Exploring the Use of SSL/TLS Certificates for Identity Assertion and Verification in Ethereum

Friederike Groschupp, 18.05.2020, Final Presentation Master's Thesis

Chair of Software Engineering for Business Information Systems (sebis)
Faculty of Informatics
Technische Universität München
wwwmatthes.in.tum.de

Outline

1. Motivation, Problem Statement, Research Questions

2. RQ1: How can we enable on-chain decisions on identity using TLS certificates?

3. RQ2: How can we use SSL/TLS certificates to endorse Ethereum addresses?

4. Evaluation of the System

5.Conclusion and Future Work

Motivation

Authenticating Ethereum address owners **enables new applications** and on **promotes trust** in information and services provided

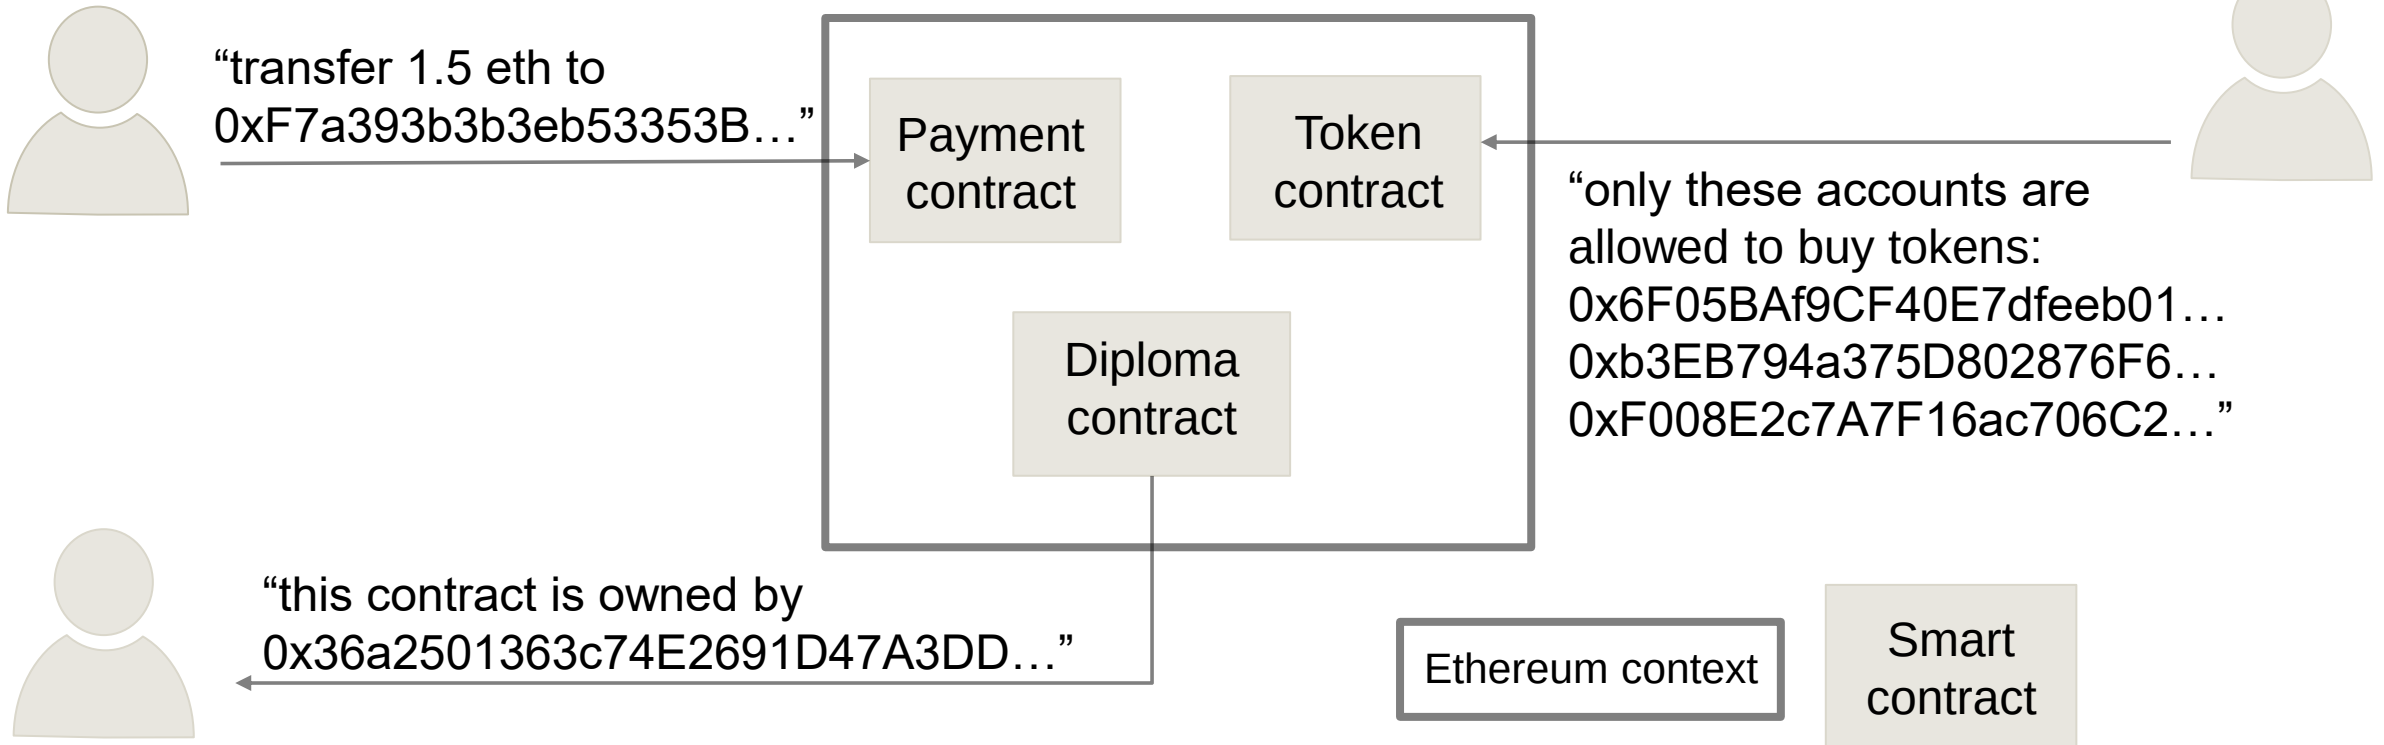

Motivation

Authenticating Ethereum address owners **enables new applications** and on **promotes trust** in information and services provided

