A Mechanism for Discovery and Verification of Trust Scheme Memberships: The LIGHTest Reference Architecture

Sven Wagner[1], Sebastian Kurowski[1], Uwe Laufs[2], Heiko Roßnagel[2]

Abstract: Electronic transactions are an integral component of private and business life. For this purpose, a certification of trustworthy electronic identities supported from authorities is often required. Within the EU-funded LIGHTest project, a global trust infrastructure based on DNS is built, where arbitrary authorities can publish their trust information. A high level description of the LIGHTest reference architecture is presented. Then, the Trust Scheme Publication Authority, which enables discovery and verification of trust scheme memberships is introduced.

Keywords: trust infrastructure, trust scheme, trust scheme verification, electronic transaction, trust management, identity management

1 Introduction

Traditionally, we often knew our business partners personally, which meant that impersonation and fraud were uncommon. Nowadays, an ever-increasing number of transactions are conducted virtually over the Internet. As a result electronic transactions are an integral component of private and business life. Thereby, it is important to know, who the partner on the other side is and if it can be trusted. For this a certification of trustworthy electronic identities is required.

Authorities can assist in this matter. For example, the EC and Member States have already legally binding electronic signatures. However, the query of such authorities in a secure manner is currently comparatively complicated due to the lack of a standard for publishing and querying trust information on a global scale. Without this standard, a high number of different protocols and formats need to be queried during the verification process. This is especially cumbersome if more than a single trust domain is involved.

To address this problem, the EU-funded LIGHTest project (http://lightest.eu/) attempts to build a global trust infrastructure. LIGHTest is the acronym for Lightweight Infrastructure for Global Heterogeneous Trust Management in support of an open Ecosystem of Stakeholders and Trust schemes. The LIGHTest infrastructure makes use of the Internet Domain Name System DNS with its existing global infrastructure,

[1] University Stuttgart IAT, Identity Management, Allmandring 35, 70569 Stuttgart, {firstname.name}@iat.uni-stuttgart.de
[2] Fraunhofer IAO, Identity Management, Nobelstr. 12, 70569 Stuttgart, {firstname.name}@iao.fraunhofer.de

organization, governance and security standards. With the LIGHTest infrastructure, arbitrary authorities can publish their trust information. For example, the EC and Member States can use the LIGHTest infrastructure to publish lists of qualified trust services (e.g. business registrars). Further examples are the establishment of trust in the private sector (e.g. international trade, shipping, credit rating). With the help of the open source LIGHTest infrastructure, companies, administrations, and citizens can then easily query the requested trust information, e.g. for the verification of a signed document in the simplest case.

In [BL16] a first introduction into LIGHTest has been provided. We build on this work and in this paper, we focus, after a short related work section in Chapter 2, on the reference architecture (see Chapter3) and the trust scheme publication authority (see Chapter 4), which is one of the major components of the LIGHTest reference architecture and which is used for each verification of an electronic transaction. For the reference architecture, we present the architectural principles and goals, its components and possible scenarios with a detailed description of the assumptions, trust policy and information flow of the basic scenario for trust scheme publication for qualified signatures. For the trust scheme publication authority, we outline the different types of trust scheme representation, the concept for trust scheme publication and publication and querying of trust schemes. We follow this up with a short discussion and outlook in Chapter 5, before we conclude our findings.

2 Related Work

Most of the existing trust infrastructures follow the subsidiarity principle. One prominent example is the eIDAS Regulation (EU) N°910/2014 ([EI14]) on electronic identification and trust services for electronic transactions in the internal market. This includes that each Member State establishes and publishes national trusted lists of qualified trust service providers. For the access of these trusted lists, the EC publishes a central list ("List Of Trusted Lists") which contains links to these lists. Due to the fact that for verifiers the direct use of trust lists can be very onerous, in particular for international electronic transactions, LIGHTest provides a framework that is conceptually comparable to OCSP for querying the status of individual certificates and which facilities the verification of trust.

DANE (DNS-based Authentication of Names Entities) is a standard using DNS and the DNS security extension DNSSEC to derive trust in TLS server certificates (RCF6698 [HS12] and RCF7218 [Gu14]). For this purpose, the DNS resource record TLSA was introduced which associates a TLS server certificate (or public key) with the domain name where the record is found. Within LIGHTest, the DANE standard will be used to secure network communication and where certificates are used for verifying data.

