



Input paper: XXXX-n.n.n

Input paper for the following Committee(s):

- | | | |
|--|------------------------------|------------------------------|
| <input type="checkbox"/> ARM | <input type="checkbox"/> ENG | <input type="checkbox"/> PAP |
| <input checked="" type="checkbox"/> ENAV | <input type="checkbox"/> VTS | |

Purpose of paper:

- | |
|---|
| <input type="checkbox"/> Input |
| <input checked="" type="checkbox"/> Information |

Agenda item**Technical domain/ Task number****Author(s)/Submitter(s)**

GLA, DLR, and AIVeNautice

THE DECENTRAL TRUST SYSTEM OF THE MARITIME CONNECTIVITY PLATFORM

1. INTRODUCTION

The purpose of this document is to explain the motivation, design and application of the Decentral Trust System of the Maritime Connectivity Platform. It is intended as a (not too technical) general introduction for the members of the MCP Consortium and the wider maritime sector with an need to trusted communication.

The document will provide the background to the area of trust systems by-large and highlight the drawbacks and trade-offs that is required in such a system. Based on this we will give the reasoning behind the architectural design choices of the MCP Decentral Trust System. We will finally outline the resulting architecture and the current development.

1.1. MOTIVATION

Trust management systems are frameworks for facilitating trust relationships between digital identities in digital environments. At the core, their common goal is to ensure that communication between entities in a network is secure and unmodified. In a digital network, this goal is often achieved by means of cryptographic protocols implementing processes such as authentication and authorization.

Authentication is the process of verifying the identity of an entity in a network; that is, ensuring that the entity is who it claims to be. Following this, authorization determines what operations an authenticated entity is allowed to perform. These processes are often facilitated by cryptographic protocols, which use mathematical algorithms to secure communication, maintaining its integrity and confidentiality.

The importance of trust management systems has grown steadily in the digital age, and the rise of distributed systems, cloud computing, and Internet of Things (IoT) further necessitates scalable and reliable solutions. With an increasing amount of sensitive information and critical operations being conducted digitally, the need for effective trust management has only become more crucial in recent years.

This is also seen in the maritime sector. Companies and authorities here work in a fragmented environment with many different systems governed under different rules. Different areas of the maritime sector reports to different authorities. Ships and equipment can be spread around the globe communicating to both central and local entities.

1.2. TRUST: GOAL OF THE SYSTEM DESIGN

Trust is central to all contemporary systems and concepts that follow, as well as the trust system for MCP this document outlines. Trust is a set of **axiomatic assumptions** on which all further security in the system depends on, be it organizational or technical.

One such assumption for asymmetric cryptographic systems (such making use of public / private key pairs) is the secure handling of private key material by each party to the extend that one can assume only the owner of a keypair has access to its private key. Another is the assumption that certain operation are easy to perform, but hard to reverse (e.g., calculating hashes or creating/verifying signatures), thus enabling the concept of immutability (what has been hashed and/or signed cannot be changed without invalidating the hashsum and/or signature). The goal of the trust system is to minimize degree and number of axiomatic assumptions needed to make a system work securely.

1.3. APPLICATION AND OVERALL DESIGN

The Decentral Trust System described in this document will be an extension/implementation over the MCP Maritime Identity Registry and the P3KI core trust system. This with an extension with a maritime/general trust policy language. The resulting system resulting it designed with applications within the maritime section in mind.

First use cases will especially have a focus in relation and role management between maritime authorities, actors, and services. However, this work is not limited to this. The general concept has just a much application to equipment and resource management within and without the maritime sector; though this can require extensions over this first version.

2. BACKGROUND

2.1. PKI SYSTEMS, CENTRALISATION, AND THEIR DRAWBACKS

Public key infrastructures (PKIs) are among the most widely employed systems to facilitate trust management systems. As the name implies, they do so by managing the distribution of public keys for asymmetric cryptography, enabling secure communication, identity verification, and secure access to resources in a network. However, PKIs and other established trust systems have several shortcomings:

Single point of failure. Most PKIs (not including PGP [4]) are centralized and follow a hierarchical model with a so-called Certificate Authority (CA) at the root. In these PKIs, the CA represents a single point of failure, and if the CA is compromised, the entire trust network can be at risk. A notable example of a CA being compromised is the DigiNotar breach in 2011, resulting in intercepted and decrypted secure communications on a significant scale.