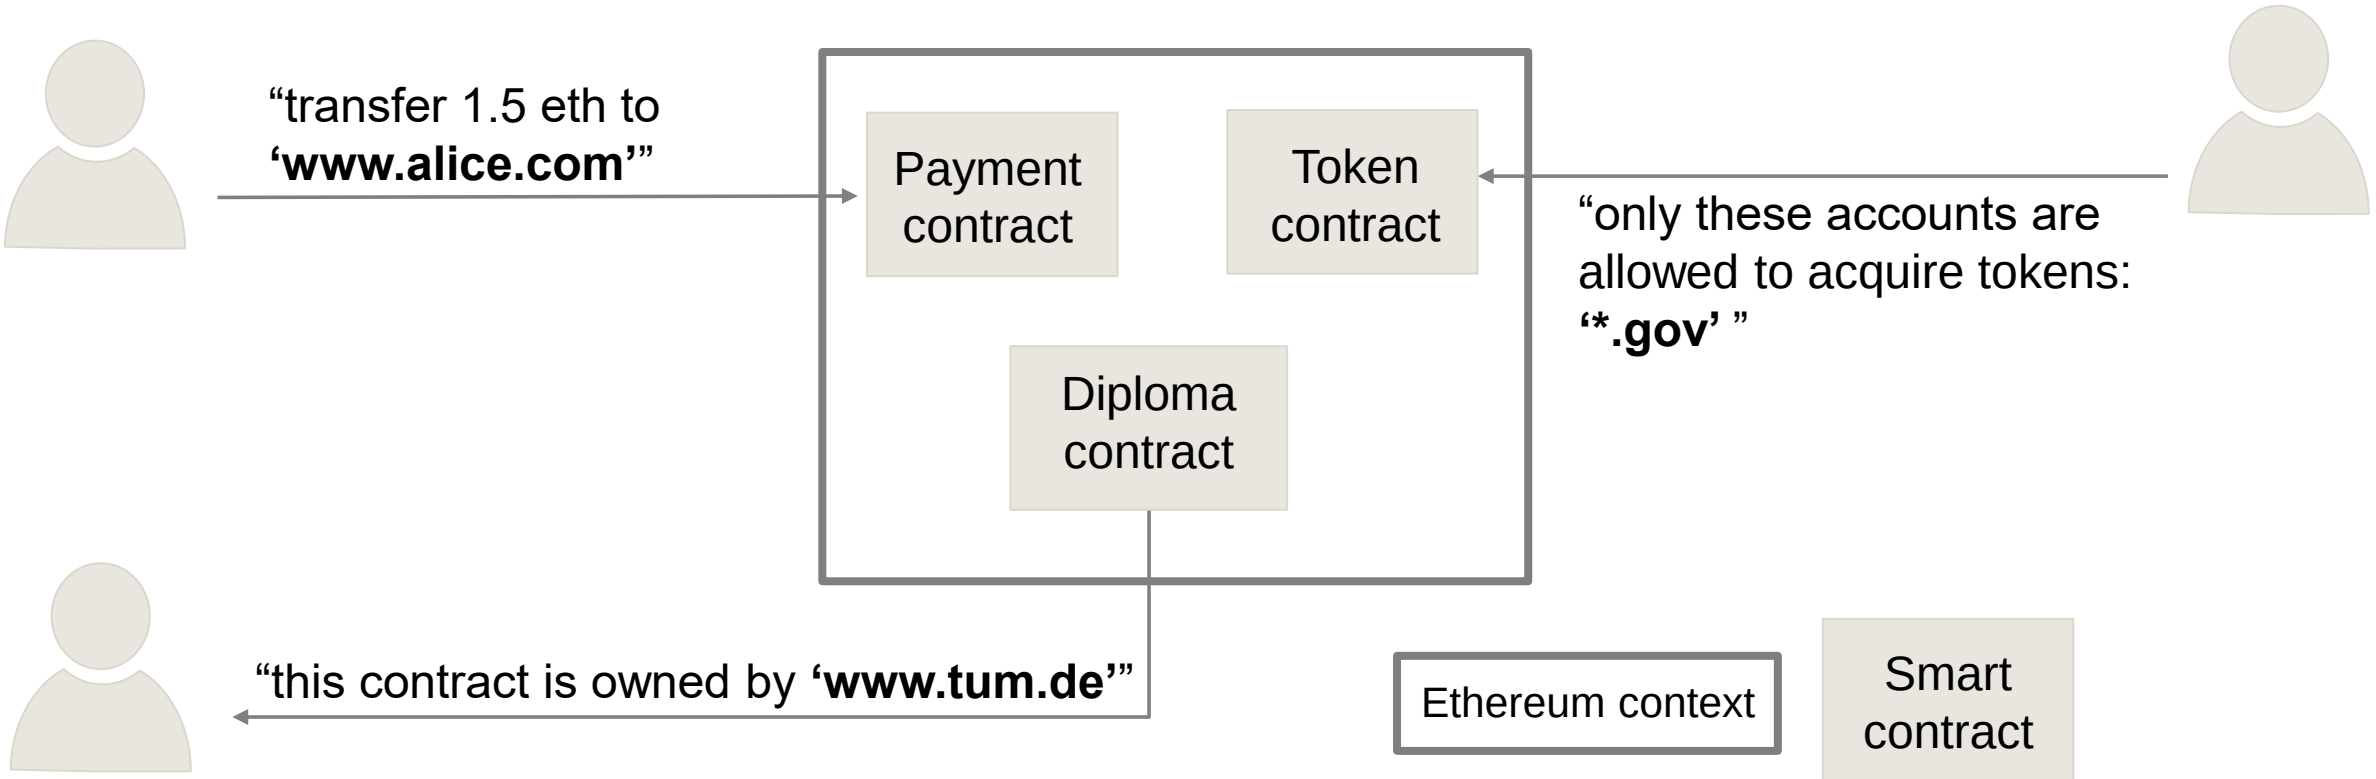

Motivation

Authenticating Ethereum address owners **enables new applications** and on **promotes trust** in information and services provided

Slow adoption of identity solutions for Ethereum due to **lack of trusted information**

Leverage **established TLS certificates and public key infrastructure** (PKI)

Basic Idea

Traditional authentication on the internet using TLS certificates

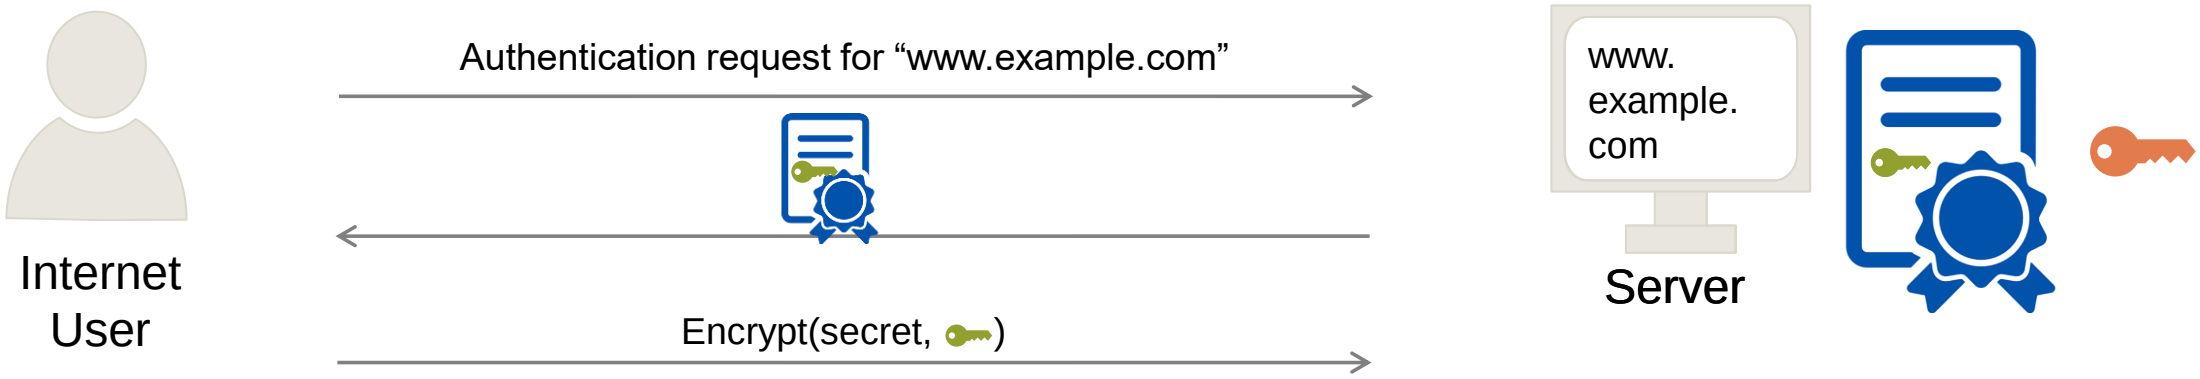

Authentication on Ethereum using TLS certificates

Problem Statement

TLS validation depends on availability of the server and subjective factors, **on-chain** identity validation requires **determinism**

TLS certificate system was **not designed with our use case in mind**

Research Questions

RQ1 How can we enable on-chain decisions on identity using TLS certificates?

RQ2 How can we use SSL/TLS certificates to endorse Ethereum addresses?

Research Questions

RQ1 How can we enable on-chain decisions on identity using TLS certificates?

RQ2 How can we use SSL/TLS certificates to endorse Ethereum addresses?

Contribution
Design, implementation, and evaluation of a TLS-based authentication framework for Ethereum

- Certificate framework, including Solidity library for parsing and validating TLS certificates
- Endorsement framework

Excerpt of the Requirements for the System

1. Support of on-chain decisions
3. Individual revocation of endorsements
5. Availability
6. Compatibility
9. Cost-efficiency for the verification of endorsements

Outline

1. Motivation, Problem Statement, Research Questions

2. RQ1: How can we enable on-chain decisions on identity using TLS certificates?

3. RQ2: How can we use SSL/TLS certificates to endorse Ethereum addresses?

4. Evaluation of the System

5.Conclusion and Future Work

RQ1: How can we enable on-chain decisions on identity using TLS certificates?