3 Reference Architecture

This section gives an overview of the LIGHTest reference architecture. It defines the macroscopic design of the LIGHTest infrastructure as well as the overall system's components, their functionality and their interaction on a high-level view.

3.1 Architectural Principles and Goals

From the requirements of enabling a globally scalable trust infrastructure that integrates the existing technical and organisational environment considering the given constraints, organisational and technical goals are defined, which needs to be addressed in the architecture.

Therefore, the architecture takes the following general principals into consideration: subsidiarity and no change of data ownership; minimisation of adoption barriers; reuse of existing software and infrastructure; and separation of concerns and abstraction. With the subsidiarity and no change of data ownership principle, participants stay in control regarding all relevant areas, such as data ownership or trust decisions. The minimisation of adoption barriers reduces organisational adoption as well as technical barriers regarding the integration into the existing, real world, technical environment. The reuse of the existing DNS infrastructure with its existing single, global trust root, its world-wide organization composed of name registries, etc. enables organizations, which intend to publish trust schemes to reuse their existing DNS servers. In addition, the reuse of the existing DNS software and protocols aims to lower the adoption barriers for users. For the separation of concerns and abstraction, a modular approach ([Di76]) is used to achieve an easier collaboration and to reduce efforts for maintenance and optimization of the different components ([Pa72]). In addition, the approach of reduction of complexity by abstraction ([Sh95]) is followed.

The technical goals for the architecture are the following: distributed system; extensibility; scalability; security screening; fault tolerance and high availability; maturity; and traceability. A distributed architecture allows the realisation of the required separations described in the general principals above (i.e. subsidiarity, distributed ownership of data, distributed control and responsibilities). The extensibility of the overall system (e.g. for the implementation of new use cases or extending the amount of users) without modifications on architecture level requires an incremental approach. Scalability is an important technical goal to meet the requirements of the large number of possible users and application fields ([TS07]). Security screening is, in addition to the security in the domain of trust infrastructure itself, very important. For this purpose, mechanisms, which enable security screening are required. Fault tolerance and high availability are central aspects which have to be addressed by the architecture ([AL12]) due to the large scale overall approach of LIGHTest. For maturity, the largely reuse of existing and sound technologies facilitates a high level of maturity. Traceability is another, important requirement in the domain of trust decisions, which enables to trace and comprehend the procedure and final decision.

3.2 Components of the Reference Architecture

The major software components of the LIGHTest reference architecture are already introduced in [BL16]. In this section, the components of the LIGHTest reference architecture are described in more detail. In Fig. 1 the components and their interactions are presented, which are required if a verifier wants to validate a received electronic transaction.