Lack of offline functionality. Most trust management systems require an active network connection for interaction. This can be limiting in maritime environments with inconsistent network connectivity, or underground scenarios such as mining operations, where internet connections are often unstable or absent.

Lack of specificity. In PKIs, trust is typically represented in a binary fashion; an entity is either trusted or not trusted, based on whether it has a valid certificate. This simple model fails to capture the more nuanced trust relationships that might be needed in certain situations. For instance, an entity might be trusted for certain actions but not others, or its trustworthiness might be different in different contexts.

Certificate revocation. Revoking certificates, especially in a timely manner, is a well-known problem in PKIs. The current mechanisms, such as certificate revocation lists (CRLs) [13] or Online Certificate Status Protocol (OCSP) [31] [30], have been criticized for their inefficiency and scalability issues.

Some, but not all, of these shortcomings can be mitigated by employing a non-hierarchical web-of-trust model instead.

2.2. MARITIME CONNECTIVITY PLATFORM

The Maritime Connectivity Platform (MCP)¹ is a decentralised platform that facilitates secure and reliable information exchange within the maritime domain and beyond [11] [12]. Beyond – because the maritime world is not isolated, but need to exchange information with other domain – for instance with other transport domains.

The information exchanged can be almost of any nature, ranging from private confidential information between a vessel and the shore office of the shipowner, to public information provided by authorities, such as the provision of navigational warnings.

As a decentralised platform, there is no single entity operating this. Several organisations are MCP service providers, and collectively they form “the Maritime Connectivity Platform”.

2.2.1. MARITIME IDENTITY REGISTRY (MIR)

The prerequisite for the digitalisation of the maritime domain is a trustworthy provision of digital services for information exchange. [27] The MCP features - as one of its core components - an identity registry, where all entities that wish to exchange information are registered and have a digital certificate issued to them. Thus, a vessel registered with the MCP identity registry (having a digital certificate issued from it), can authenticate itself (cryptographically prove its identity) to the VTS centre, and thus provide data to the VTS centre which the VTS centre can trust the origin of. The principle of authentication is a cornerstone in contemporary digital solutions.

2.2.2. MARITIME RESOURCE NAME (MRN)

In MCP MRN will be used as an unique identifier for entities. The MCP namespace is a subspace [27] of the Maritime Resource Name (MRN) space [26], which is an official URN namespace.

It must be ensured that

- each IPID refers to at most one Identity Service Provider,
- each Identity Service Provider must ensure to respect all syntax prescribed in the MRN specification, and
- each Identity Service Provider must ensure that each IPSS of their name space refers to at most one entity of their domain.

This defines uniqueness in the sense that one MCP MRN is assigned to at most one entity. This is a general requirement for any URN.

2.2.3. MCP IDENTITY PROVISIONING

There are two aspects of MCP identity provisioning [27]:

1. Identity Management: A MIR enables that each maritime entity (such as a device, human, organization, service, ship, etc.) can be registered as a participant of the MCP and be equipped with a unique identifier. The identifier is given in terms of an MRN (Maritime Resource Name [26]). While MIR governance harmonizes the MRN namespace governed by the MCP Consortium (MCC) and sets out criteria for the registration process, it is up to the MIR services to implement and have certified concrete identity registries. The following terminology

¹<https://maritimeconnectivity.net/>

- MCP entity: An entity registered at some MIR services.
 - MCP namespace: The subspace of the MRN namespace that is governed by the MCC.
2. Public Key Infrastructure (PKI): The MIR enables that each MCP entity holds a cryptographic identity in terms of a public/private key pair and a certificate bound to their MRN identifier within the MCP. The cryptographic identity of a MCP entity will change over time (due to updates of key material), but the MRN identifier must be unchanged over this certificate change.