- RQ1.1 What are possibilities to provide determinism for the validity decision?
- RQ1.2 What are the associated costs of the approaches?
- RQ1.3 How can certificates be revoked on-chain?
- RQ1.4 What are inherent problems of the SSL/TLS public key infrastructure and how can we mitigate them?

RQ 1.1: What are possibilities to provide determinism for the validity decision?

Migration Approach

Store and validate certificates on Ethereum, either

decentralized: every endorsement is stored with its complete certificate chain.

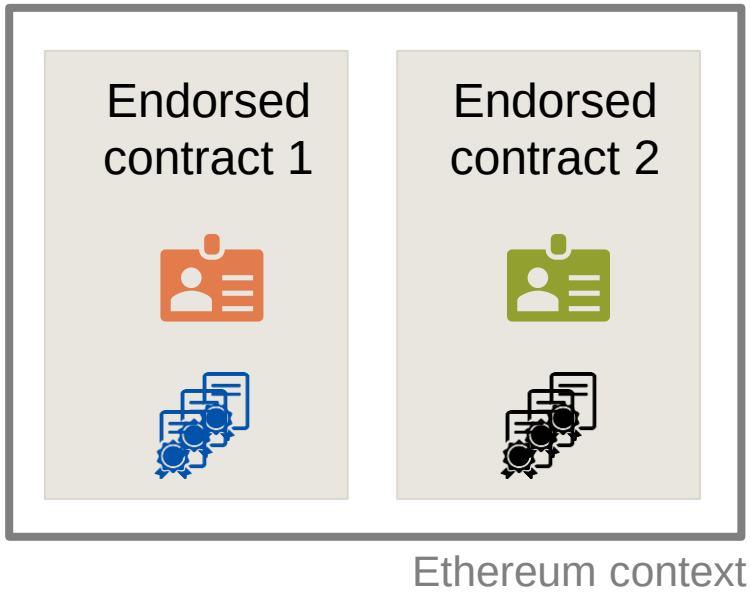

centralized: Certificates are stored in a central database, endorsements are stored separately from certificates.

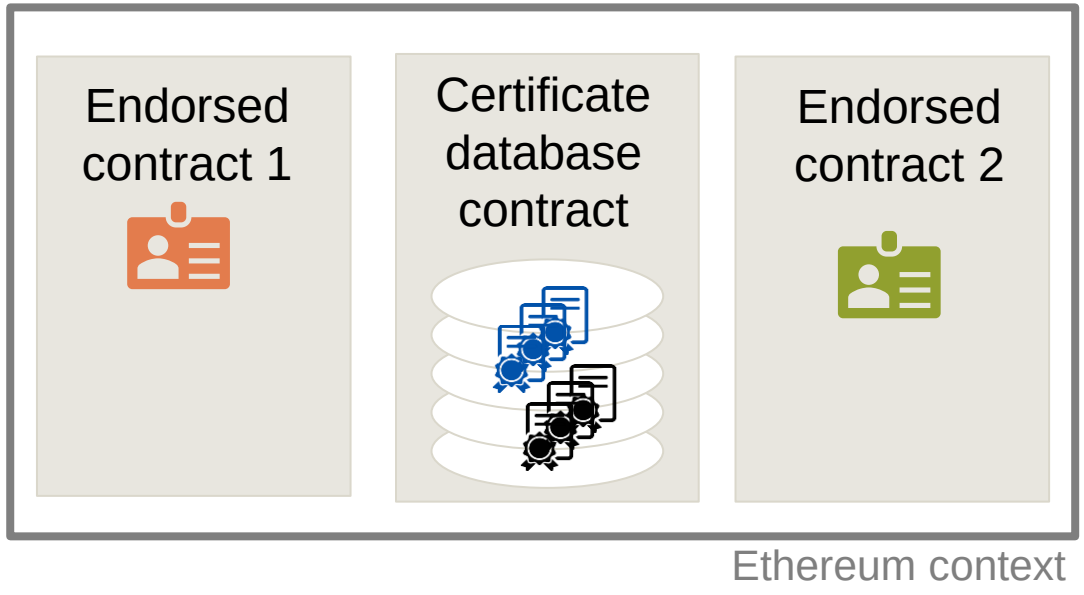

Both approaches allow on-chain decisions (R1) and guarantee availability (R5).

RQ 1.2: What are the associated costs of the approaches?

Survey of the TLS PKI: **200,079,469 domain** certificates (98%) are issued by only **37 certificate authority (CA)** certificates[1]

Decentralized migration

- **Redundant information** and is stored **on-chain**
- Submitting CA certificates for every endorsement incurs **extreme overhead**

Centralized migration

- Centralized approach profits from potential to **reuse already stored CA certificate** information for newly submitted certificates
- Central database **can perform certificate validation** upon submission
- Significantly **lower cost** for all stakeholders expected

[1]Dataset retrieved and analyzed on 21.04.2020 from Censys, https://censys.io/

Outline

1. Motivation, Problem Statement, Research Questions

2. RQ1: How can we enable on-chain decisions on identity using TLS certificates?

3. RQ2: How can we use SSL/TLS certificates to endorse Ethereum addresses?

4. Evaluation of the System

5.Conclusion and Future Work

RQ2: How can we use SSL/TLS certificates to endorse Ethereum addresses on-chain?

- RQ2.1 How can already deployed contracts and externally owned accounts be endorsed?
- RQ2.2 How can identity endorsements be revoked?
- RQ2.3 What measures can an identity owner take to increase trust in their identity claim?

RQ2: How can we use SSL/TLS certificates to endorse Ethereum addresses on-chain?

Endorsement links an Ethereum **address to a domain name**

Endorsement content

- **Address** of the endorsed account
- **Domain name**
- Identifier of the **certificate** used to create the endorsement
- **Signature** of the information above

RQ2.1: How can already deployed contracts and externally owned accounts be endorsed?

Introducing the **endorsement database:**

- A central database can be used to store endorsements and look them up on-chain
- Endorsements in contract can coexist with endorsements in database

