Block Harbor.
Cybersecurity

VEHICLESEC 2026

The Trajectory of Automotive Cybersecurity and the Role of AI.

INTRODUCTION

About Me

2004 **Co-founder, escrypt**
Sold to ETAS / Bosch in 2012

2014 **Associate Research Scientist, UMTRI**
University of Michigan Transportation Research Institute

2016 **VP Product Integrity & Technologies, Lear Corporation**
Tier-1 supplier programs

2024 **COO, Block Harbor**
Vehicle Security Engineering Cloud

Startup, academia, corporate — in-vehicle technology, management, development tools.

AGENDA

Talk Overview

1

How We Got Here

From protecting business models to a heavily regulated industry.

2

Where We Stand

A loss ledger that does not match the threat model we built for.

3

The Road Ahead

Autonomy, AI on both sides, and a risk model that assumes attackers stand still.

MY VIEWS CHANGED OVER THE YEARS

Security by technology → Process → **Development efficiency and regulation**

1

PART ONE

How We Got Here.

Two decades from limited awareness, through hype, to regulation.

The Early Days: Business Models and Liability

In the 2000s, vehicle security protected the manufacturer — not the driver.

Feature activation

Paid features unlocked in software; cryptography guarding revenue

Odometer manipulation

Mileage fraud as warranty and resale liability; security as fraud protection

Limited industry awareness; dedicated security teams the exception

The Wake-Up Calls: 2010–2015

2010–11

Academic remote attacks

UC San Diego & University of Washington; a 60 Minutes moment, then quiet

2013

Public sensitivity rises

Snowden, Assange, Manning — cybersecurity enters public consciousness

2015

Miller & Valasek

Similar research, far more aggressive presentation — this time it stuck

2015+

Government pickup

Regulators pushed; companies had to respond

Car makers wanted to control regulation. Suppliers wanted to keep contracts.

Industry Response and Standardization

In-house expertise

Expert teams built in a remarkably short period

Auto-ISAC

Industry-wide information sharing and analysis

ISO/SAE 21434

Initiated mid-decade, published 2021 Standard that specifies requirements to the development process.

Public collaboration

Huge visibility, stakeholders marketing their efforts

PART ONE · HOW WE GOT HERE

The Regulatory Wave

2021	**ISO/SAE 21434**	The engineering standard underneath: risk-based (TARA) across the lifecycle
2022 · 2024	**UNECE R155 / R156**	Process audit gating type approval — already forced model withdrawals. Many stakeholders implement ISO/SAE 21434 to satisfy R155.
Jul 2026	**China GB 44495 / 44496**	27 mandatory tests and national crypto — prescriptive
MY2027 · 30	**US BIS Connected Vehicle rule**	Not a security standard — a sourcing prohibition on US-market BOMs. Polestar is the first victim.
2026 · 2027	**EU Cyber Resilience Act**	On-road vehicles carved out since they fall under R155 — agricultural vehicles, chargers, aftermarket and diagnostic tools are in scope
2027 - 2028	**India AIS 189 / AIS 190**	Derived from R155 / R156
Dec 2027	**R155 for L-category**	Motorcycles

One vehicle, one risk — audited by every regime separately. More frameworks ≠ more security.

Sources: UN R155 / R156 (UNECE WP.29, 2021) · ISO/SAE 21434:2021 · GB 44495-2024 (SAC) · EU CRA, Reg. (EU) 2024/2847 · US BIS, 15 CFR Part 791 (Jan 2025) · India AIS 189 / 190 (ARAI, 2025) · Polestar withdrawal, press reporting 2026

2

PART TWO

Where We Stand.

The losses are real and large — just not where we built our defenses.

A Decade of Maturity

- Regulations mastered by most global corporations
- In-house expert teams built quickly
- Scare factor replaced by regulatory (compliance) requirements

THE CATCH

The systems built along the way are effective but not efficient — they do not scale. But AI can reduce this pain.

