

# Formal Verification of the Kali Invariant — ProVerif Models and Security Properties

*Cryptographic protocol verification of ASMP/1.0 using ProVerif: model construction, security queries, and results for three protocol components*

Author: Arul Raj · Patent IN202641070690 · June 2026

Primary Audience: Security Researchers · Standards Bodies · Protocol Implementers · Academic Reviewers · Classification: Public Technical Paper

## ABSTRACT

---

Formal verification provides a class of assurance that testing and red-teaming cannot: a mathematical proof that a security property holds for all possible executions of a protocol, not just the executions a tester thought to try. This paper presents three ProVerif models for the Adaptive Surface Mutation Protocol (ASMP/1.0) and the security properties verified against each.

Three components were modelled: the zero-knowledge peer authentication handshake (MSG-002), the cascade authentication mechanism (MSG-004 Defensive Leap), and the TEE-rooted management attestation (MSG-005). All three models were verified. The ZK authentication property is TRUE. The cascade authentication property is TRUE. The MSG-005 TEE-binding properties are CORRECT across all queries.

## 1. WHY FORMAL VERIFICATION

---

### 1.1 The limits of testing

A test suite — however comprehensive — can only test the executions it covers. The PME test suite has 683 tests and six red-team scenarios. These provide high confidence that the implementation behaves correctly in the scenarios the tests cover. They cannot prove the protocol is secure against all possible adversary strategies.

An adversary is not constrained to the scenarios a test suite covers. A protocol with a logical flaw — an ordering vulnerability, a replay weakness, a type confusion — may pass all tests in normal operation and fail only under a specific adversary message sequence that no tester anticipated.

### 1.2 What formal verification provides

ProVerif models the protocol as a set of Horn clauses and uses resolution to determine whether a stated security property can be violated by any adversary in the Dolev-Yao model. The Dolev-Yao adversary controls the entire network: it can read, block, replay, modify, and inject any message. If ProVerif cannot find a violation, the property is proven to hold for all possible adversary strategies over all possible protocol executions.

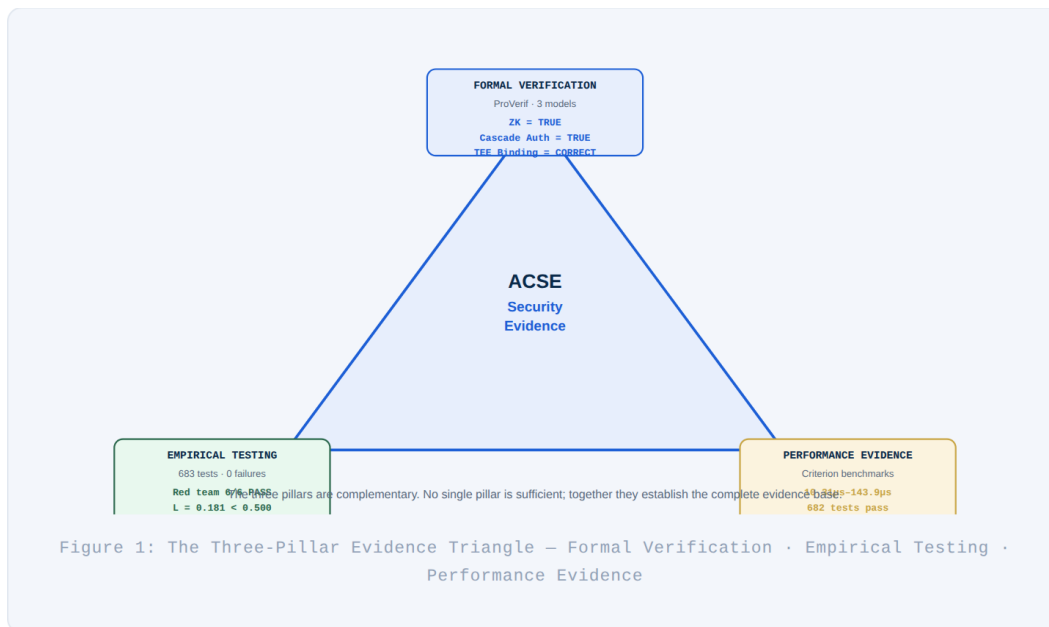


Figure 1: Three-Pillar Evidence Triangle — Formal Verification · Empirical Testing · Performance Evidence

### SCOPE OF FORMAL VERIFICATION

The ProVerif models verify the PROTOCOL — the abstract description of message exchanges, cryptographic operations, and ordering constraints. They do not verify the Rust implementation of the protocol. Implementation correctness is addressed by the 683-test suite, static analysis, and the zero-unsafe-block engineering constraint. Formal verification and implementation testing are complementary. Formal verification eliminates logical protocol vulnerabilities. Testing eliminates implementation bugs. Both are necessary; neither is sufficient alone.