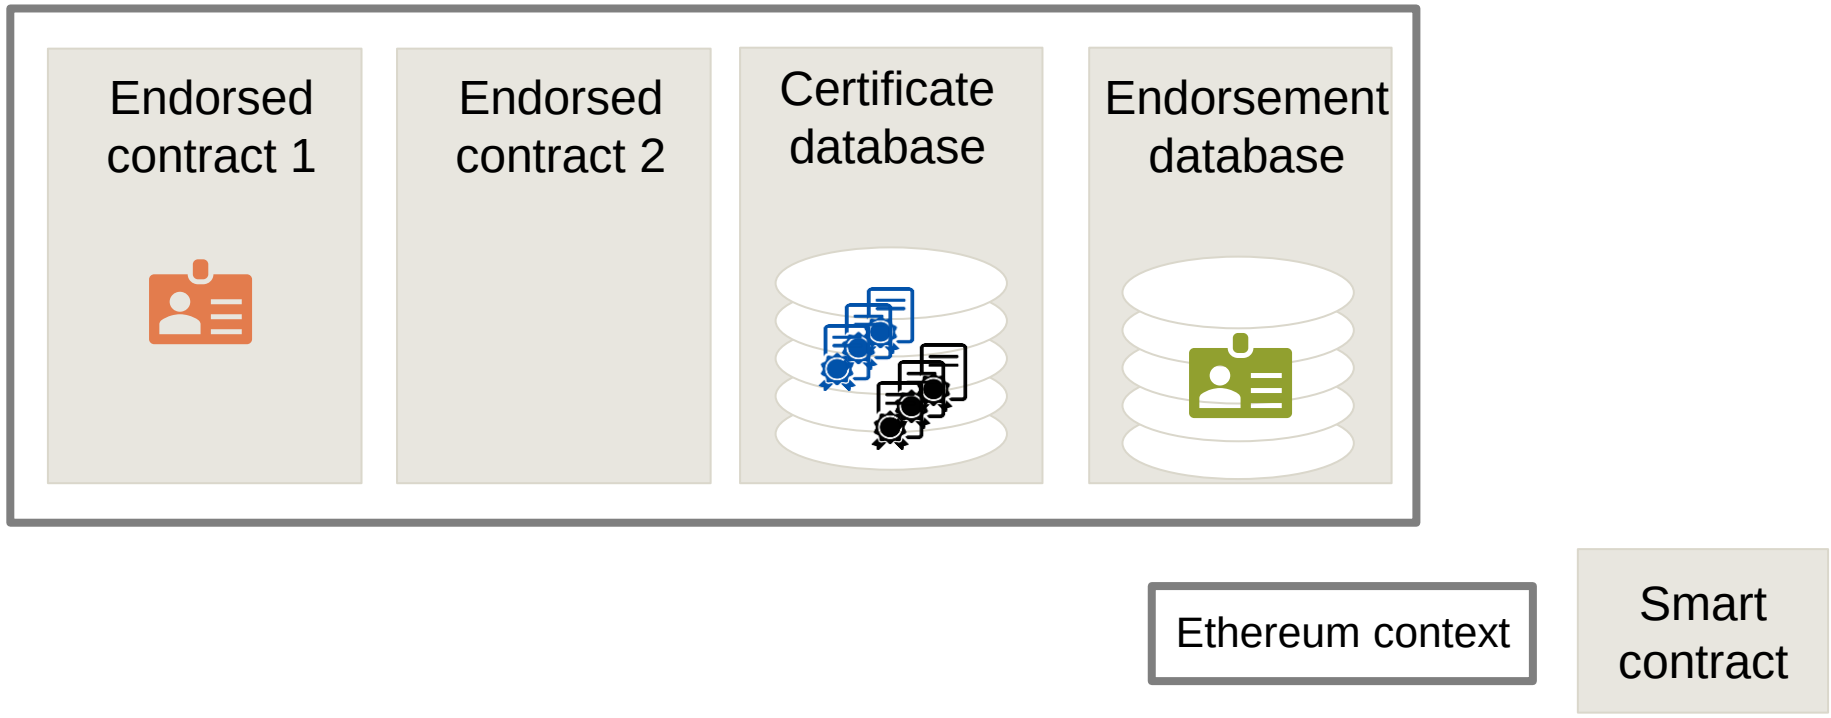

RQ2.2: How can identity endorsements be revoked?

Revocation

- Announces that a specific endorsement cannot be trusted anymore
- Must be possible **without revoking other endorsements or the TLS certificate**

Solution

- Private key owners can create “**revocation messages**”
- Revocation message details account address, domain name, certificate identifier
- Revocation message is **stored in database**

Outline

1. Motivation, Problem Statement, Research Questions

2. RQ1: How can we enable on-chain decisions on identity using TLS certificates?

3. RQ2: How can we use SSL/TLS certificates to endorse Ethereum addresses?

4. Evaluation of the System

5.Conclusion and Future Work

Prototype Performance

Certificates[1]

- **Average cost** per domain certificate:
  1,038,042 gas ≈ **2,35$**[2], average cost sinks for a higher number of submitted domain certificates
- **Mean cost** for adding one domain certificate:
  793,954 gas ≈ **1.81$**[2]

Endorsements

- **Submitting** an endorsement:
  446,900 gas - 577,219 gas ≈ **1.02$ - 1.32$**[2]
- **Retrieving** an endorsement:
  35,000 gas - 446,800 gas ≈ **0.08$ - 1.02$**[2]

[1]Submission of 576 domain certificates from the Top 1000 domains by daily visits, 68 CA certificates required

[2]Conversion rates: Gas fee of 11 Gwei, 206$ per ether, source: https://ethgasstation.info/, accessed 30.04.2020

Security Considerations

Security of the implementation

- Supporting only essential functionality keeps attack surface small

Mapping between domain names and real-world identites

- Internet users are used to domain names
- Threat: Typosquatting

Security of the TLS PKI

- Concept of TLS PKI under criticism of security researchers[1]
- But also a well-studied, closely monitored system[2,3]

[1]Vratonjic, Nevena, et al. "The inconvenient truth about web certificates." *Economics of information security and privacy iii*. Springer, New York, NY, 2013. 79-117.
[2]Holz, Ralph, et al. "The SSL landscape: a thorough analysis of the x. 509 PKI using active and passive measurements." *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. 2011.
[3]Amann, Johanna, et al. "Mission accomplished? HTTPS security after DigiNotar." *Proceedings of the 2017 Internet Measurement Conference*. 2017.

Outline

1. Motivation, Problem Statement, Research Questions

2. RQ1: How can we enable on-chain decisions on identity using TLS certificates?

3. RQ2: How can we use SSL/TLS certificates to endorse Ethereum addresses?

4. Evaluation of the System

5.Conclusion and Future Work

Fulfillment of Requirements

1	Support of on-chain decisions	✓	By storing certificates on-chain
3	Individual revocation of endorsements	✓	Revocation scheme for endorsements
5	Availability	✓	By storing certificates on-chain
6	Compatibility	✓	By introducing endorsement database
9	Cost-efficiency for the verification of endorsements	!	Costs depend strongly on scheme

Conclusion

Great advantage of using TLS certificates for identity assertion and verification on Ethereum: massive amount of trusted data available

On-chain decisions are possible by migrating the necessary parts of the PKI on-chain

Centralizing the validation and storage of certificates as well as of endorsements avoids redundancy and reduces the total costs for all stakeholders

Caveats: not a fully-fledged identity management system as only certificate owners can be authenticated, Ethereum is not cost-optimized for TLS certificates

Future Work

- Extension and improvement of prototype implementation
- Development of a more elaborate endorsement framework
- Investigation of ways to combine a TLS-based authentication framework with an identity systems specifically designed for Ethereum

Friederike Groschupp

friederike.groschupp@tum.de

Technische Universität München
Faculty of Informatics
Chair of Software Engineering for Business Information Systems