2.3. P3KI

P3KI is short for either P3KI GmbH, a Berlin, Germany based company or P3KI Core, its main technological solution. P3KI Core is a web-of-trust authentication and authorization system. It offers to flexibly delegate mandates in the form of permissions, capabilities, or roles and is aimed to be operated in fully decentralized and offline scenarios.

- P3KI Explained: Decentralized Offline Authorization for IoT (v1.3)
- P3KI Core: Decentralized Access Delegation for Critical Infrastructure (v1.2)

3. USE CASES

This section will describe possible overall application and potential use cases relating to the trust system in a MCP context. This is in no way intended to be an exhaustive list, should only be used as inspiration.

3.1. VIRTUAL ATONS

A general application would be to define virtual AtoNs as an component of the MCP Trust System.

The general alternative approach would be to issue policy-based certificates. This means that the issuer have to provide information in the metadata of the certificate of how the certificate is to be used. All applications must then know this specific semantics of using the certificate (of class of certificates). This comes with the following limitations:

- All certificate policies must be clearly standardised across all possible virtual AtoNs.
- Applications and devices must be implemented with these specific policies.
- Updates to policies (new areas where the usage can be beneficial)
- Dynamic usage of the certificates is impossible. A virtual AtoN that is created for an emergency can be used only for that specific situation. There will be a need for many certificates for many possible situations.

Using the MCP trust system can alleviate many of these limitations. Here a maritime authority (or others) can issue a few certificates and their usage (roles), time-limits, geographical limits, etc. will then be defined by policies in the MCP trust system.

This does not mean that IALA (or other organisations) cannot publish guidelines that define what a virtual AtoN is and how it works. But the guidelines can focus more on the policies and requirements, than on technical implementations. These will be handled by the definition of simple trust policies.

4. ARCHITECTURE

The establishment of a global communication platform for international maritime shipping presents unique challenges, particularly in implementing Public Key Infrastructure (PKI) to ensure non-repudiation, message integrity, and confidentiality. Classic PKI systems, while widely used for secure communications, face several limitations when applied to such a global and diverse context.

4.1. OUTLINE

Describing and formalising the governance and technical challenges of maritime shipping is complex. No other environment is built on trust and carefully forged relationships over many years to comparable extents. Lifting this trust and the associated relationships into the digital domain to enable secure automation without losing flexibility and sovereignty is, therefore, a major undertaking rife with pitfalls and dead-end strategies.

Adopting established and best-practice solutions proven within large-scale IT systems into the maritime environment is bound to fail. Similar constraints exist in industrial automation, where IT solutions cannot be directly adopted to OT (operation technology) environments, because they violate fundamental security assumptions of the target environment. However, with maritime shipping, the key impediment is a mismatch in governance realities first, and operational realities second. The main causes, though, are very similar:

- Introduction of centrally or online-managed technologies (e.g., IAM).
- Introduction of technologies that are theoretically applicable but can only be effectively maintained through centralized automation (e.g., PKI).
- Technological category errors.

Maritime shipping is an inherently pluralistic environment defined by relationships between many sovereign parties. Therefore, it will be impossible to find a single party everyone will ultimately and forever trust. And this is exactly what is required when setting up a IAM and classic PKI.

While most participants in the maritime sector will agree that they trust bodies like IMO, IALA, or similar, the implicit question here is always “with what exactly and to what extend?”. And while those bodies are perfect venues to figure out how to talk to each other, no sovereign will subjugate themselves under these bodies allowing them to decide who they have to trust. And this is exactly what would be required if those bodies were to operate central security related systems for the maritime sector.

We, therefore, propose a decentralized trust system concept, outlined in this paper. Such a system will be able to allow each sovereign to selectively decide who to trust and to what extend, without sacrificing either security or flexibility.

4.2. CONCEPTS

4.2.1. GLOBAL MARITIME CONTEXT

Classic PKI systems, such as X.509, transform the fundamental trust problem into a key distribution problem. PKI systems are known for their centralized management and need for online connectivity to validate certificates fully. The only feasible way of scaling such a PKI is centralization and automation, implying regular connectivity to all participants. The likely even bigger implication is the trust required towards that central organizational structure. It is unlikely one is able to find a single global party trusted by everyone with effectively their identities and communication security.

