Mission-Aware Attestation Envelopes for Time-Critical Autonomous Action:

A Hardware-in-the-Loop V2I Study

Dimitrios Nikou[1] Nikolaos Kekatos[2] Sophia Petridou[3]
Stylianos Basagiannis[1]

[1]International Hellenic University [2]Clone Systems [3]University of Macedonia

MESAS 2026 · Prague · 21–23 October 2026

The situation

An ambulance asks for a green light

- The ambulance sends a request to the roadside unit at the junction
- The roadside unit decides: switch the light, or not
- It must decide **before the ambulance arrives**

Before saying yes, it checks two things

- Is this a **registered** vehicle unit?
- Has **its own computer** been tampered with? A hacked roadside unit could say yes to anyone.

Why it matters

Same pattern in defence: a military vehicle asking a checkpoint barrier to open.

How the computer check works today

1. Record

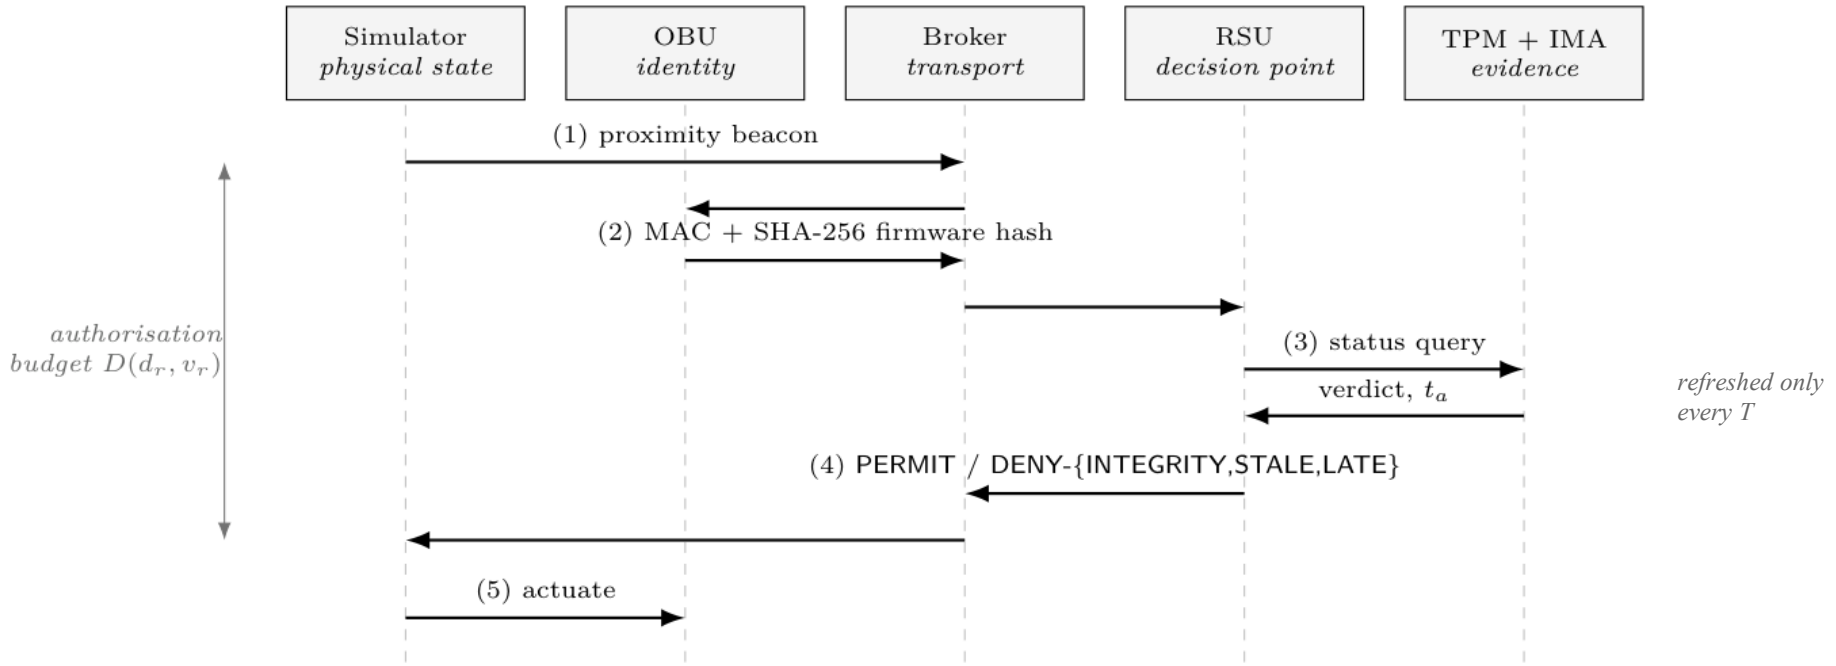
A security chip (TPM) records which software the roadside unit is running.