Boltzmannstraße 3
85748 Garching bei München

Backup Slides

Background Information: TLS, Certificates, and PKI

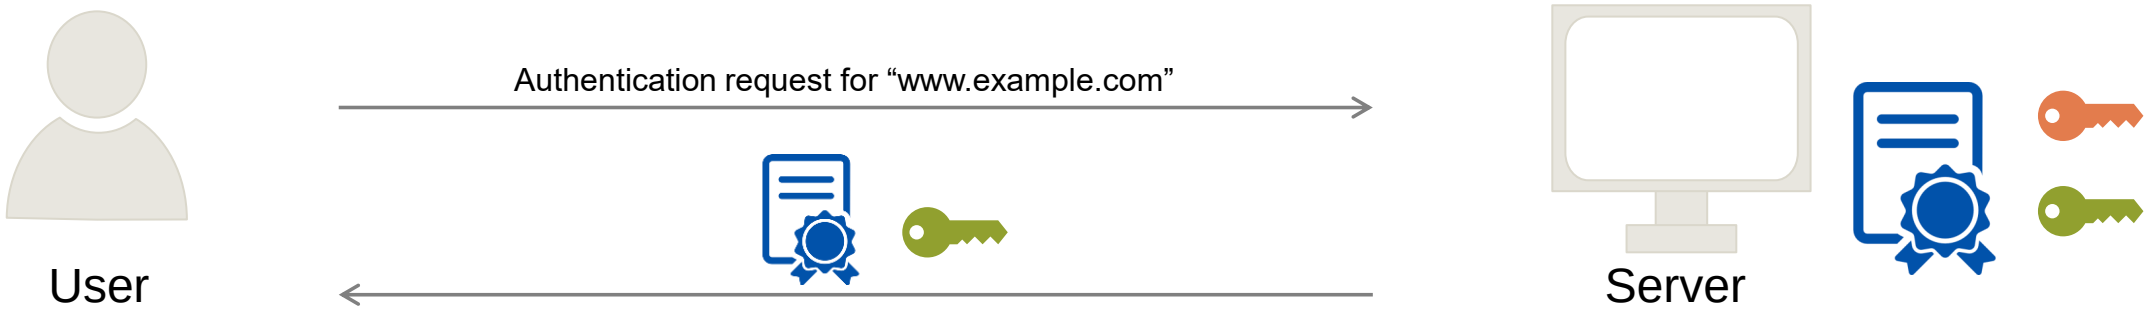

Background Information: TLS, Certificates, and PKI

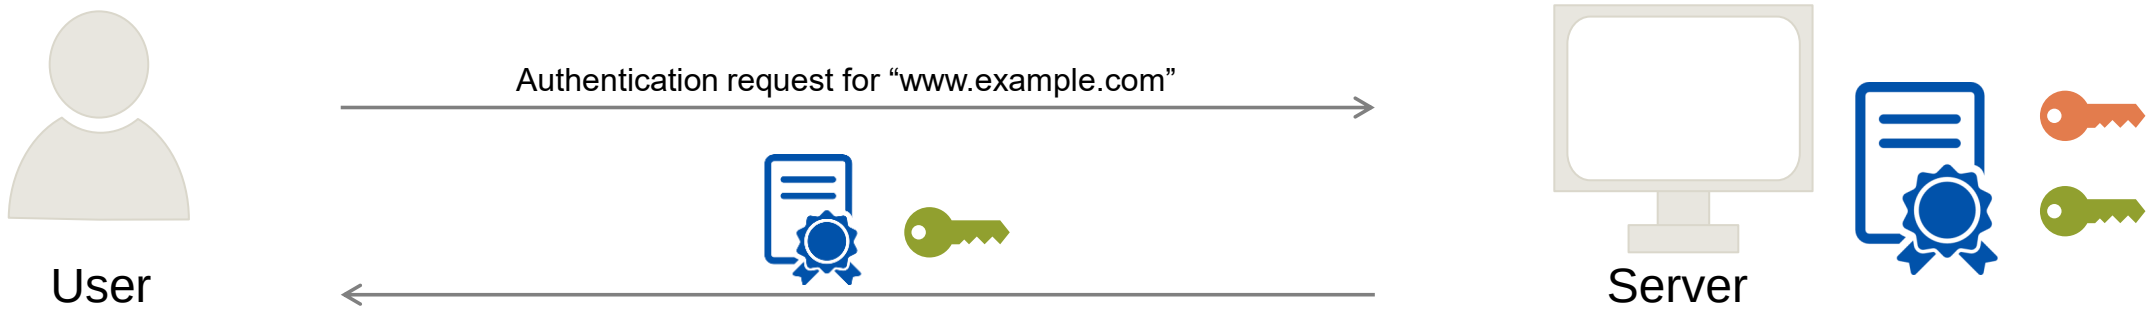

- TLS certificates bind a public key to a domain name
- TLS certificates are issued by certification authorities (CAs)
- Server proves that it “is” the domain by producing a valid signature with the private key
- User decides whether the certificate is valid based on
  - the time of validation
  - the integrity of the certificate and its certificate chain
  - whether they trust the root certificate

Fulfillment of Requirements

R1	Support of on-chain decisions	✓	By storing certificates on-chain
R2	Unambiguous endorsements	✓	Signature of domain and address
R3	Individual revocation of endorsements	✓	Revocation scheme for internal and external endorsements
R4	Open participation	✓	Everybody can determine their set of trusted roots
R5	Availability	✓	By storing certificates on-chain
R6	Compatibility	✓	By introducing endorsement database
R7	Flexible design	✓	No assumptions made on decision policies
R8	Enable independent and fast adoption	✓	No action by other stakeholders required, huge amount of trusted information
R9	Cost-efficiency for the verification of endorsements	~	Cost efficiency for external endorsements

RQ 1.1 What are possibilities to provide determinism for the validity decision?