Even if the number of global parties that would need to be trusted is limited, PKI is not well equipped to express exactly what those parties are trusted with. Furthermore, inter-operation across multiple, disparate PKI hierarchies, each with their distinct roots of trust, is one of the hard problems of our time.

In a global maritime communication platform, these requirements can lead to scalability and manageability problems due to the vast number of entities, the distributed nature of operations, and general mistrust between the participants in global trade.

Pajoo et al. [19] have discussed how the use of a multi-layer blockchain architecture, to mitigate some of these issues by providing a decentralized approach to identity and permission management, which highlights the scalability challenges faced by traditional PKI systems in expansive networks.

However, similar to classic PKI systems, blockchain and distributed ledger technologies (DLT) necessitate online connectivity for nodes to participate in the network, refresh their copy of the ledger, or add to it. Also, cryptographic keys need to be exchanged up front, requiring a high degree of planning and inflexibility in the field to react to unforeseen changes. While blockchain does address some of the issues around the need of roots of trust, its other trade-offs, like an ever growing data structure representing a shared global truth, can become prohibitive.

The problem class itself is not unknown and several approaches have been discussed and implemented in the past. Garfinkel (1995) documented PGP as a system for making identities around email addresses verifiable using a web-of-trust approach. With Blaze et al. (1996) identifying the shortcomings of PKI in the context of decentralized trust management and providing their own solution by introducing a policy-based permission system.

What might be perceived as a possible shortcoming of web-of-trust approaches, the lack of central management, is at the same time enabling large scale operation and alleviating the need for centrally trusted peers.

Historically, global trade always and heavily relied on humans trusting each other, as is outlined in the seminal works of Francis Fukuyama (1995). With trust being one side of the coin and vulnerability the other, it is paramount to precisely scope the degree and context of a given trust relationship.

We therefore propose a web-of-trust approach to address the challenges of securing communications for global maritime exchange. As with all technologies applied to problems, there is never a solution, but instead a mapping of one problem to a hopefully lesser or at least better manageable problem. A key element of this paper is to enumerate the trade-off decisions made while addressing the requirements and weighing their implication.

4.2.2. TRUST VS. AUTHORIZATION

Trust in the context of PKI, web-of-trust, or IT systems in general is a combination of many factors. Fundamentally it describes the axioms on which we build our systems. We trust the mathematical principles of cryptography to hold true. We trust anyone using anything that needs to stay secret is also sufficiently protecting it, so others cannot access it. We trust known rules and procedures are followed by parties issuing certificates.

If and only if those axioms are fulfilled, can we use the capabilities of the system created as such to do actual work. Capabilities such as verifiable authorization of a third party to act on ones behalf or assigning an approved and verifiable identity to a party. Trust and authorization are so closely linked on a conceptual level, that they are often used interchangeably in the form of “I trust you to do X” to express, what is actually an authorization.

What the decentralized trust system aims to offer, is the capability to make such statements third-party verifiable in an environment, where the party that made the original statement cannot be present nor actively contribute to the verification of said statement.

Some of these axioms are closely tied to properties and our understanding of the universe, while others are organizational and procedural. The latter, we can influence. This in turns enables us to adjust the level of trust we need to put into any party to a usable minimum within the given context.

