S-EL2

Secure Systems Group, Aalto University

Ayoub Chouak, Lachlan J. Gunn

Trustworthy Data Provenance for Enclaves in Heterogeneous Distributed Systems

- Confidentiality and integrity of data is at risk as mobile enclaves migrate through many systems
- Trustworthy data provenance aids in detecting confidentiality and integrity breaches of data
- Tamper-evident, cryptographically-protected logs provide trustworthy provenance

System model

- Trusted Execution Environment (TEE)-based enclave platform
- Append-only logs hold provenance data

- **Challenges: Compromised TEE**, provenance **forging**, **truncation** and/or **replay**

Objectives

- Tamper-evidence: alterations must be detectable
- Provenance attestation: provenance data must be convincing to observers
- Simplicity: no expensive network algorithms (e.g., consensus)

Solution

- WebAssembly-based enclave platform: enclaves sandboxed from each other, and TEE/runtime
- wasmi WebAssembly interpreter running under TrustZone using OP-TEE trusted OS
- Tamper-evident logs using Hypercore library stored in trusted storage
- Logs IO operations and data migrations

Conclusion

- Simple and highly scalable design
- Availability: Should the logs be replicated?

- **Applications**: information flow tracking, detection of access control policy violations

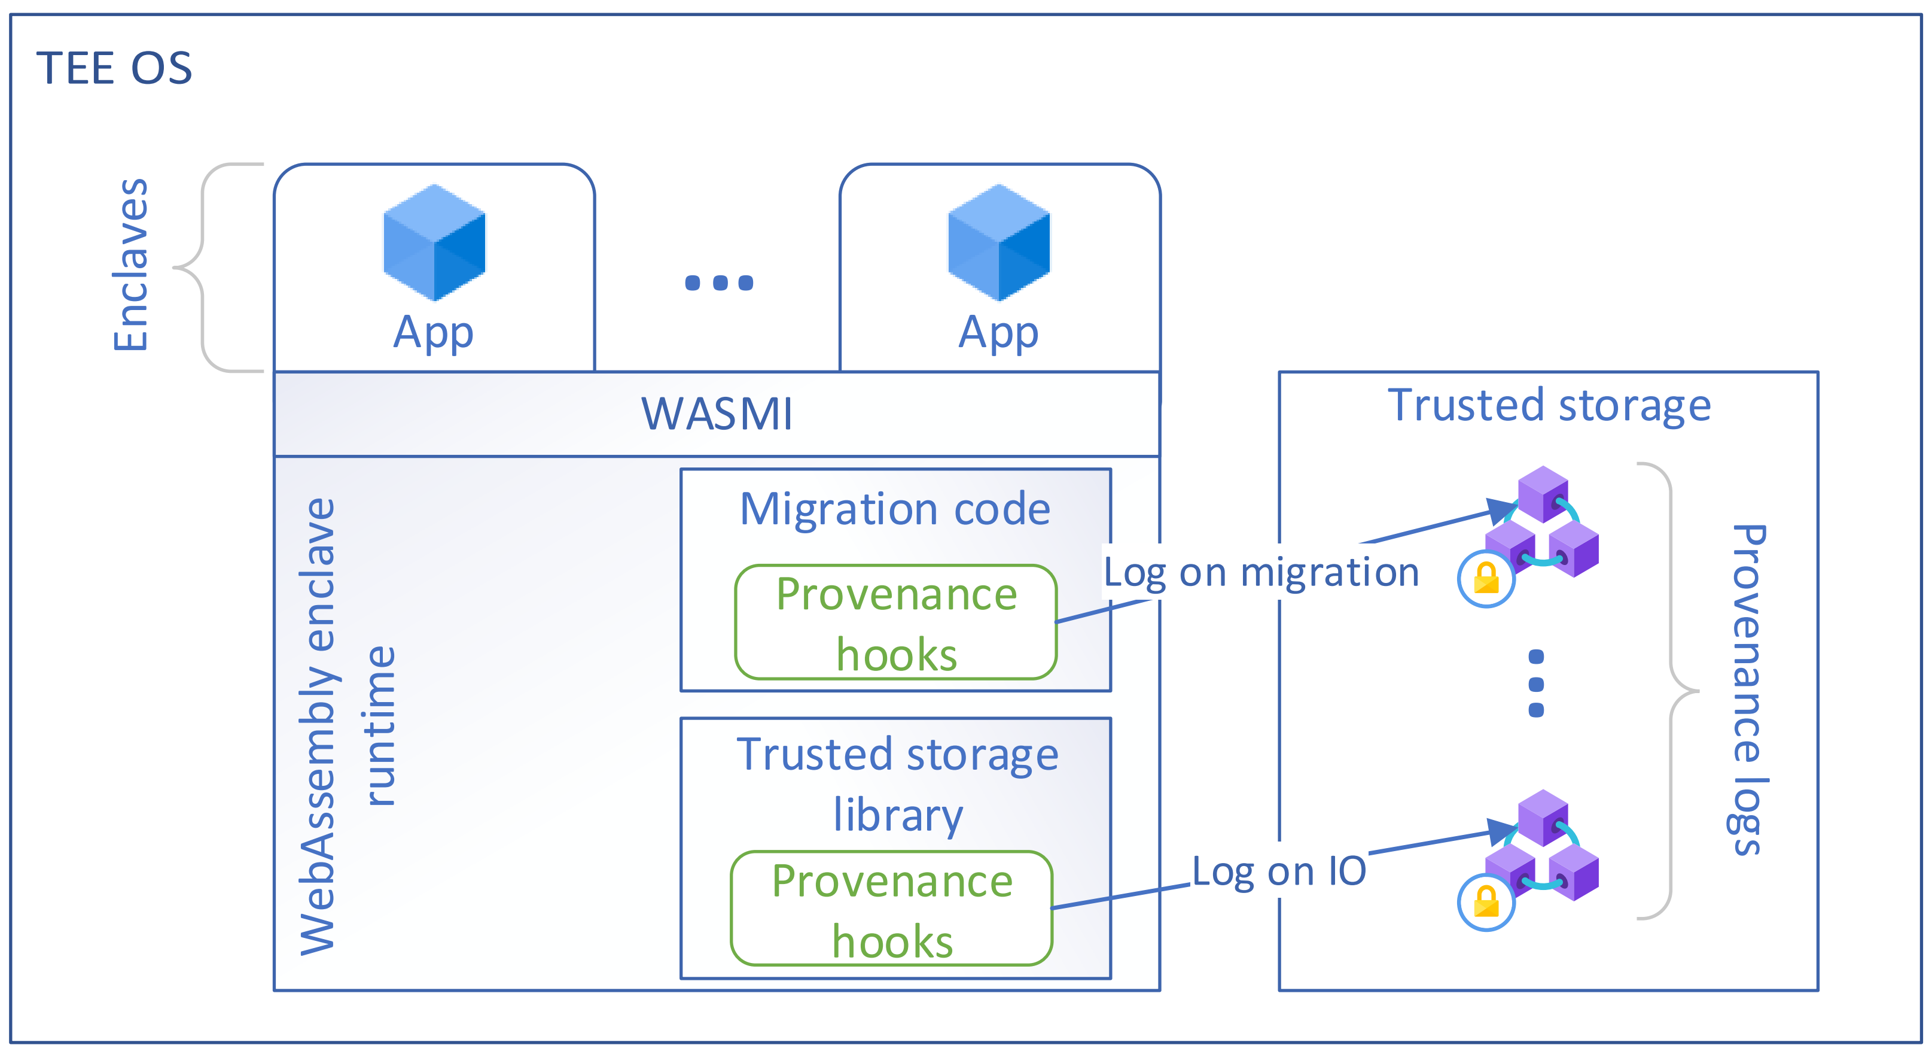

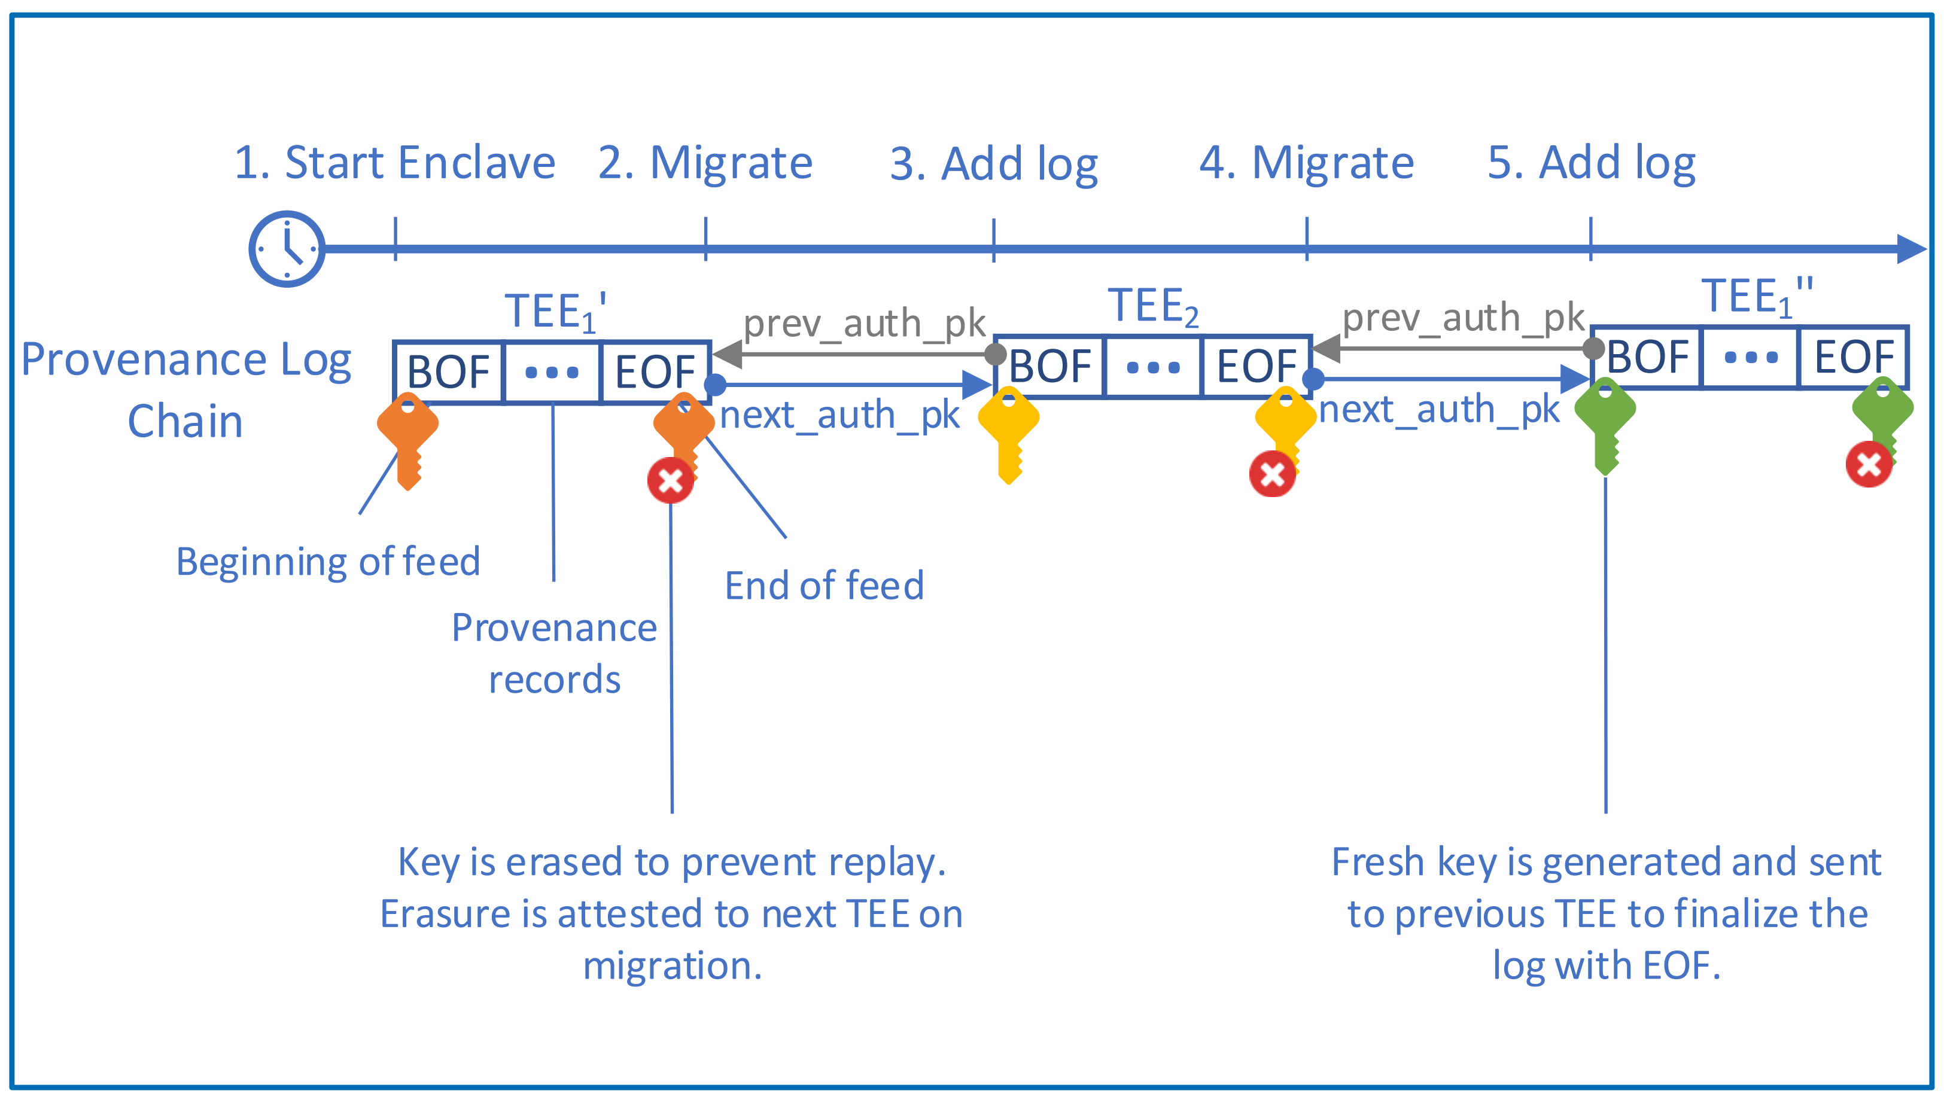

Key is erased to prevent replay. Erasure is attested to next TEE on migration.

Fresh key is generated and sent to previous TEE to finalize the log with EOF.