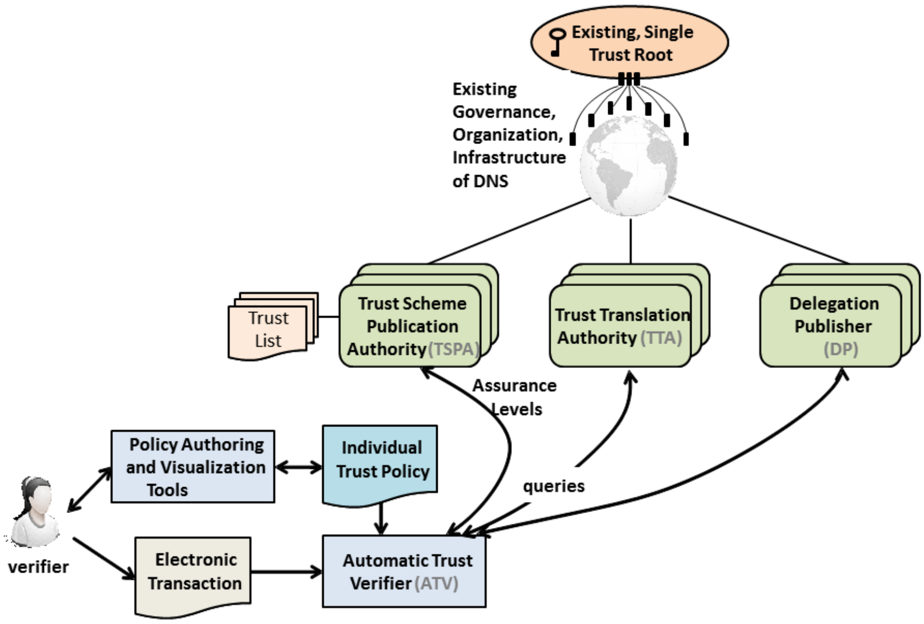

Fig. 1: The LIGHTest Reference Architecture (see also [BL16])

The verifier interacts with the Policy Authoring and Visualization Tools (e.g. desktop or web applications). These tools also facilitate non-technical users the visualization and editing of trust policies, which can be individual and specific for each transaction. The role of the trust policy is the provision of formal instructions for the validation of trustworthiness for a given type of electronic transaction. For example, it states which trust lists from which authorities should be used.

The Automatic Trust Verifier (ATV) takes the electronic transaction and trust policy as input and provides as output if the electronic transaction is trustworthy or not. In addition, the ATV may provide an explanation of its decision, in particular if the transaction was considered as not trustworthy.

The Trust Scheme Publication Authority (TSPA) uses a standard DNS Name Server with DNSSEC extension. A server publishes multiple trust lists under different sub-domains of the authority's domain name. The TSPA enables discovery and verification of trust scheme memberships. In Chapter 4, the TSPA is described in more detail.

The Trust Translation Authority also uses a standard DNS Name Server with DNSSEC extension. Here, a server publishes trust data under different sub-domains of the authority's domain name. In addition, trust translation lists express which authorities from other trust domains are trusted.

The Delegation Publisher also uses a DNS Name Server with DNSSEC extension. Here, a server publishes multiple delegations under different sub-domains of the organization's domain name.

3.3 Scenarios

In this section, examples of usage scenarios are presented. There are basic scenarios for trust publication, trust translation, and trust delegation, which can be used for qualified signatures, qualified seals, qualified identities, or qualified timestamps. The functionality (publish, translate, delegate) of the basic scenarios can be used to realise a wide range of more sophisticated scenarios. These scenarios can be either variants of the basic scenarios or a combination of different basic scenarios. A combination can be composing two trust services in a chaining process where the output level of the inner trust service becomes the input level of the outer trust service. For example, qualified delivery services, where E-registered delivery can be realised using a combination of the scenarios signature and timestamps. Another example is qualified website authentication, where trust publication with qualified identities is the basic scenarios and additionally, trust translation could be used to e.g. authenticate third party users/things.

As an example for a basic scenario, a successful trust scheme publication for qualified signatures is presented. For this example, the following preconditions and assumptions for the electronic transaction and trust policy are made:

1. As preconditions, it is assumed that the verifier and signer are both located in the EC/eIDAS trust domain and that the eIDAS trust domain contains the actual eIDAS trust scheme. This means that trust translation is not required in this scenario. This could for example be managed in the following domain name structure: trust.ec.europa.eu - signature - TrustScheme - actual eIDAS trust scheme for qualified signature.
2. For the electronic transaction, it is assumed that the transaction is simply a signed document. Furthermore, the certificate used to sign the document contains a link to the trust list (Trust Membership Claim) for easier discovery such as "Issuer Alt

Name: XYZ.qualified.trust.admin.ec" that points to the DNS resource records of the native trust scheme for qualified signatures. In addition, this trust scheme lists the certificate as qualified.

3. For the trust policy, it is assumed that trust policy simply states that the signature of the document is trusted if the issuer of the certificate is listed in TrustScheme.signature.trust.ec.europa.eu. Hence it is published as a Boolean trust scheme publication (see Section 4.1 for the definition of Boolean trust scheme publication).

For the basic scenario of a successful trust scheme publication for qualified signatures with the preconditions and assumptions mentioned above, the corresponding information flow in the architecture is described in the following and depicted in Fig. 2.