2. Compare

A checker compares that record with a list of approved software.

3. Answer

The result is a single answer: **trusted** or **not trusted**.



Two checks, two different jobs

Vehicle admission (ESP32)

Registered vehicle identifier

Firmware digest included with request

Not TPM-backed vehicle attestation

Roadside integrity (Linux + TPM)

IMA records software measurements

Keylime verifies the evidence periodically

The verifier status can become stale

Why the distinction matters

An admitted vehicle does not make a compromised roadside decision point trustworthy.

Why a single yes-or-no answer is not enough

The answer is a snapshot

"Trusted" means trusted *when the last check ran*, up to 2 seconds ago.

Deciding takes time

About 0.4 seconds from request to decision on our hardware.

The vehicle keeps moving

At 110 km/h an ambulance covers about 61 metres in 2 seconds.

So the real question is

Is the "trusted" answer still **recent enough** and does the decision arrive **in time**, for this vehicle, at this speed and distance?

The road sets the deadline

20 metres	110 km/h	0.65 seconds
Remaining distance	Vehicle speed	Approximate time to junction

Compare with the measured decision time

95th-percentile decision: about 0.56 s. Only about 0.09 s remains, before signal actuation or clearance time.

Our idea: three questions, four answers

Untampered?
The check reports no violation.

Recent enough?
The last check is not too old.

In time?
The decision lands before the vehicle passes.

Answer	Meaning
PERMIT	all three hold: switch the light
DENY-INTEGRITY	tampered computer or unregistered vehicle unit
DENY-STALE	the last check is too old for this vehicle
DENY-LATE	the decision would arrive too late

A finding: sometimes “too late” cannot be seen

The rule

If the limit on how old a check may be is **no longer than** the time the vehicle allows, then every late decision is also too old.

Why

The check was made *before* the request, so by the time a decision is late the check is older still.

Consequence

The system would always report “too old”, never “too late”. The age limit must be chosen **together with** the vehicle’s time budget.

The ambulance, step by step

110 km/h, request 20 m before	Our answer	Yes-or-no gate
Check just made	PERMIT	yes
Check made 2 s earlier	DENY-STALE	**yes**
Roadside computer tampered with	DENY-INTEGRITY	no

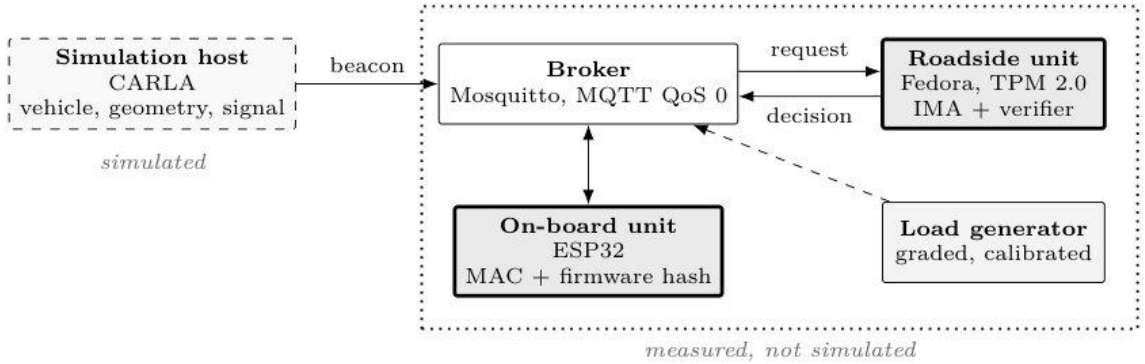
The numbers

Time available: 0.65 s. Our slowest measured decision without network load: 0.62 s, so the answer is in time; only the age of the check decides.

The fix

Start the request about 78 m before the junction, and even a 2-second-old check is acceptable.