References

- [1] The specification of e-navigation technical services, 2018. <https://www.iala-aism.org/product/g1128-specification-e-navigation-technical-services/>, edition 1.2.
- [2] ISO/TC 22/SC 31. Road vehicles - security certificate management, July 2006. <https://www.iso.org/standard/41891.html>.
- [3] IEC 63173-2:2022. Maritime navigation and radiocommunication equipment and systems - data interfaces - part 2: Secure communication between ship and shore (secom). <https://webstore.iec.ch/publication/64543>.
- [4] D. Atkins, W. Stallings, and P. Zimmermann. PGP message exchange format, August 1996. <https://www.rfc-editor.org/rfc/rfc1991>.
- [5] M. Blaze, J. Feigenbaum, and J. Lacy. Decentralized trust management. In *Proceedings 1996 IEEE Symposium on Security and Privacy*, pages 164–173, 1996.
- [6] C. Bormann. Well-known URIs for the websocket protocol. <https://www.rfc-editor.org/rfc/rfc8307>.
- [7] S. Bradner. Key words for use in RFCs to indicate requirement levels. <https://www.rfc-editor.org/rfc/rfc2119>.
- [8] Tim Bray. The JavaScript object notation (JSON) data interchange format. <https://www.rfc-editor.org/rfc/rfc8259>.
- [9] J. Callas, L. Donnerhake, H. Finney, D. Shaw, and R. Thayer. OpenPGP message format, November 2007. <https://www.rfc-editor.org/rfc/rfc4880>.
- [10] David Chadwick, Gansen Zhao, Sassa Otenko, Romain Laborde, Linying Su, and Tuan Anh Nguyen. Permis: a modular authorization infrastructure. *Concurrency and Computation: Practice and Experience*, 20(11):1341–1357, 2008.
- [11] MCP Consortium. Maritime identity registry of the maritime connectivity platform. <https://maritimeconnectivity.net/>.
- [12] MCP Consortium. The mcp concept document – conceptual overview. <https://maritimeconnectivity.net/mcp-documents/>.
- [13] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housley, and W. Polk. Internet X.509 public key infrastructure certificate and certificate revocation list (CRL). <https://www.rfc-editor.org/rfc/rfc5280>.
- [14] C. Ellison. SPKI requirements, September 1999. <https://www.rfc-editor.org/rfc/rfc2692>.
- [15] C. Ellison, B. Frantz, B. Lampson, R. Rivest, B. Thomas, and T. Ylonen. Spki certificate theory, September 1999. <https://www.rfc-editor.org/rfc/rfcSDSI>.
- [16] S. Farrell and R. Housley. An internet attribute certificate profile for authorization, April 2002. <https://www.rfc-editor.org/rfc/rfc3281>.
- [17] I. Fette and A. Melnikov. The WebSocket Protocol. <https://www.rfc-editor.org/rfc/rfc6455>.
- [18] K. Hamilton-Duffy, R. Grant, and A. Gropper. Use cases and requirements for decentralized identifiers, March 2021. <https://www.w3.org/TR/did-use-cases/>.

- [19] Houshyar Honar Pajoo, Mohammad Rashid, Fakhru Alam, and Serge Demidenko. Multi-layer blockchain-based security architecture for internet of things. *Sensors*, (3), 2021.
- [20] Gregor Jehle. Cross-pki web-of-trust als enabler für zusammenarbeit, 21. September 2023. https://www.teletrust.de/fileadmin/user_upload/07_TeleTrust-EBKA_PKI-WS_Jehle_P3KI.pdf.
- [21] Michael Jones. JSON web algorithms (JWA). <https://www.rfc-editor.org/rfc/rfc7518>.
- [22] Michael Jones and J. Bradley. JSON web signature (JWS). <https://www.rfc-editor.org/rfc/rfc7515>.
- [23] Michael Jones, J. Bradley, and N. Sakimura. JSON web token (JWT). <https://www.rfc-editor.org/rfc/rfc7519>.
- [24] P. Leach, M. Mealling, and R. Salz. A Universally Unique Identifier (UUID) URN namespace. <https://www.rfc-editor.org/rfc/rfc4122>.
- [25] Google LLC. Protocol Buffers Documentation. <https://protobuf.dev/>.
- [26] Management of Maritime Resource Name Organization Identifiers. IALA guideline G1164, ed 1.1. <https://www.iala-aism.org/content/uploads/2022/09/G1164-Ed1.1-Management-of-Maritime-Resource-Name-Organisation-Identifiers-December-2021.pdf>.
- [27] IALA Guideline on the Provision of MCP Identities. IALA guideline G1183, ed 1.0, Jun 2024. <https://www.iala-aism.org/>.
- [28] E. Rescorla. The transport layer security (TLS) protocol version 1.3, 2018. <https://www.rfc-editor.org/rfc/rfc8446>.
- [29] Ronald L. Rivest and Butler Lampson. SDSI – a simple distributed security infrastructure. <http://people.csail.mit.edu/rivest/pubs/RL96.ver-1.1.html>, 1996.
- [30] Mohit Sahni. Online Certificate Status Protocol (OCSP) Nonce Extension, 2020. <https://www.rfc-editor.org/info/rfc8954>.
- [31] S. Santesson, M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. X.509 internet public key infrastructure online certificate status protocol - OCSP. <https://www.rfc-editor.org/rfc/rfc6960>.
- [32] M. Sporny, D. Longley, M. Sabadello, D. Reed, O. Steele, and C. Allen. Decentralized identifiers (dids) v1.0, July 2022. <https://www.w3.org/TR/did-core/>.
- [33] Charikleia Zouridaki, Brian L. Mark, Kris Gaj, and Roshan K. Thomas. Distributed ca-based pki for mobile ad hoc networks using elliptic curve cryptography. In Sokratis K. Katsikas, Stefanos Gritzalis, and Javier López, editors, *Public Key Infrastructure*, pages 232–245, Berlin, Heidelberg, 2004. Springer Berlin Heidelberg.