In step 1, the verifier feeds both, the Trust Policy and the Electronic Transaction into the ATV. The ATV parses the electronic transaction and yields the document, the signer certificate and the issuer certificate (step 2). In step 3, the ATV validates the signature on the document to make sure it is signed by the signer certificate. Next, the ATV validates that the signer certificate is signed by the issuer certificate (step 4). In step 5, the ATV searches the signer certificate and the issuer certificate for discovery information. The ATV finds a Trust Membership Claim in the signer certificate: "Issuer Alt Name: XYZ.qualified.trust.admin.ec". Hence, the issuer name is extracted from the certificate. In step 6, the ATV contacts the TSPA for retrieving the associated trust scheme. Therefore, the ATV issues a DNS query for all relevant resource records for boolean trust schemes for XYZ.qualified.trust.admin.ec. In step 7, the ATV verifies the chain of signatures from the DNS trust root of the DNS response using a validating resolver and stores the response as a "receipt" for future justification of its decision. Next, the ATV converts the resource records of the response into a boolean value (step 8). In the final step, the ATV looks at the trust policy and detects that the trust scheme, TrustScheme.signature.trust.ec.europa.eu is trusted (step9). Hence, the overall result of applying the trust policy to the electronic transaction is trusted and sent back to the verifier (step 10).

The basic structure of the information flow for the other basic scenarios is similar. For qualified seals, qualified identities, or qualified timestamps it is mainly the domain name structure which differs. For trust translation, and trust delegation there are in addition some additional steps required using the Trust Translation Authority and the Delegation Publisher, respectively.

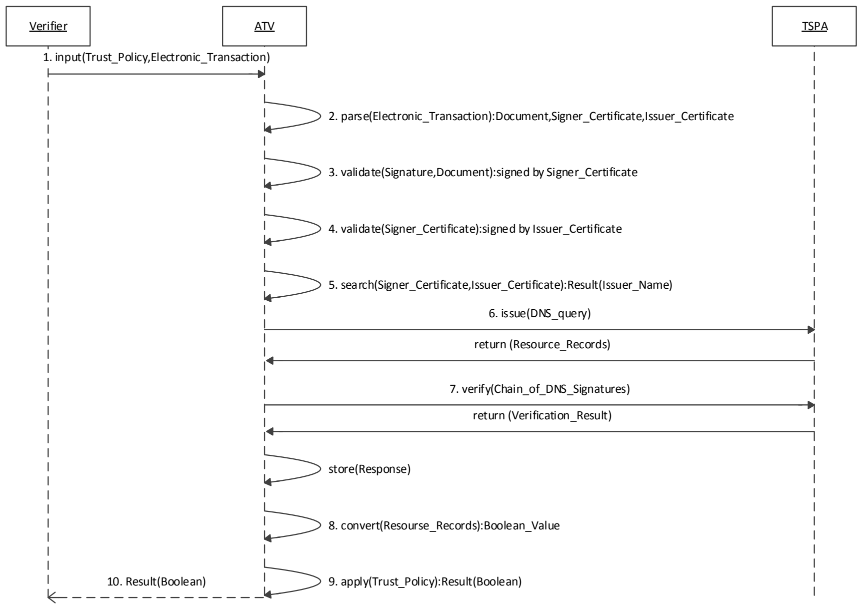

Fig. 2: Sequence Diagram for Trust Publication of a Qualified Signature (Boolean)

4 Trust Scheme Publication Authority

Knowing which trust scheme the issuer of the signers' certificate complies to is critical, in order to be able to verify whether an electronic transaction complies with the users' trust policy. It shows which security controls, and security requirements are fulfilled by the certificate issuer and thus indicate the security quality of the certificate that is used, e.g. for signing a document. The Trust Scheme Publication Authority (TSPA) is therefore an important component of the LIGHTest reference architecture. It enables discovery and verification of trust scheme memberships. Trust scheme publications are always associated with lists that indicate the membership of an entity with the referred to trust scheme. We refer to these lists in the following as trust lists. The described setup aligns well with existing trust list standards, which involve a trust list and a trust list provider (ETSI TS 119 612 [ET13]). In the sense of LIGHTest, the trust scheme provider can also be the trust list provider, or a party that is trusted by the trust scheme provider.