Oracle Approach

Use an oracle to retrieve certificates from the internet and validate them

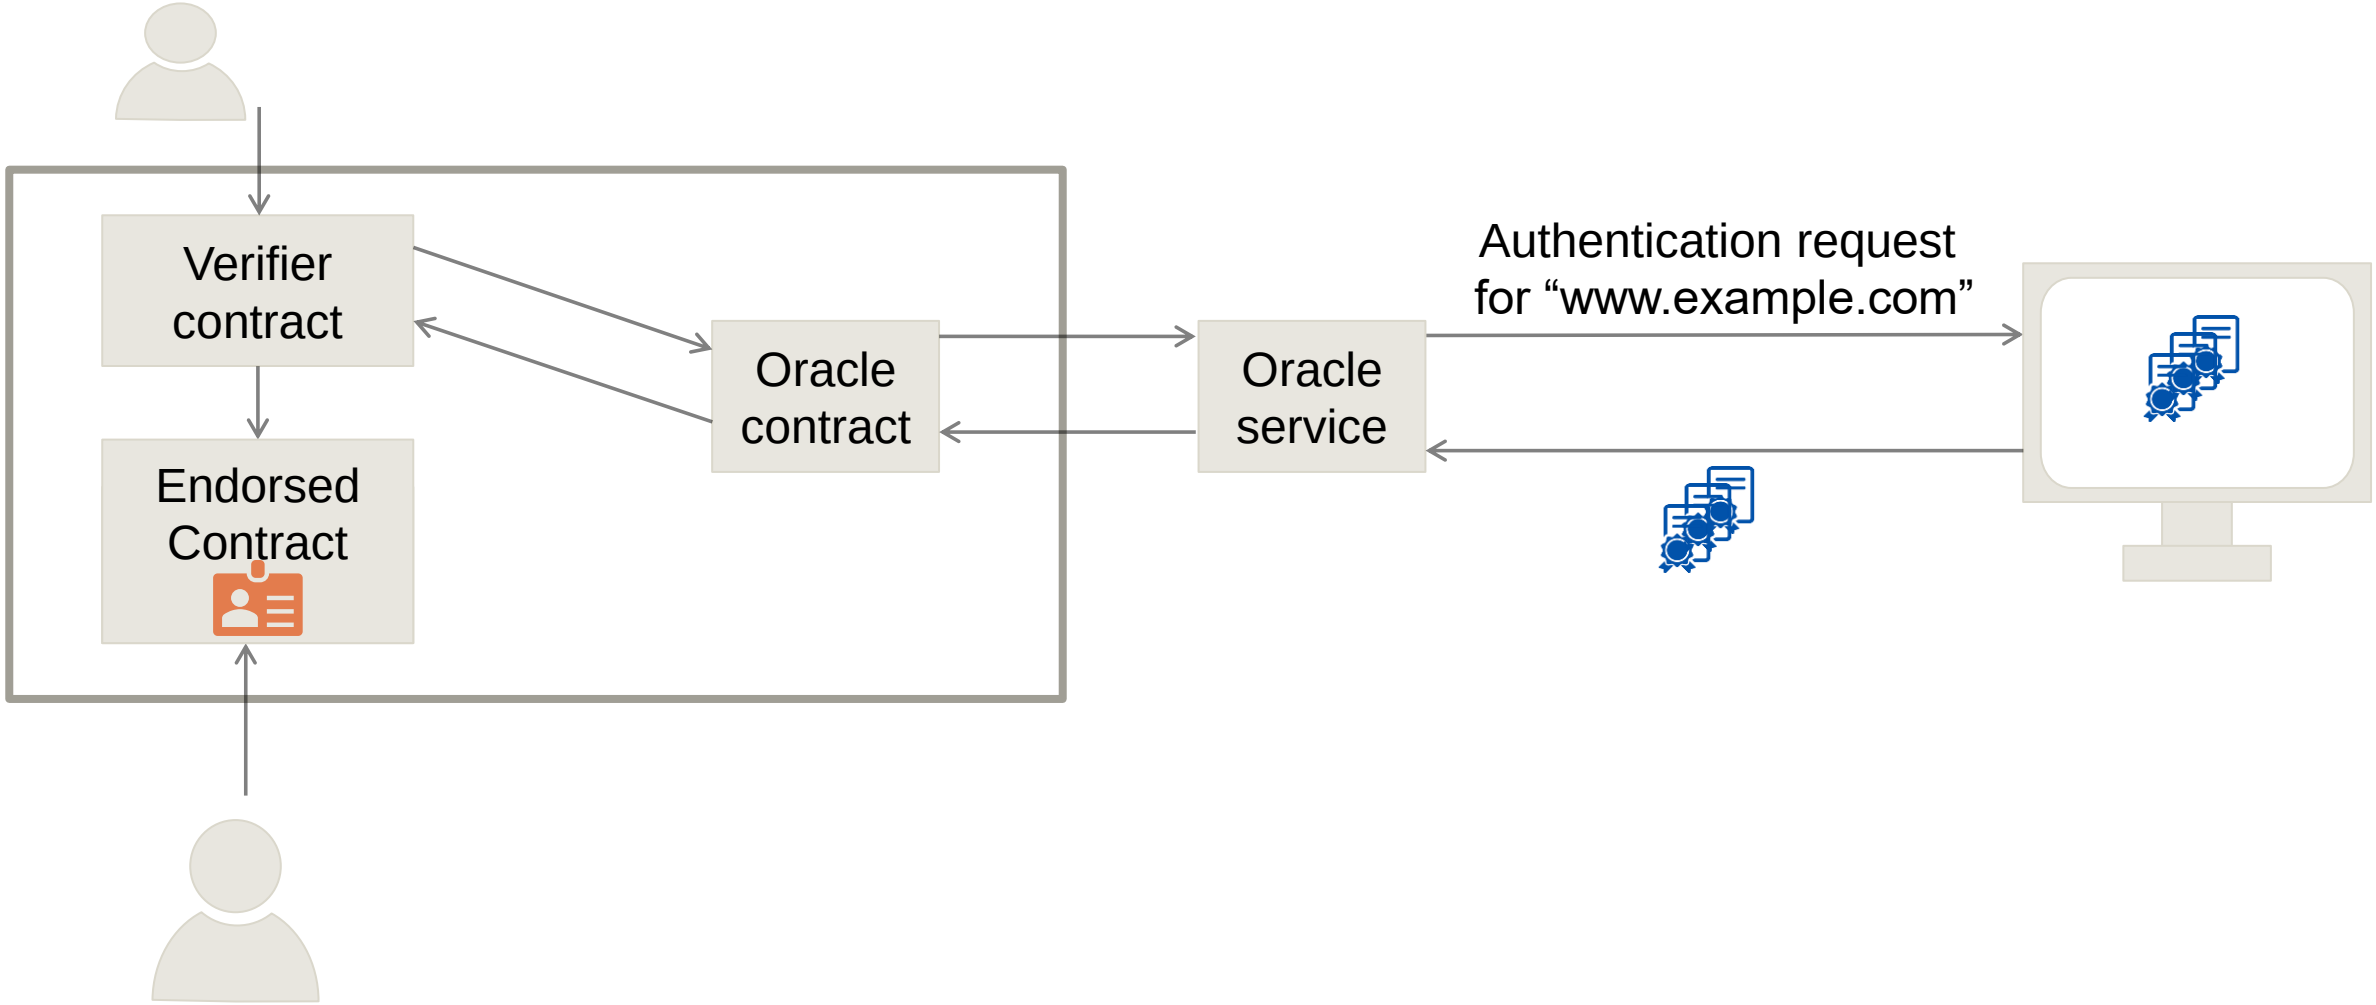

RQ 1.1 What are possibilities to provide determinism for the validity decision?

