EMPIRICAL EVALUATION OF INTEGRITY ATTESTATION PROCEDURES IN PUBLIC CLOUD ENVIRONMENT

Shubh Gupta - Supervised by Dr. Ryan Ko and Omar Jarkas

REIT4841: Research and Development Methods and Practice

Abstract

Goals:
- Evaluate the resilience and industry readiness of the Keylime remote attestation framework under simulated attack scenarios

Background Info:
- Remote Attestation (RA) is crucial in cloud computing for verifying the integrity of virtual machines and protecting against tampering in distributed environments. Keylime is an open-source framework that enables continuous, dynamic attestation using TPM-based trust to ensure system integrity.

Outcomes:
- A Man-in-the-Middle attack successfully intercepted initial attestation data by exploiting Keylime's database trust model. However, Keylime's continuous attestation process detected the intrusion, demonstrating its strong resilience and robust security design.

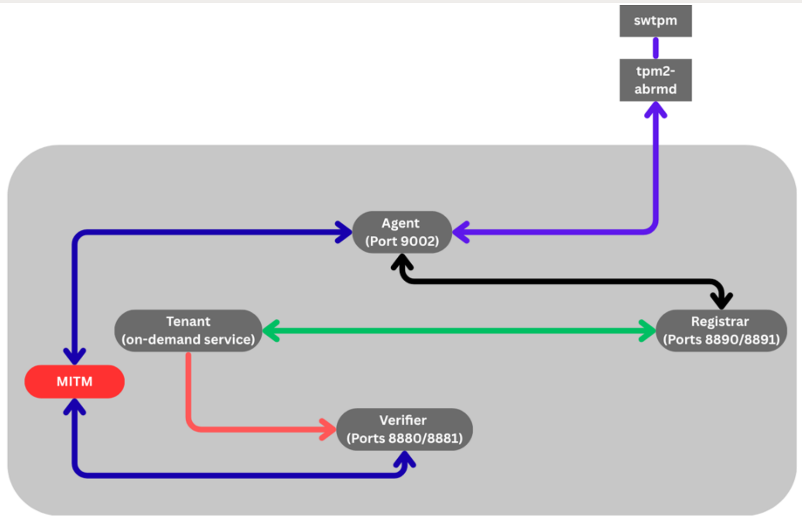

Figure 1: Keylime MITM Attack Architecture

Attack Methodology

1. Start a clean state
   - Stop/remove any old containers and start the services
2. Establish baseline
   - Add agent to verifier once to confirm normal operation
3. Poison the state of truth
   - Poison the registrar DB so agent contact points to mitm-verifier:9002
   - Poison verifier's sqlite cache for the same agent
4. Ensure MITM is in place
   - mitm-verifier listens on port 9002 and forwards to keylime-agent:9002 with TLS
5. Trigger attestation using poisoned database
   - Expect Keylime tenant operations to be redirected through the MITM proxy due to database poisoning, allowing interception of initial attestation quotes. Background attestation may fail if TLS certificate validation is enforced