4.1 Trust Schemes and Trust Scheme Publications

A trust scheme itself can for example be constituted by requirements to information security processes, processes for issuance or revocation, requirements towards used technologies, or simply one single one-dimensional requirement, e.g. the geographical location of an entity. While some trust schemes, such as ETSI_EN_319_401 [ET16], just flatly lay out managerial requirements, trust schemes such as ISO/IEC 29115:2013 [IS13] further use different level of assurances to define which requirements must be met to comply with the trust scheme. In summary this all means, that a trust scheme can be published as a boolean trust scheme publication (e.g. [ET16]), and a ordinal trust scheme publication (e.g. [IS13]) (see Tab. 1). Boolean trust scheme publications indicate the entities that comply with the requirements of the trust scheme, and thus are a member of the trust scheme. Ordinal trust scheme publications indicate the entities that comply with the requirements of an ordinal aspect (e.g. a level of assurance) of the trust scheme.

Type of Trust Scheme Publication	Example	Verifiable Information
Boolean	ETSI_EN_319_401	Compliance of an entity to a trust scheme
Ordinal	LoA4.ISO29115	Compliance of an entity to an ordinal value of a trust scheme
Tuple-Based	{(authentication:2Factor), (identityProofing:inPerson)}	Requirements of a trust scheme

Tab. 1 Types of Trust Scheme Publications in LIGHTest

Both, Boolean and ordinal trust scheme publications do not provide any information on the requirements of the trust scheme, or the ordinal value (e.g. Level of Assurance) of the trust scheme that is represented by the trust scheme publication. In order to fill this gap, tuple-based trust scheme publications provide the requirements of a trust scheme in the form of attributes and values. For this purpose, the TSPA development foresees the development of a data model for trust schemes, that is able to provide a unified view on the requirements of trust schemes.

4.2 Concept for Trust Scheme Publication

The concept of the TSPA in LIGHTest consists of two components. It uses an off-the-shelf DNS Name Server with DNSSEC extension, in order to enable discovery of the Trust Scheme Provider that operates a Trust Scheme. The Trust Scheme Provider constitutes the second component of the TSPA. It provides a signed Trust List which indicates that a certificate Issuer is trusted under the scheme operated by the Trust Scheme Provider. It further provides the Tuple-Based representation of a Trust Scheme. As the DNS Name Server is only used to provide pointers to location of resources rather than storing the respective resources as DNS resource records directly, the TSPA is well-aligned with existing DNS practices. The use of pointers ensures the limited size of DNS messages, which is required for fast response times in the discovery process.

4.3 Publishing and Querying Trust Scheme Publications with the TSPA

The use of the DNS Name Server system by LIGHTest enables easy and widespread adoption of the approach. We assume that the trust scheme of a certificate issuer is unknown, upon receiving an electronic transaction. The TSPA therefore provides the capability to discover a trust scheme membership claim for a certificate issuer, and verify this claim. The discovery of a trust scheme membership claim is done by using the domain name resolution capabilities of the DNS Name Server. Fig. 3 provides an overview on the representation of trust scheme publications in the TSPA. The left side of the figure shows the data that is provided by the DNS Name Server.