Oracle Approach

Use an oracle to retrieve certificates from the internet and validate them

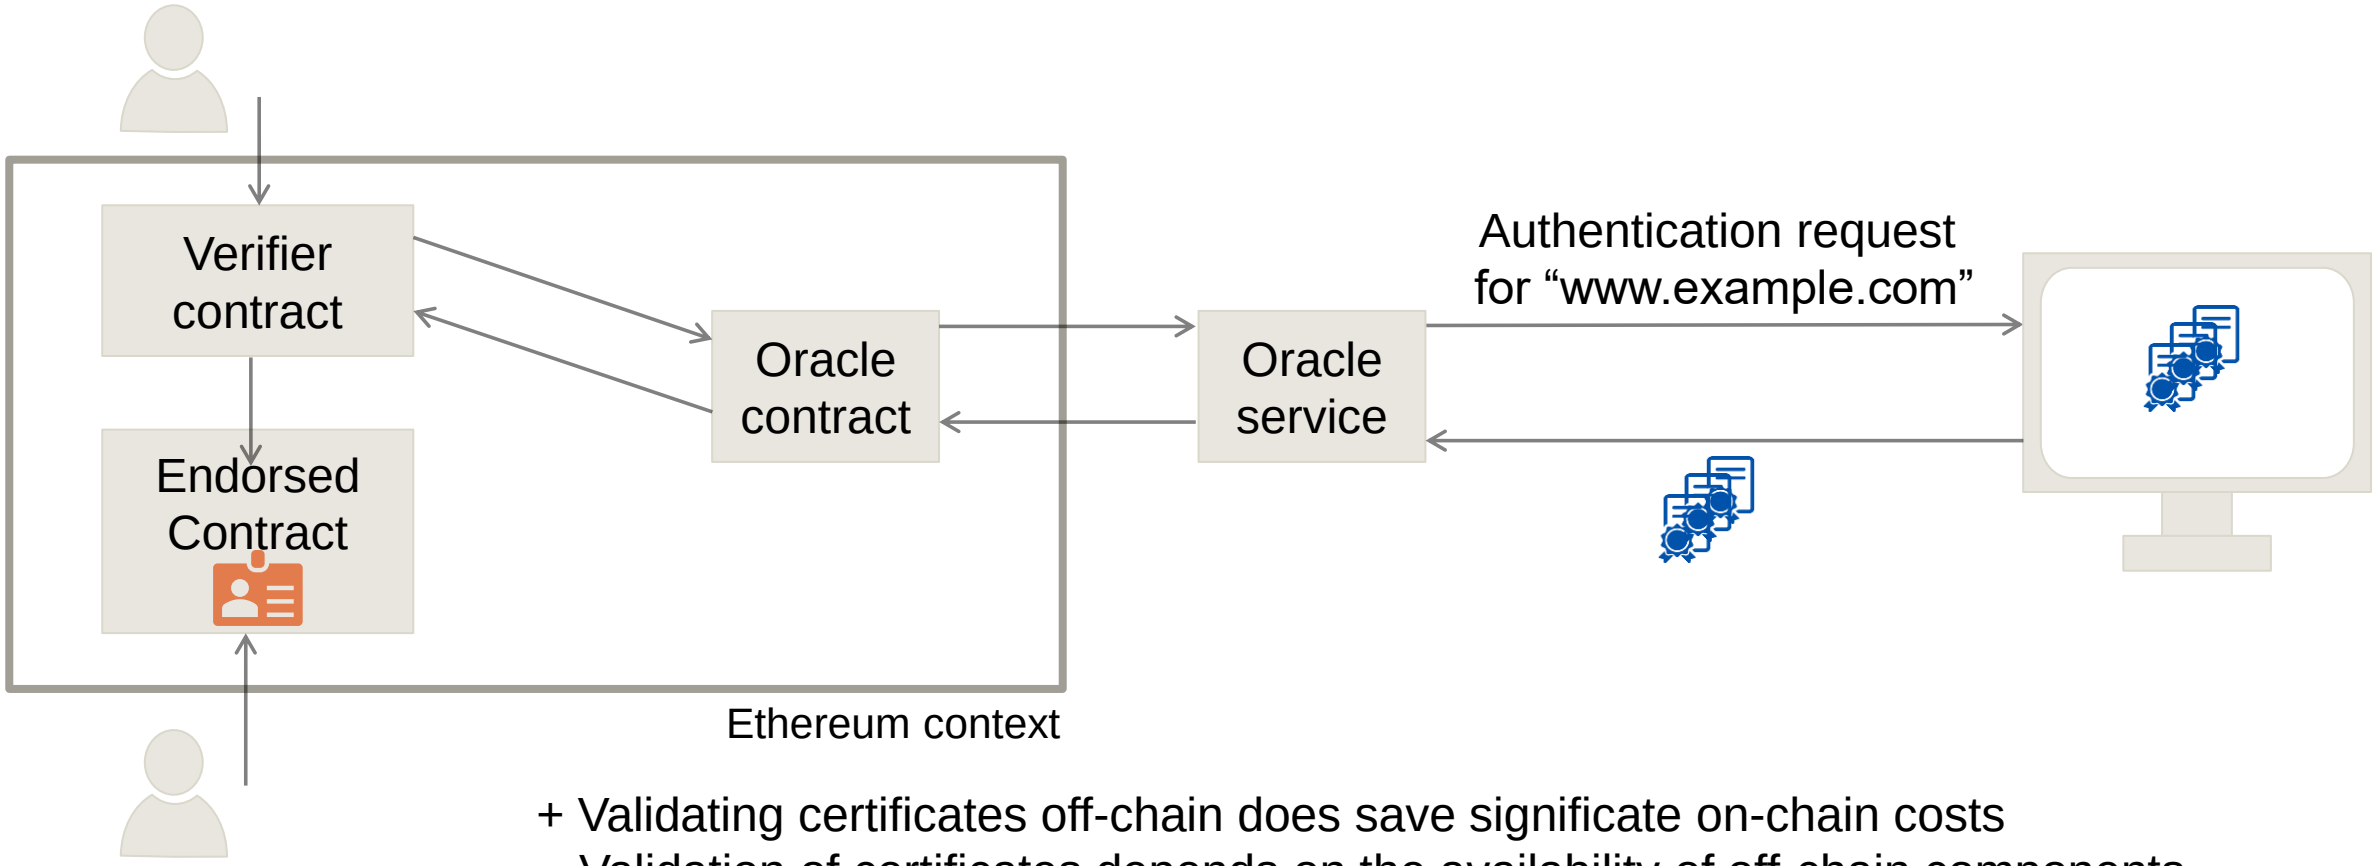

+ Validating certificates off-chain does save significate on-chain costs
- Validation of certificates depends on the availability of off-chain components

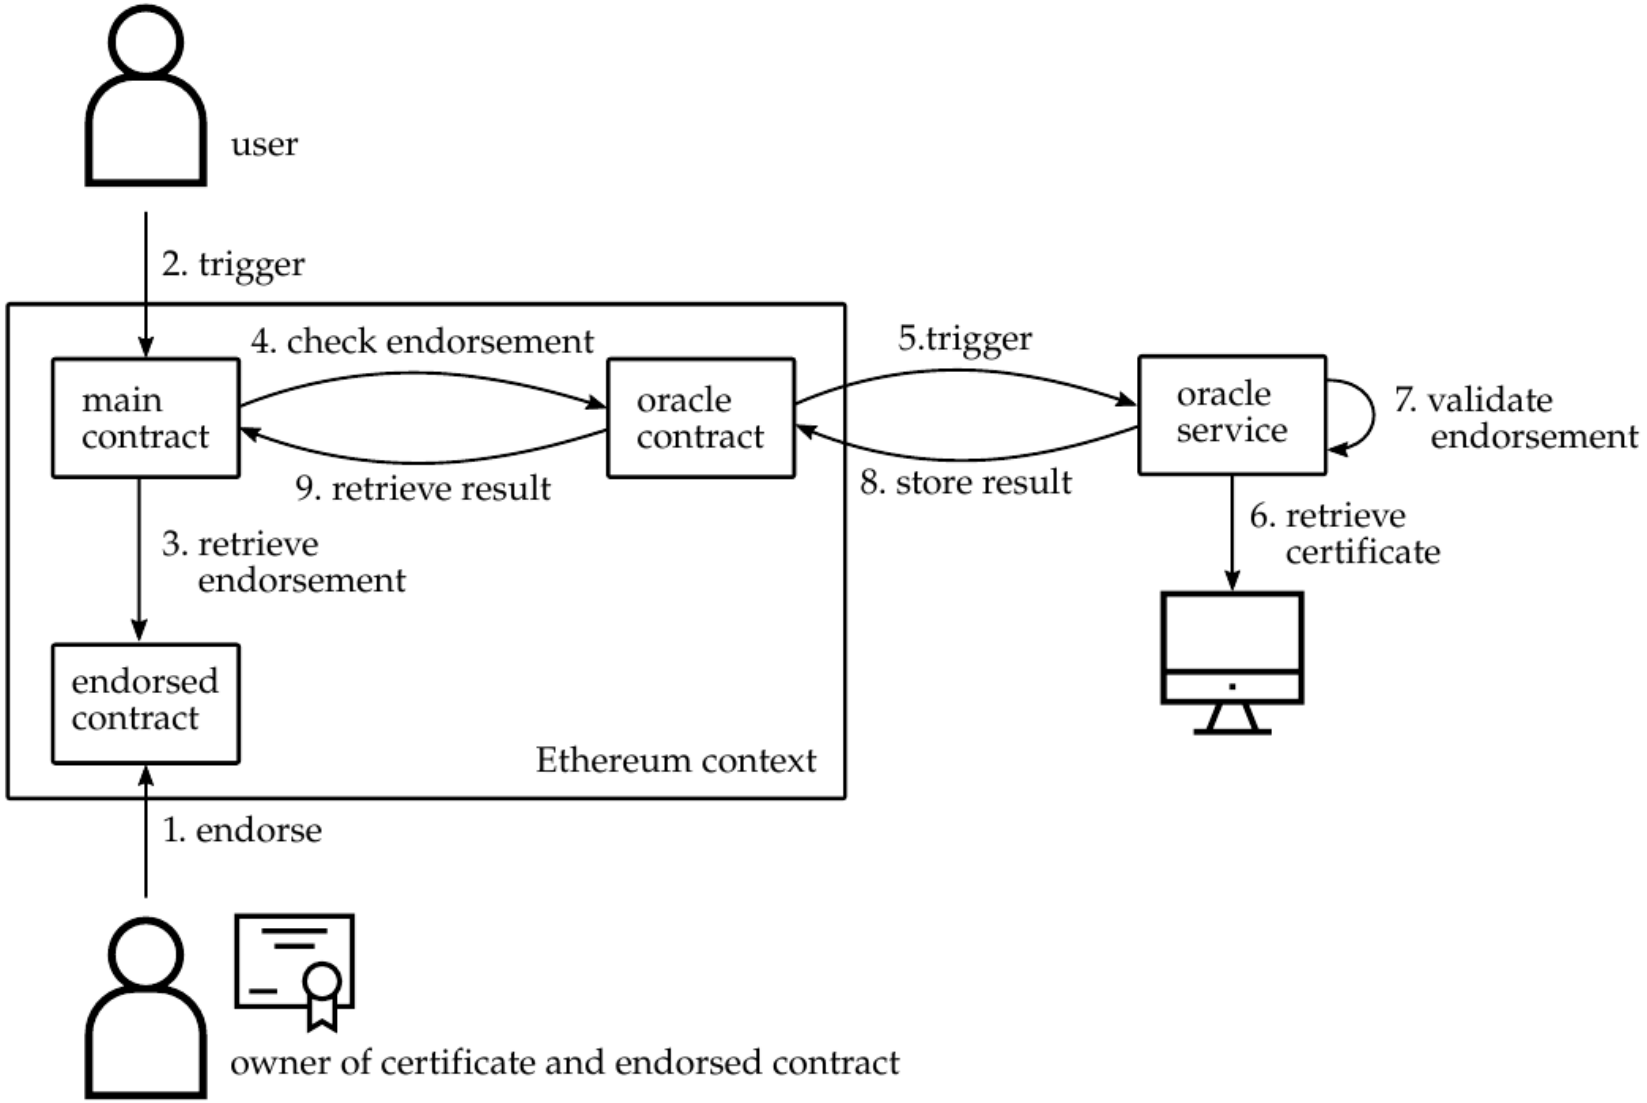

Figure 5.2: Depiction of the workflow when using the oracle approach for certificate validation.