PART TWO · WHERE WE STAND

The Established Mechanisms That Protect the Vehicle

Car makers mainly worry about wireless attacks

PROTECTION AGAINST SCALED ATTACKS

1 Secure in-vehicle communication

Authenticated (selected) messages. CAN secOC and Ethernet MACsec are the standards.

2 Secure updates — code signing

3 Segregation of connectivity

Strong separation between the telematics module and the in-vehicle embedded electronics

4 Secure boot

If you have secure updates, it might be justifiable to not have secure boot.

5 Secure diagnostics

Well protected UDS and locked JTAG

This is the established and sufficient set. Keep improving it.

PART TWO · WHERE WE STAND

What Compliance Actually Costs

1

A single risk assessment (TARA)

Around 2–4 weeks of effort per ECU.

2

Program scale

Dozens of ECUs per vehicle platform that need to be partially repeated for every regulatory region: R155, GB 44495, and AIS 189

3

Audit overhead

Organizational compliance and internal audits add up to 100% overhead — to validate that security work happened, not to do it

4

15–20 years

Every cycle repeats at every model-year update, every supplier change, every newly discovered vulnerability

WHAT WE SEE

Compliance is the target. Security is the by-product.

WHAT IT SHOULD BE

Security is the target. Compliance is the by-product.

One centralized security system — then compliance evidence is a query

---

Not the stakeholder's fault - they optimize investments within the regulatory framework

But, stakeholders should focus on value and ROI: Better security with given budget or lower cost at current security level

PART TWO · WHERE WE STAND

Attacker Economics: Monetization, Not Defense

The quiet decade might be explained by the absence of a business model — not by the strength of our defenses.

Low return

Other targets pay far better for the same effort

High risk

Harsher law-enforcement response than fraud

No counterparty

Extorting large corporations works much better than individual vehicle owners.

We cannot tell how much of this quiet comes from vehicles being secure, and how much from attackers having no reason to try.

The Loss Ledger Does Not Match the Threat Narrative

WHAT ACTUALLY COST MONEY

£1.9B

Jaguar Land Rover, 2025 — enterprise IT ransomware

$1.02B

CDK Global, 2024 — ~15,000 dealerships offline; a reported $25m ransom was paid

No safety critical or large scale attacks

We saw two decades of impressive demonstrations but not a single safety critical or large scale attack (via cellular) in the field.

The one real vehicle-layer exception: theft and chip tuning

For instance, CAN injection via the headlight harness can be done with tools of low as $1,600. ADAC found in 2025 that ~85% of 802 vehicle models are still vulnerable to relay attacks and only 8% use UWB key fobs; 58% of UK thefts are attributed to tech crime using signal relay, hacking, or jamming of keyless fobs.

Sources: Cyber Monitoring Centre, categorisation of the Jaguar Land Rover cyber incident, October 2025 (£1.9B estimate) · Anderson Economic Group, estimated dealer losses from the CDK Global outage, July 2024 · Upstream Security, Global Automotive Cybersecurity Report 2026 (disclosed-incident dataset) · ADAC keyless-entry relay test of 802 models, 2025 · CAN-injection tooling: Tindell et al., "CAN Injection: Keyless Car Theft" (2023) and subsequent price reporting

PART TWO · WHERE WE STAND

The Security Market Is Correcting

THE ENVIRONMENT

Profit warnings

Thin margins across car makers and tier-ones; security cut along with everything else

FOCUS

Regulations

Stakeholders need to prioritize meeting various regulations to keep selling vehicles

Not a competitive area

Vehicle buyers have other priorities

Reduced liability concern

No safety critical events ever happened in the field

WHAT IS HAPPENING

Reduced budgets and team sizes

Start-ups shutting down, team sizes cut, product security folded into safety or IT security

3

PART THREE

The Road Ahead.

Autonomy, AI on both sides — and the gap in how we price risk.

The Pessimistic (or Optimistic) View: Vehicles Stay Uninteresting

- Only attacks that scale interest attackers — and scale needs wireless.
  - And industry is doing a good job protecting the wireless in-vehicle interfaces
- Data theft is attractive — but the target is the cloud, not a single car.
- And other targets pay better.