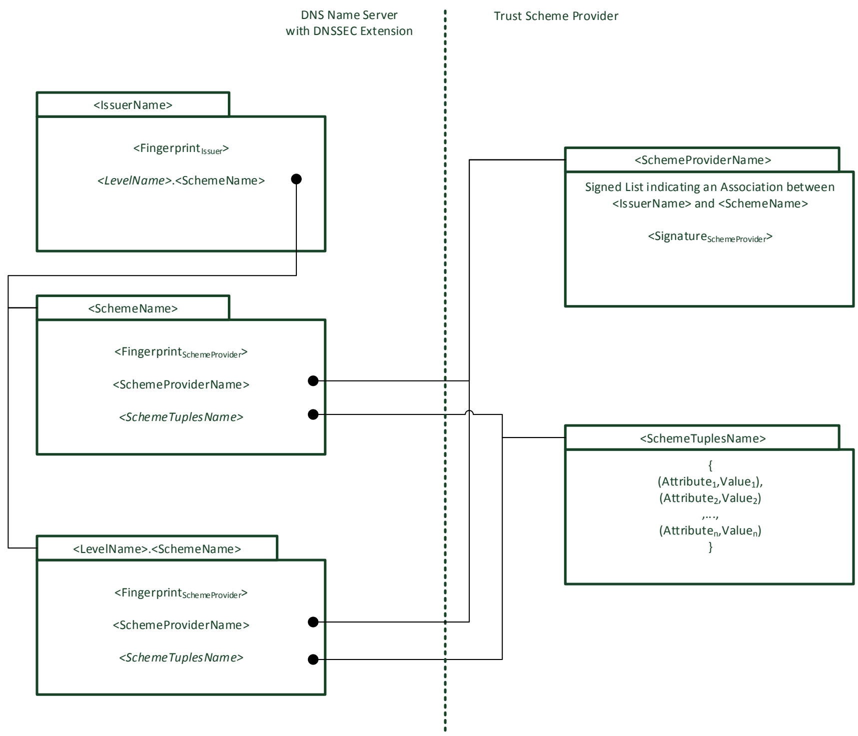

Fig. 3: Representation of Trust Scheme Publications in the TSPA

The DNS Server provides a pointer from the Issuer, indicated by <*IssuerName*> to a boolean (indicated by <*SchemeName*>) and/or ordinal trust scheme publication (indicated by <*LevelName*>.<*SchemeName*>). We refer to this information as the trust scheme membership claim of the issuer in the following, as it indicates the trust scheme publication that an issuer claims to comply with. The claim can be verified as to whether it is associated with the issuer certificate, by providing a fingerprint (e.g. a hash value) of the issuer certificate (indicated by <*Fingerprint*$_{Issuer}$>).

The DNS records of the boolean and ordinal trust scheme publication are always associated with the respective trust scheme provider. Therefore, these records include a fingerprint of the trust scheme provider certificate, indicated by <Fingerprint$_{SchemeProvider}$>. It further provides a pointer to the trust list that is operated by the trust scheme provider (<*SchemeProviderName*>). This trust list is signed with the trust scheme provider certificate and provides the verification of the trust scheme membership claim of the issuer. The claim is only true, if the issuer is listed on the trust list that is signed with the trust scheme provider certificate, and only if the fingerprint of the trust scheme provider certificate matches with the fingerprint stored in the DNS records of the boolean or ordinal trust scheme publication.

The tuple-based trust scheme publication is also accessible by querying the DNS Name Server, and retrieving the pointer to the tuple-based trust scheme publication (<*SchemeTuplesName*>). As the tuple-based trust scheme publication requires storing pairs of attributes and values, rather than pointers to records, it is stored on a different web component (e.g. a web server). The same of course holds for the trust list. This way, LIGHTest does not interfere with usual DNS usage, and thus with operational good practices of the DNS Name Server System.

As previously mentioned, tuple-based trust scheme publications require a unified data model, in order to be able to automatically query and process these trust scheme publications. Therefore, a still ongoing consolidation effort aims at consolidating the requirements of trust schemes in order to retrieve a data model that is using restricted and fully specified attribute domains. The latter means, that all possible values of an attribute must be previously known (e.g. integers in a certain range, or a finite set of strings).

5 Discussion and Outlook

The LIGHTest reference architecture and trust scheme publication authority (TSPA) support the implementation of the eIDAS Regulation ([EI14]). It enables the integration of existing trust lists using the global DNS infrastructure. Furthermore, it even expands eIDAS towards a global market and multi-users from the public and private sector. For the demonstration of the functionality of the LIGHTest infrastructure, two real world pilots are conducted within LIGHTest: In the first one, LIGHTest is integrated in an existing cloud based platform for trusted communication. In the second one, LIGHTest is integrated in an existing e-Invoicing infrastructure and application scenario. The DANE standard will be used to secure the transport protocols for retrieving information from the trust scheme provider.

The next steps will aim at the conceptualization of the DNS components, and the implementation of the TSPA, and the remaining infrastructure. Hereby, application of LIGHTest with currently available trust schemes and compliance with legal regulations is an important aspect for the uptake of the LIGHTest infrastructure. Therefore, future work will include the validation of the TSPA data models' capability to represent available trust schemes, beyond those used for modelling, and application of the LIGHTest infrastructure to scenarios including FIDO, and eIDAS.