The test rig: simulated road, real security hardware



Real hardware

- A small vehicle unit (ESP32 board)
- A roadside computer with a security chip (TPM)
- The network between them, with and without heavy traffic



simulated approach

Real vehicle unit (ESP32): what it sends

```
V2X_OBU: RSU Beacon Received! Triggering
Attestation...
V2X_OBU: Sent Attestation Quote, msg_id=7924
DATA SENT: {"mac":"EC:E3:34:65:7F:6C",

 "hash":"7a204e7a0319a02e2bb9…"}
```

What we measured

Decision time in milliseconds: 50 runs without network load

Configuration	n	mean	median	p95	max obs.
End-to-end, no contention	50	422.8	419.8	559.2	624.2
Attestation component	50	33.3	20.0	93.9	147.3

Heavy network traffic

Slows the **whole** decision; no evidence that it slows the security check.

Caution

50 normal runs and 7 runs under heavy traffic: enough to show that traffic has an effect, not how large it is.

Where a “yes” is still usable

Simulation only vs. real hardware

Simulation only, with instant and always-fresh checks: **60 of 60** situations look fine.
With the real hardware: **46 of 60** work; the fast, close-range ones fail.

Start distance needed, by speed and age of the check

configuration	slope	50 km/h	80 km/h	110 km/h
$\phi = 0$, $T = 2\,\mathrm{s}$	0.56 s	7.8 m	12.4 m	17.1 m
$\phi = 0.5$, $T = 2\,\mathrm{s}$	1.56 s	21.7 m	34.6 m	47.6 m
$\phi = 1$, $T = 2\,\mathrm{s}$	2.56 s	35.5 m	56.9 m	78.2 m
$\phi = 1$, $T = 10\,\mathrm{s}$	10.56 s	146.7 m	234.6 m	322.6 m

Reading

The age of the check costs far more distance than the decision time does.

The setting that matters most

Check less often: timer from 0.5 s to 10 s

interval	feasible	gap	infeas.	note
0.5 s	54 / 60	10.0%	6	below the configured period
1.0 s	52 / 60	13.3%	8	below the configured period
2.0 s	46 / 60	23.3%	14	configured period
5.0 s	35 / 60	41.7%	25	between the two
10.0 s	24 / 60	60.0%	36	maximum the verifier tolerates

Make every decision slower: 1× to 5×

		infeasible	
factor	feasible	age-only	age + latency
1.00	46 / 60	9	5
1.27	44 / 60	9	7
1.50	43 / 60	9	8
2.00	41 / 60	9	10
3.00	35 / 60	9	16
5.00	27 / 60	9	24

Reading

The two have a similar effect, but the timer is a **setting** you can change today, while a 5× slowdown would mean very different hardware or network.

A rule of thumb for engineers

Choose the check timer from the road, not from a default

timer $\leq$ (distance $\div$ speed) $-$ decision time

Longest safe timer	50 km/h	80 km/h	110 km/h
request 20 m before	0.9 s	0.3 s	**0.1 s**
request 50 m before	3.0 s	1.7 s	1.1 s
request 100 m before	6.6 s	3.9 s	2.7 s

Decision time: 0.56 s (95% of our measured decisions). Where no timer fits, check on demand instead.

What we did not show

Scope

- One junction, one vehicle unit, one vehicle at a time
- A simulator, not a real road
- An ordinary network instead of a vehicle radio

Not tested

- An actual attack
- On-demand checking
- Larger test runs under heavy traffic

What carries over

The method: take the deadline from the vehicle's situation, measure the real security delays, and judge the age of the check and the decision time together.

Take-away

In one sentence

A security “yes” is usable only while the check is recent enough and the decision arrives in time for the vehicle that asked, so the check timer and the age limit must be chosen from the mission, not from security settings alone.

Three questions

untampered, recent, in time

Four answers

permit, or deny with the reason

One setting

the check timer

Questions?

Dimitrios Nikou

International Hellenic University

diminiko14@ihu.gr

Backup: glossary

Term	Meaning
Roadside unit (RSU)	the computer at the junction that decides
Vehicle unit (OBU)	the small device in the vehicle that sends the request
TPM	a security chip that records which software is running
Remote attestation	proving to a checker which software a computer runs
Check timer (attestation interval)	how often the check is repeated
Time budget	distance divided by speed: how long the decision may take
HIL (hardware-in-the-loop)	real devices connected to a simulated world