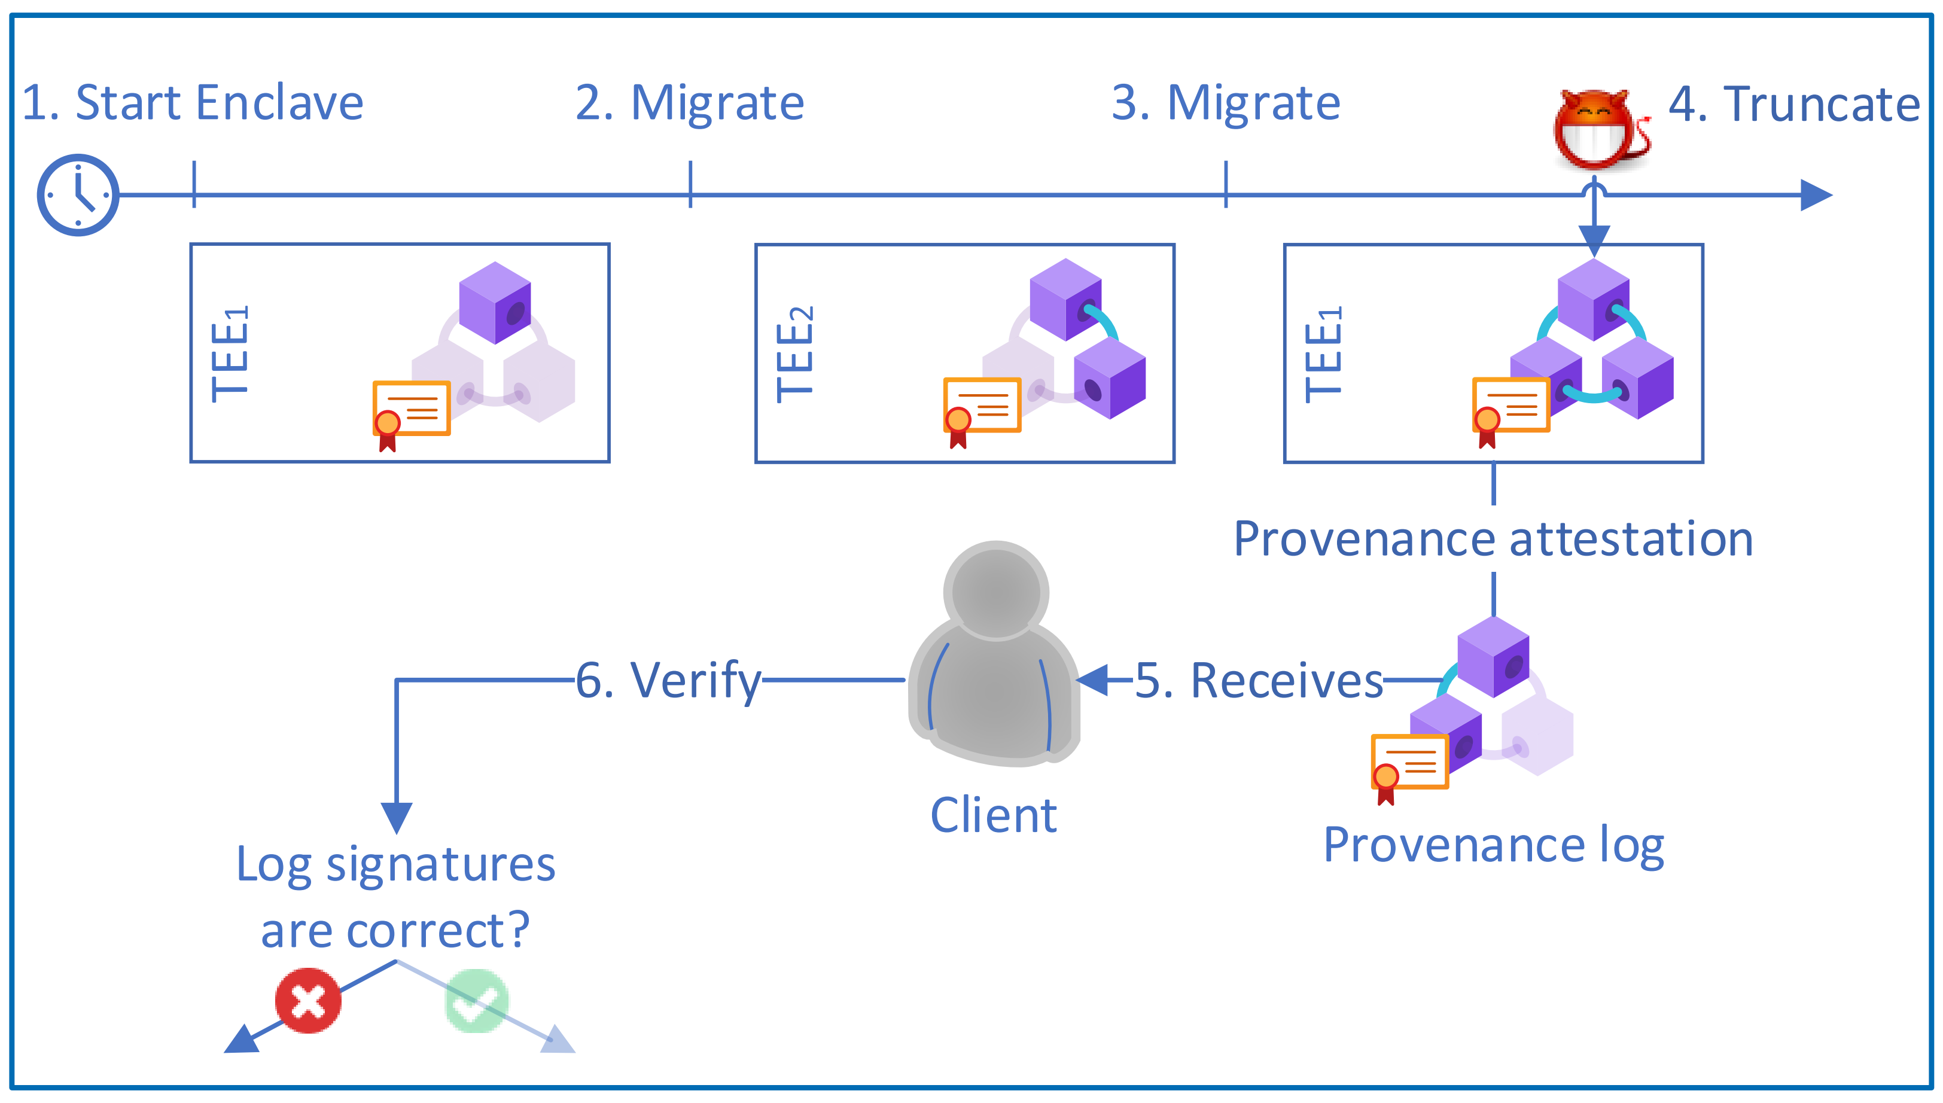

Acknowledgements

This work was funded by Huawei Technologies Oy.