A. A BRIEF HISTORY OF CONCEPTS AND TECHNOLOGIES

No idea exists in a vacuum and many novel approaches are standing on the shoulder of giants. We give a brief history of decentral trust systems.

A.1. X.509 AND THE BIRTH OF PKI (1988 TO 1997)

X.509 [13] is the oldest technology discussed here and to date the industry standard for managing keys, assigning identities. Such systems are generally known as public-key infrastructure (PKI). X.509 generally addresses the matter of **authentication**, answering the question of “who are you?”. This is done based on a strongly hierarchical model with trusted certificate authorities at the top, intermediate certificate authorities in between, and ultimately certificate holders who cannot issue certificates of their own. X.509 works well within an organization and can address large-scale systems effectively if centralized automation and orchestration is possible. While the authenticity of certificates can be verified easily, regardless of the context one is in, full verification of certificate validity is a process strongly tied to central infrastructure, making it effectively an online-only operation in most cases. Trust is based in certificate authorities to the extent that they follow known and published procedures for verifying identities of parties they issue certificates to as well as the availability of parties required for verification processes (e.g. certificate revocation lists (CRL) or open certificate status protocol (OCSP) servers).

A.2. PGP INTRODUCES THE IDEA OF WEB-OF-TRUST (1991 TO 1996)

PGP [4] and later OpenPGP [9] introduced the concept of a web-of-trust for **authentication**. Contrary to the hierarchical organization approach of X.509, PGP allowed everyone to make cryptographically signed statements about the identity of another party, forming a mesh or web structure. Its main use-case is a peer-to-peer driven attestation of ownership over an identity, usually in the form of an email address. Trust is based in each peer that they perform identity checks based on known and published procedures to an acceptable degree. PGP’s model establishes a implicit transitive trust concept where the more trusted peers attesting your identity you have and the more trusted peers they have in turn attesting their identity, the more trusted you are. Ideally, there is an explicit chain of trust connecting you to whatever party you want to authenticate.

A.3. SDSI / SPKI ADDRESSES DECENTRALIZATION (MID-1990’S TO 1999)

For the first time SDSI [29] stepped away from authentication as the main goal and focused primarily on **authorization**, answering the question of “what are you allowed to do?”. To enable authorization in decentralized contexts, a trade-off needed to be made by relaxing rules around naming. Instead of depending on an authority to assign unique names (as is done with X.509), local namespaces were introduced. These dynamically scoped names into identifiers that made sense in specific contexts, but were not guaranteed to be globally unique. SDSI / SPKI also introduced the concept of flexible delegation, allowing everyone to hand off permissions to their peers in a traceable and verifiable manner. Trust is based in locally selected peers, based on context and risk is reduced by employing a policy language specifying what permissions are granted to which resource.

A.4. DISTRIBUTED CA FOR MANET APPLICATIONS (2004)

Zouridaki et al [33] tried to tackle the decentralized and ad-hoc nature of MANETs (mobile ad-hoc networks) in their 2004 paper. Their proposal is based on distributing the actual PKI over multiple nodes and clusters of nodes, then using threshold cryptography to determine whether the issuance or revocation of a certificate was legitimate. This reduces the trust that needs to be put into a single node over a larger number of participating nodes, assuming that compromising a large enough number of nodes to negatively effect the PKI is low enough. The fundamental concepts, however, are similar to X.509, with the PKI still purely focused on **authentication**.