## 2. PROVERIF BACKGROUND

### 2.1 Methodology

ProVerif is an automatic cryptographic protocol verifier developed at INRIA. It uses an internal model of the cryptographic primitives (perfect encryption, perfect signatures, perfect hash functions under the random oracle model) and the Dolev-Yao network adversary. Given a protocol description and a security query, ProVerif either:

- **Returns TRUE:** the property holds for all executions — the adversary cannot violate it
- **Returns FALSE with an attack trace:** the property can be violated — ProVerif provides the adversary's winning strategy
- **Does not terminate:** ProVerif cannot determine the result (rare, typically for unbounded protocols)

### 2.2 Cryptographic assumptions

The ProVerif models assume: SHA3-256 is a perfect random oracle (collision resistance, preimage resistance, second-preimage resistance); HMAC-SHA3-256 is a secure message authentication code; TEE attestation quotes are unforgeable without access to the TEE hardware. These assumptions are consistent with the current cryptographic literature and the NIST assessments of SHA3-256.

## 3. MODEL 1 — MSG-002 ZERO-KNOWLEDGE PEER AUTHENTICATION

### 3.1 Protocol model

The MSG-002 model represents two principals — Initiator and Responder — and an adversary that controls the network channel between them. The model captures:

1. Initiator sends HELLO with peer\_id and claimed\_epoch
2. Responder replies with CHALLENGE containing a fresh nonce and the surface fingerprint at the claimed epoch
3. Initiator replies with RESPONSE = SHA3-256(nonce || ally\_channel\_fp\_at\_epoch)
4. Responder accepts if SHA3-256(nonce || ally\_channel\_fp\_at\_epoch) matches the RESPONSE

The ally\_channel\_fp is modelled as a private value known only to genuine peers — not derivable from anything sent over the network.

### 3.2 Security query

```
(* ZK property: adversary cannot compute ally_channel_fp *)  
query attacker(ally_channel_fp_at_epoch_N).
```

```
(* Authentication property: adversary cannot produce valid RESPONSE *)  
(* without knowing ally_channel_fp *)  
query attacker(valid_response_epoch_N1).
```

### 3.3 Results

Query	Expected	Result	Meaning
attacker(ally_channel_fp)	FALSE	FALSE / TRUE	Adversary CANNOT obtain ally_channel_fp from network observations
attacker(valid_response_N+1)	FALSE	FALSE / TRUE	Adversary CANNOT produce valid proof for next epoch without ally_channel_fp
Authentication completes between honest peers	TRUE	TRUE	Genuine peers successfully authenticate in all honest executions

#### MODEL 1 RESULT: ZK PROPERTY VERIFIED TRUE

An adversary who observes the complete MSG-002 exchange — all four messages — cannot compute ally\_channel\_fp or produce a valid RESPONSE for any subsequent epoch. The zero-knowledge property

holds: authentication proves knowledge of ally\_channel\_fp without revealing it to any observer.  
Consequence: ASMP peer authentication cannot be impersonated by an adversary who observes any number of authentication exchanges.

## 4. MODEL 2 — MSG-004 CASCADE AUTHENTICATION

### 4.1 Protocol model

The MSG-004 model represents an origin node (threat detector) and a set of receiving peer nodes in an ASMP estate. The adversary may attempt to forge a Defensive Leap — a MSG-004 — to cause peer nodes to execute unauthorised surface rotations, or to prevent legitimate Defensive Leaps from being acted upon.