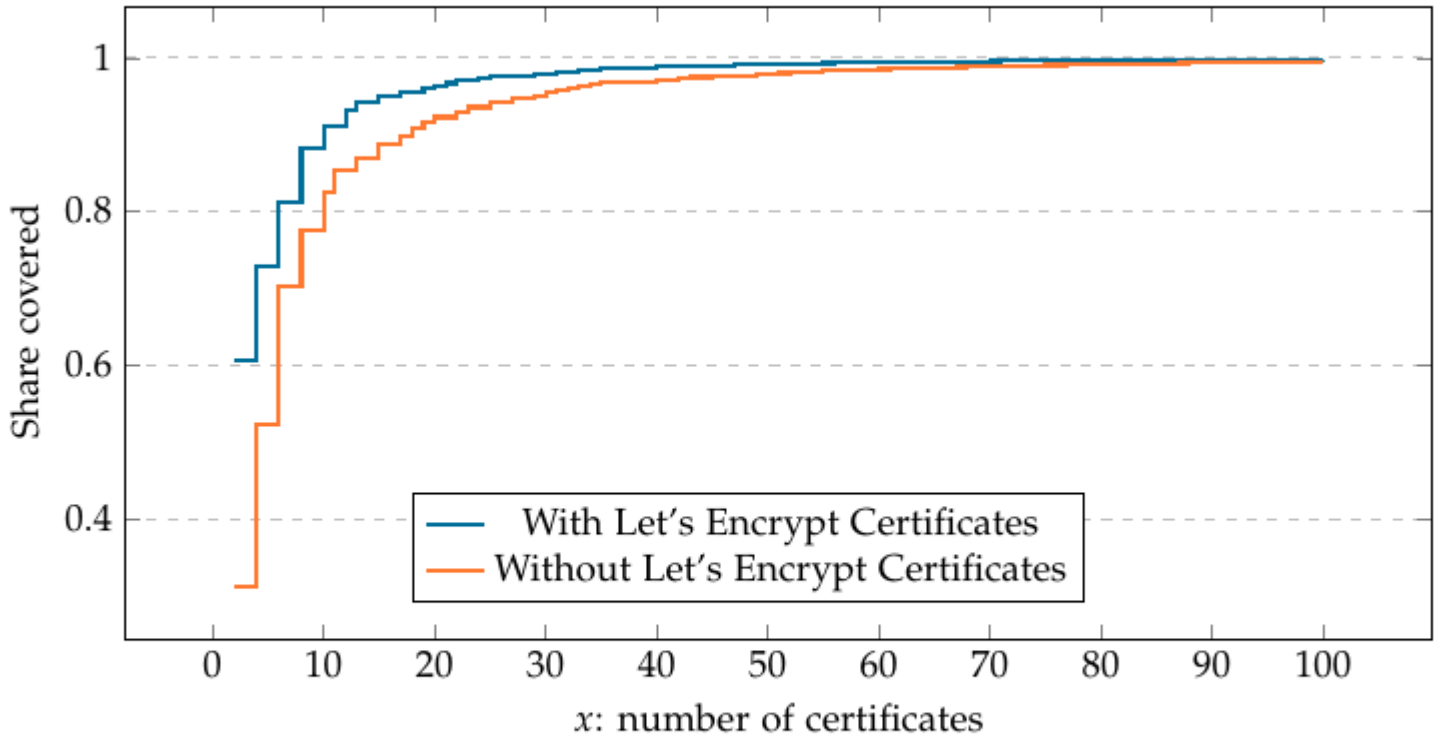

Figure 3.2: Tight lower bound on achievable share of certificate domain when choosing x CA certificates. Graph is truncated after $x = 100$.

RQ1.3 How can certificates be revoked on-chain?

- The revocation status of a certificate is checked by the verifier by requesting a
  - Certificate Revocation List (CRL) or
  - Online Certificate Status Protocol (OCSP) response

- These are documents cryptographically signed by the certificate issuer either attesting the
  - Continued validity of a certificate
  - Revocation of certificate

- CRLs and OCSP responses can be validated on-chain to revoke or “refresh” a certificate

RQ1.4 What are inherent problems of the SSL/TLS public key infrastructure and how can we mitigate them?

- One **trusted rogue certification authority** (CA) can **undermine the security** of the whole system
  - A central certificate database on a blockchain hinders split-world attacks
  - CAs cannot issue fraudulent certificates unnoticed
  - Trust in rogue CA can be revoked

- TLS ecosystem is very diverse, TLS certificate parser had **significant security-relevant bugs** in the past
  - Support only the most essential functionality
  - Keep the attack surface small

RQ2.2: How can identity endorsements be revoked?

External Endorsements:

- Owner of the private key  can create and sign a revocation for any endorsement created with 
- Revocation can be **submitted to and validated by the endorsement database**
- Verifiers can retrieve the revocation status together with the endorsement

Internal Endorsements: What if initial owner loses control over an endorsed contract?

- Only the owner of a contract can store an endorsement in it
- Anybody can submit a revocation to a contract, but only, if they can **produce a valid revocation** message signed with the private key  of a **certificate that was previously used to endorse** the contract
- Enables previous owners to revoke an endorsement, but **prevents spamming** of contract with revocations

RQ2.3: What measures can an identity owner take to increase trust in their identity claim?

The credibility and **trustworthiness of an endorsement depends on the certificate** that was used to create it.

Identity owner can take several measures to increase trust in their endorsements:

- Buy certificates from **well-trusted CAs** with high security standards
- Buy **higher-grade certificates** such as “Extended Validation” Certificates
- **Refresh the validity status** of their certificate by updating it with respective CRLs and OCSP responses

Prototype Performance

Submission of 576 certificates from the Top 1000 domains by daily visits

Average costs per certificate:

- 1,038,042 gas / 2,35 $ average cost per domain certificate, average **cost sinks for a higher number** of submitted domain certificates
- Centralized approach: 2,762,042 gas / 6.25 $ average cost, does **not decrease** with more domain certificates

Cost of submitting endorsement:

- Internal: 446,900 gas / 1.02 $
- External: 577,219 gas / 1.32 $

Retrieving (and validating) an endorsement:

- Internal: 446,800 gas / 1.02 $
- External: ~ 35,000 gas / 0.08 $

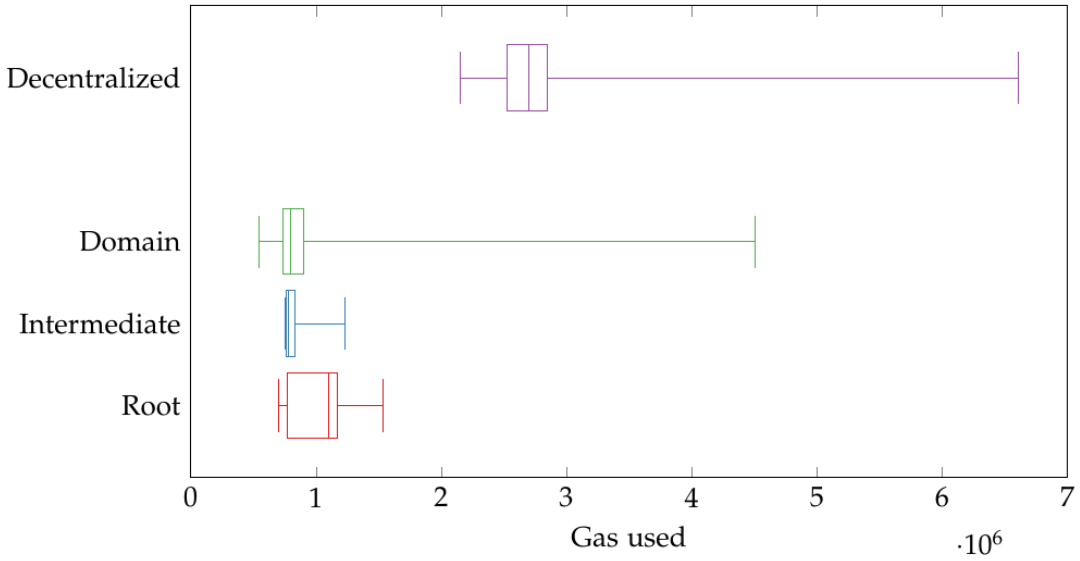

Amount of gas used for submission of root, intermediate, and domain certificates. Decentralized approach for comparison.

	Root certificate			Intermediate certificate			Domain certificate		
	gas	ether	$	gas	ether	$	gas	ether	$
min	705,035	0.0078	1.60	750,584	0.0083	1.70	544,777	0.0060	1.23
1st	770,455	0.0086	1.77	762,129	0.0085	1.75	733,073	0.0081	1.66
med	1,105,114	0.0123	2.53	783,324	0.0087	1.79	793,954	0.0088	1.81
3rd	1,170,981	0.0130	2.67	832,031	0.0092	1.89	903,813	0.0100	2.06
max	1,537,513	0.0171	3.52	1,233,724	0.0137	2.82	4,503,213	0.0500	10.3

Table 7.1: Cost of certificate submission in gas usage, ether, and US dollar.