shubh.gupta@student.uq.edu.au
linkedin.com/in/shubhgupta2510
uq.edu.au

School of Electrical Engineering and Computer Science

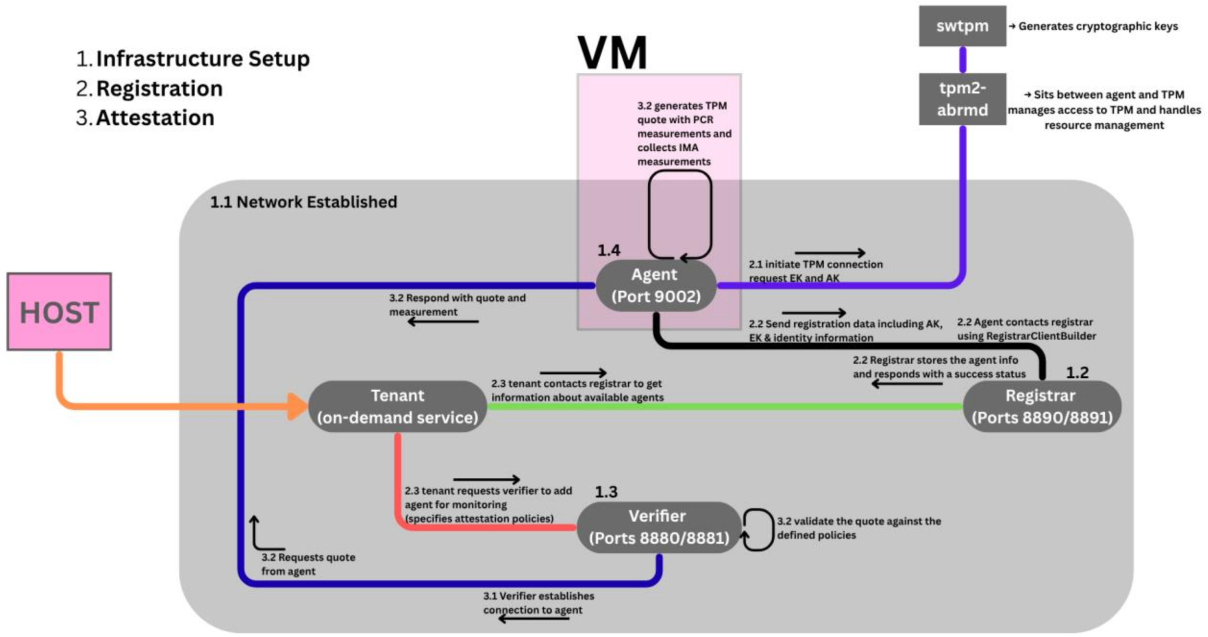

Figure 2: Keylime Architecture

Figure 3: Keylime Logo

Threat Model and Security and Threat Analysis

Simulated MITM attack intercepting verifier–agent traffic to test Keylime's real-world attestation resilience.

Threat Model:
- Defined a threat model simulating a MITM adversary to evaluate Keylime's ability to preserve confidentiality, integrity, and authenticity of attestation exchanges in untrusted network environments.

Security and Threat Analysis:
- Analysed Keylime's resilience against MITM attacks through its mTLS, TPM quote validation, and nonce-based attestation mechanisms.
- Identified residual weaknesses, including PKI dependence, mTLS misconfiguration risks, and attestation timing gaps.
- Proposed extensions such as continuous attestation, cryptographic IMA binding, and verifier isolation for enhanced assurance.


```
2025-10-15 09:36:11.967 - keylime.tenant - INFO - Agent Info from Verifier (keylime-verifier:8881):
{"d432fbb3-d2f1-4a97-9ef7-75bd81c00000": {"operational_state": "Get Quote", "v": null, "ip": "mitm-verifier", "port": 9002, "tpm_policy": "{\"mask\": \"0x0\"}", "meta_data": "{}", "has_mb_refstate": 0, "has_runtime_policy": 0, "accept_tpm_hash_algs": ["sha512", "sha384", "sha256"], "accept_tpm_encryption_algs": ["ecc", "rsa"], "accept_tpm_signing_algs": ["ecschnorr", "rsassa"], "hash_alg": "", "enc_alg": "", "sign_alg": "", "verifier_id": "default", "verifier_ip": "keylime-verifier", "verifier_port": 8881, "severity_level": null, "last_event_id": null, "attestation_count": 0, "last_received_quote": 0, "last_successful_attestation": 0}}
```


Figure 4: Keylime remote attestation logs with the Man-in-the-Middle attack

Acknowledgements

I would like to sincerely thank Dr. Ryan Ko and Omar Jarkas for supervising this research and providing guidance throughout the project. Their insights, feedback, and encouragement were instrumental in shaping both the direction and the execution of this work. I would especially like to acknowledge the thoughtful discussions and technical advice they provided, which helped me navigate challenges and refine my approach. I am also grateful to the technical and administrative staff at the School of Electrical Engineering and Computer Science, University of Queensland, for their support and for providing access to the necessary resources. Finally, I would like to thank my family and friends for their patience, encouragement, and understanding throughout this project. Their support made the completion of this thesis possible.

Results

The attack exploited Keylime's database trust model to redirect verifier traffic and intercept sensitive attestation data.

The attack successfully compromised initial attestation by modifying database entries to point to mitm-verifier:9002, allowing interception of TPM quotes and public keys. While continuous attestation failed due to TLS validation, the initial compromise demonstrates a critical vulnerability in Keylime's database trust model.

Conclusion

- Keylime demonstrated robust cryptographic resilience under most network-level threats, successfully detecting active MITM interference during continuous attestation.
- However, the attack revealed a weakness in Keylime's database trust model, allowing partial compromise of initial attestation through verifier redirection.
- These findings highlight the importance of securing backend trust anchors, not just communication layers, in remote attestation frameworks.
- Future work will focus on strengthening database integrity, implementing continuous attestation, and improving PKI trust validation to achieve full end-to-end assurance.

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

CREATE CHANGE