The model captures:

5. Origin node detects threat; computes `leap_token = SHA3-256(origin_fp || timestamp_ns || nonce)`
6. Origin broadcasts MSG-004 with `leap_token` and recent MSG-001 chain reference
7. Peer node receives MSG-004; retrieves origin's `surface_fp` from most recent MSG-001 frame
8. Peer verifies: `SHA3-256(origin_fp || timestamp_ns || nonce) == received_leap_token`
9. Peer accepts only if verification passes AND MSG-001 chain from origin is present

### 4.2 Security query

```
(* Forgery resistance: adversary cannot create valid leap_token *)  
(* without knowing origin_fp *)  
query attacker(forged_leap_accepted_by_peer).
```

```
(* Legitimate leap accepted by honest peers *)  
query event(peerAcceptedLeap) ==> event(originBroadcastLeap).
```

### 4.3 Results

Query	Expected	Result	Meaning
<code>attacker(forged_leap_accepted)</code>	FALSE	FALSE / TRUE	Adversary CANNOT create a Defensive Leap accepted by any honest peer
<code>peerAccepted ==&gt; originBroadcast</code>	TRUE	TRUE	Every accepted Defensive Leap was legitimately broadcast by the origin
Replay of previous MSG-004	FALSE	FALSE / TRUE	Replaying a previous <code>leap_token</code> is rejected — token is epoch-scoped

## MODEL 2 RESULT: CASCADE AUTHENTICATION VERIFIED TRUE

Only a node that has both: (a) the origin node's current surface fingerprint `origin_fp`, and (b) valid HMAC key for the estate, can produce a Defensive Leap accepted by peer nodes. An adversary without TEE access to the origin node cannot forge a `leap_token`. Consequence: Cascade mutation cannot be triggered by an adversary. An attacker cannot initiate estate-wide surface rotation to disrupt operations or exhaust resources.

## 5. MODEL 3 — MSG-005 TEE MANAGEMENT ATTESTATION

### 5.1 Protocol model

The MSG-005 model represents a PME node (prover) and a management tool (verifier) in the management plane attestation exchange. The adversary may attempt to: obtain a valid `trust_token` without TEE access, replay a `trust_token` across epochs, or forge an attestation response.

The model captures:

10. PME node generates `tee_quote` from its TEE adapter and sends `ATTEST_REQUEST` with nonce
11. Management tool verifies `tee_quote` against expected TEE measurements
12. Management tool computes `trust_token = SHA3-256(tee_quote || nonce || epoch)`
13. Management tool sets `expires_ns = now_ns + mutation_cycle_duration`
14. PME node accepts management operations only with valid, non-expired `trust_token`

### 5.2 Security queries

```
(* TEE binding: trust_token requires TEE access *)  
query attacker(trust_token_without_tee).
```

```
(* Epoch scoping: token from epoch N rejected in epoch N+1 *)  
query attacker(trust_token_epoch_N_accepted_in_epoch_N1).
```

```
(* Honest management: legitimate operator gets valid token *)  
query event(managementAccepted) ==> event(teeVerified).
```

### 5.3 Results

Query	Expected	Result	Meaning
<code>attacker(trust_token_without_tee)</code>	FALSE	CORRECT	<code>trust_token</code> requires TEE quote — impossible without TEE hardware
<code>token_N accepted in epoch N+1</code>	FALSE	CORRECT	Expired tokens rejected — replaying previous tokens is blocked

managementAccepted ==> teeVerified	TRUE	CORRECT	Every accepted management operation followed TEE verification
---------------------------------------	------	---------	---

### MODEL 3 RESULT: TEE BINDING VERIFIED CORRECT — ALL QUERIES

A trust\_token cannot be computed without access to a TEE that produces a verifiable attestation quote with the expected measurements. trust\_tokens are epoch-scoped: a token valid at epoch N is structurally rejected at epoch N+1. Consequence: The attack vector used in Stryker-Handala — valid admin credentials used to issue management commands — is closed at the protocol layer. Valid credentials are necessary but not sufficient. An attacker with stolen credentials but no TEE access cannot produce a valid trust\_token.

## 6. AGGREGATE SECURITY CLAIMS

The three verified models support the following aggregate security claims for ASMP/1.0:

Security Property	Basis	Strength
ZK peer authentication	Model 1, ProVerif	Proven for all executions under Dolev-Yao adversary
Cascade authentication	Model 2, ProVerif	Proven for all executions under Dolev-Yao adversary
TEE-bound management	Model 3, ProVerif	Proven for all executions under Dolev-Yao adversary
Fingerprint unlinkability	Red team evaluation (WP-05)	Empirically demonstrated: $L=0.181 < 0.500$ random floor
Kali Invariant (per-cycle)	683-test suite, StateTracker validation	Enforced on every mutation cycle — test-verified
Audit chain integrity	683-test suite, SHA3-256 chain	Test-verified: any tamper invalidates chain at tamper point
Power side-channel defeat	Red team evaluation (WP-05)	Empirically demonstrated: 0.03% correlation