6 Summary

There is a high need for assistance from authorities to certify trustworthy electronic identities due to the worldwide increasing amount of electronic transactions. Within the EU-funded LIGHTest project, a global trust infrastructure based on DNS is built, where arbitrary authorities can publish their trust information. In this paper, a high level description of the LIGHTest reference architecture, its components and their interactions are presented. In addition, the Trust Scheme Publication Authority, which enables discovery and verification of trust scheme memberships is introduced.

The reference architecture and the concept for Trust Scheme Publication Authority fulfil the main general principles and goals, which are required to develop a globally scalable trust infrastructure. Furthermore, it is well aligned with existing standards (e.g. ETSI TS 119 612) and fulfil the requirements using DNS name servers to build a global trust infrastructure.

Acknowledgments

This research is supported financially by the LIGHTest (Lightweight Infrastructure for Global Heterogeneous Trust Management in support of an open Ecosystem of Stakeholders and Trust schemes) project, which is partially funded by the European Union's Horizon 2020 research and innovation programme under G.A. No. 700321. We acknowledge the work and contributions of the LIGHTest project partners.

Bibliography

[AL12] Lee, P. A.; Anderson, T.: Fault tolerance: principles and practice. Springer Science & Business Media, 2012.

[BL16] Bruegger, B. P.; Lipp, P.: LIGHT^{est} – A Lightweight Infrastructure for Global Heterogeneous Trust Management. In: Hühnlein D. et al (Hg.): Open Identity Summit

2016, Rome: GI-Edition, Lecture Notes in Informatics. S. 15-26.

[Di76] Dijkstra. E. D.: A discipline of programming. Prentice Hall, ISBN 978-0-13-215871-8, p. 56, 1976.

[EI14] European Parliament, 'Regulation (EU) No 910/2014 of the European Parliament and of the Council of 23 July 2014 on electronic identification and trust services for electronic transactions in the internal market and repealing Directive 1999/93/EC', European Parliament, Brussels, Belgium, Regulation 910/2014, 2014.

[ET13] ETSI: Electronic Signatures and Infrastructures (ESI); Trusted Lists. Sophia Antipolis Cedex, France, Technical Specification ETSI TS 119 612 V1.1.1, 2013; http://www.etsi.org/deliver/etsi_ts/119600_119699/119612/01.01.01_60/ts_119612v010101p.pdf.

[ET16] ETSI: Electronic Signatures and Infrastructures (ESI); General Policy Requirements for Trust Service Providers. ETSI, Sophia Antipois Cedex, France, European Standard ETSI EN 319 401, 2016; http://www.etsi.org/deliver/etsi_en/319400_319499/319401/02.01.01_60/en_319401v020101p.pdf.

[IS13] ISO/IEC: Information technology -- Security techniques -- Entity authentication assurance framework. ISO/IEC, Geneva, CH (2013).

[Pa72] Parnas, D. L.: On the criteria to be used in decomposing systems into modules. In: Communications of the ACM. Vol. 15, Nr. 12, ISSN 0001-0782, p. 1053–1058, 1972.

[HS12] Hoffman, P.; Schlyter J.: The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 0.17487/RFC6698, 2012, http://www.rfc-editor.org/info/rfc6698, 01.06.2017.

[Gu14] Gudmundsson, O.: Adding Acronyms to Simplify Conversations about DNS-Based Authentication of Named Entities (DANE)", RFC 7218, DOI 10.17487/RFC7218, 2014, http://www.rfc-editor.org/info/rfc7218, 01.06.2017.

[Sh95] Shaw, M.; DeLine, R.; Klein, D. V.; Ross, T. L.; Young, D. M.; Zelesnik, G.: Abstractions for software architecture and tools to support them. In: IEEE Transactions on Software Engineering. Vol. 21, Nr. 4, ISSN 0098-5589, p. 314–335, 1995.

[TS07] Tanenbaum, A. S.; Van Steen, M.: Distributed systems. Prentice-Hall, p.3-16, 2007.