A.5. PERMIS AIMS TO ENABLE AUTHORIZATION ON TOP X.509 (2008)

Chadwick et al. [10] identified clearly, that X.509 is meant for **authentication** and lacks means to properly handle **authorization** tasks. They designed the PERMIS system which uses X.509 attribute certificates [16] containing the authorization policy statements which are in turn stored in an LDAP system or similar means for making the attribute certificates accessible. While addressing the challenge of authorization, implementation is based on the same limiting factors of centralized systems and orchestration as the original X.509, since attribute certificates are handled using the same procedures.

A.6. DISTRIBUTED IDENTIFIERS AIM TO ADDRESS DECENTRALIZED NAMING (MID-2010'S TO 2021)

Distributed identifiers (DIDs) are honourably mentioned here as one of the more recent approaches to **authentication**, or more specifically the age old problem of “naming things is hard, naming things uniquely is harder”, but in decentralized contexts. Decentralized in this context specifically means independence of a central authority for issuing names. Instead the canonical approach is to publish claims to a name along side some cryptographic material into an authoritative database. This makes it an inherently online-based system for all practical purposes. Trust is based into the immutability of the information exchange medium and its availability.

B. P3KI DECENTRAL TRUST

B.1. MOTIVATION

P3KI's design and ideas are deeply rooted in years of real-world cyber security experience and a deep understanding of common failure classes. A key inspiration was – the then new, now withdrawn – ISO 20828 [2], which excellently identified the challenge of decentralized use-cases, without fully specifying a way to address them. While not a direct inspiration originally, SDSI / SPKI found similar answers to common questions, confirming P3KI's approach.

P3KI's aim is also not to be a panacea to all possible issues with PKI. Instead, as with all technical solutions, a careful selection of trade-offs have been made, making it a perfect solution to a number of use-cases, but a suboptimal fit to others. Target scenarios are defined by usually large number of nodes (scalability), ad-hoc communications partners (hard to plan ahead and orchestrate), unreliable connectivity up to fully offline nodes (no central provisioning or central authoritative services possible) and a general lack of a single trusted party.

B.2. EXPLICIT TRADE-OFFS

Naming things uniquely at a global scale is a fundamentally different problem to granting authorizations. Therefore, both challenges should be addressed with tailored approaches. The former is an okay use-case for classic identity system based on PKI. Identities do not usually change that often. The most common reason for revoking a certificate is compromised key material leading to identity impersonation.

Authorizations usually are shorter lived than the identities they are attributed to. If the authorization system allows authorizations to be granted, modified, and revoked more easily or more reliably than using the analogous processes for identity management, this not only makes regular operations more efficient, but also enables the authorization system to act as a mitigation in the failure case of the identity system. An impersonated identity without authorization can do little harm.

B.3. HYBRID AUTHENTICATION AND AUTHORIZATION SYSTEM

This enables a hybrid authentication / authorization approach [20]. A classic PKI and certificates are used for assigning longer lived identities to cryptographic keys. The cryptographic keys are then used in a web-of-trust overlay to dynamically grant authorizations.

With this setup, verification happening in offline scenarios will be primarily focused on authorization. Authentication based on a given authorization is still possible and required. However, the authentication focuses on the question whether the party can demonstrate control of the expected private key material. Whether the assignment of a specific name to that key is also valid, is a secondary question.

Furthermore, such a system does not need to be based around a single PKI hierarchy. For as long as all participating PKIs have consensus on the cryptographic algorithms used, the hybrid system enables authorizations to be granted back and forth dynamically between participants under different PKI hierarchies.

While this addresses the general case of authorization and the ability to “pick and choose” who you want to trust for specific authorizations it opens the door to address one of the original challenges posed to the MCP trust system: who do I trust to assemble the collection of trustworthy PKI root certificates? This use-case is just a specific scenario with an authorization for “providing roots for identity verification”, enabling flexible automation of the process after an initial bootstrap.