented Architectures", ACM SIGCOMM Workshop on Information-centric Networking, 2011, .

[I-D.seedorf-icn-disaster] Seedorf, J., Arumaithurai, M., Tagami, A., Ramakrishnan, K., and N. Blefari-Melazzi, "Using ICN in disaster scenarios", draft-seedorf-icn-disaster-03 (work in progress), March 2015.

[Nom2014] Seedorf, J., Kutscher, D., and F. Schneider, "Decentralised Binding of Self-Certifying Names to Real-World Identities for Assessment of Third-Party Messages in Fragmented Mobile Networks", 2nd Workshop on Name Oriented Mobility (NOM), 2014, .

[Survey] Xylomenos, G., Ververidis, C., Siris, V., Fotiou, N., Tsilopoulos, C., Vasilakos, X., Katsaros, K., and G. Polyzos, "A Survey of Information-Centric Networking Research", IEEE Communications Surveys and Tutorials, Vol. 16, No. 2, pp 1024-1049, 2014, .

Appendix A. Acknowledgment

This document has been supported by the GreenICN project (GreenICN: Architecture and Applications of Green Information Centric Networking), a research project supported jointly by the European Commission under its 7th Framework Program (contract no. 608518) and the National Institute of Information and Communications Technology (NICT) in Japan (contract no. 167). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the GreenICN project, the European Commission, or NICT.

Author's Address

Jan Seedorf
NEC
Kurfuerstenanlage 36
Heidelberg 69115
Germany

Phone: +49 6221 4342 221
Fax: +49 6221 4342 155
Email: seedorf@neclab.eu