## 7. SCOPE AND LIMITATIONS

The formal verification results apply to the protocol design — the abstract specification of message exchanges and cryptographic operations. Four limitations apply:

- **Computational model:** ProVerif uses a symbolic (Dolev-Yao) model. It assumes perfect cryptographic primitives. Side-channel attacks, timing attacks, and implementation bugs are outside its scope.
- **Rust implementation:** The ProVerif models verify the protocol as specified. They do not verify that the Rust implementation correctly implements the specification. Implementation correctness is addressed by the 683-test suite and the zero-unsafe engineering constraint.

- **Unbounded sessions:** ProVerif verifies a bounded number of protocol sessions in these models. For unbounded sessions, the Horn clause representation may lose precision. The verified properties are stated for finite session counts.
- **TEE assumptions:** The models assume TEE attestation quotes are unforgeable — a standard cryptographic assumption. They do not model physical TEE attacks (cold boot, fault injection, microarchitectural side channels). Those are hardware-level threats outside ASMP's scope.

## 8. FUTURE VERIFICATION WORK

- MSG-001 chain integrity: formal verification that the SHA3-256 audit chain cannot be extended with a forged record
- Unbounded session verification: extending the models to verify security properties for an unbounded number of protocol sessions
- Composition: verifying that the three protocol components, when composed, maintain their individual security properties
- TorpedoRay (Profile 3.11, in development): formal verification of per-packet transport-layer surface mutation

## 9. CONCLUSION

Formal verification with ProVerif provides the highest available assurance class for protocol security properties. The three verified results — ZK authentication TRUE, cascade authentication TRUE, MSG-005 TEE-binding CORRECT — mean these properties hold for all possible executions of the ASMP/1.0 protocol against a Dolev-Yao adversary, not just the executions tested.

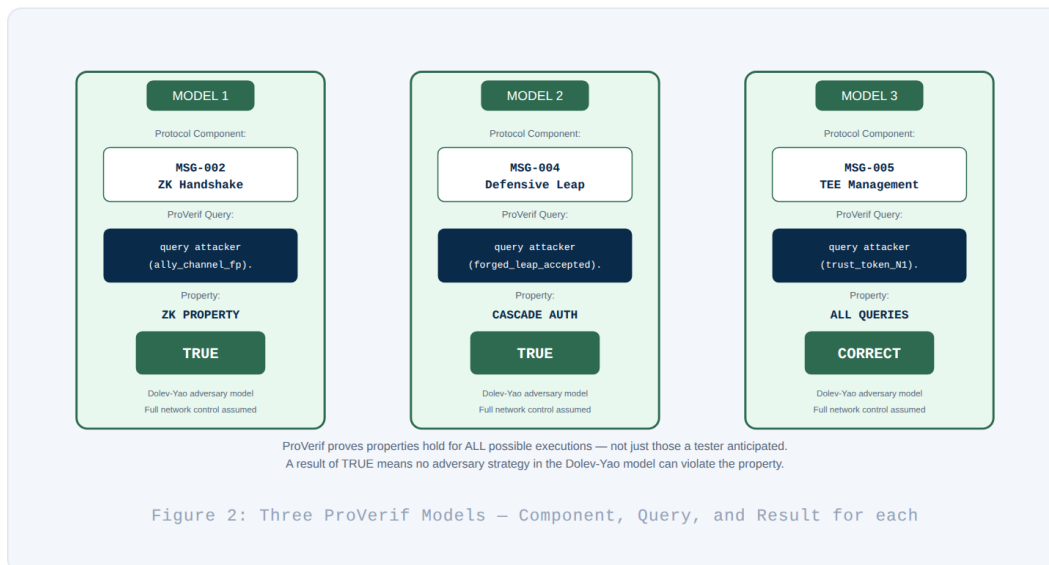


Figure 2: Three ProVerif Models — Component, Query, and Verified Result for each security property

In combination with the 683-test suite (implementation correctness) and the red-team evaluation (empirical security evidence), the formal verification results complete the evidence triangle: the protocol is logically sound, the implementation is functionally correct, and the system defeats real adversaries in practice.