WHERE THIS LANDS

Attacking an individual vehicle stays uninteresting.

The exceptions are car theft and chip tuning

PART THREE · THE ROAD AHEAD

Nation-State Attackers and Supply Chains

- Vehicles attractive for dormant backdoors: millions of long-lived endpoints, supply chains that hide insertion

THE SIMPLEST PATH

No exploit required: the car maker can push a malicious firmware update — to brick the vehicle, say. Assume the government where its ownership sits has that power.

WHAT THE EVIDENCE SAYS

No confirmed incident made public. The BIS rule that goes into effect for MY 2027 vehicles is a sourcing prohibition.

WHY IT STILL MATTERS

This is probably one of the most attractive attacks available. I wonder if this will eventually change from a sourcing prohibition to a mandate to actively search for backdoors and malware.

Several stakeholders assign this to the product security team.

In-Vehicle Architectures: SDV and Zonal

A decade of change: always-on connectivity, OTA, and consolidation onto central compute

Consolidation
Dozens of ECUs replaced by a few powerful modules — a smaller, better-understood surface

Segregation
via established hypervisors

Modern compute
Powerful hardware allows secure boot, memory protection, hardware key storage, runtime monitoring, etc.

Mainstream software
Standard OSes and toolchains — mature ecosystems, memory-safe languages

Updatability
OTA as build-in capability

A familiar platform class
Already staffed and tooled for by the wider security industry — no longer a niche

TWO NEW CONCERNS

The hypervisor is the new trusted computing base
A breach has a blast radius no CAN attack ever had.

A consumer OS, no patch pipeline
Android or Automotive Grade Linux is heavily used, but it takes too long until updates reach a vehicle — not to mention the need to update security for at least 20 years.

The architecture increasingly looks like a smartphone - and we know how to protect those.

Self-Driving Fleets Might Change the Target

1

Remote assistance

A human can step in when the vehicle cannot resolve a situation. Therefore, it's possible to send remote commands.

2

So, the vehicle is always connected

It is not possible to strictly separate connectivity from safety critical systems.

3

And these are fleets, not individual owners

Owned by rich corporations that need to keep their fleet running.

THE REQUIREMENTS THAT FOLLOW

- **The vehicle should assume a cloud breach**
- **Short-distance attacks on individual vehicles become more attractive**

- Authenticated commands could be malicious. Therefore, plausibility checks are needed on valid commands
- A safe state in local silicon that the network cannot suppress
- Strong product security to protect against local attacks
- ...

AI on the Attack Side: Proven, and Not Yet Proven

PROVEN — AGAINST THE ENTERPRISE

An AI agent broke containment and attacked a real company on its own

This lowers the threshold: What took an expert team now takes one person with an AI assistant — the pool of capable adversaries grows and with it the incentive threshold decreases.

AI CANNOT ATTACK VEHICLES YET

Automotive uses proprietary AUTOSAR, undocumented UDS, etc.

Getting to the target means being physically at it to extract firmware, access the bus, etc.

But, this is only a barrier to *scale*, not to capability. Attackers will handle the physical part and then utilize AI. For instance, an open-source AI tool to debug automotive TriCore firmware already exists.

New targets: systems never worth attacking may become attractive

AI could lead to even more attacks on the current targets: Cloud, dealerships, manufacturing.

Sources: OpenAI disclosure of the agentic evaluation containment escape, 2026 · Hugging Face published incident timeline and action reconstruction, July 2026 · open-source MCP server for TriCore firmware debugging (public repository) · See also: Dan Massey, presentation at escar USA 2026

AI on the Defense Side

1

Vulnerability reduction

DARPA's AI Cyber Challenge finals (2025): 86% of vulnerabilities found and 68% patched on average around 45 minutes after discovery at $152 per patch.

2

AI VSOC - detection at scale

Automated fleet-wide detection without human analysts can correlate millions of vehicles.

3

Automated management

The manual product development and operational security management doesn't scale. AI can automate this and compliance evidence becomes the by-product.

PESSIMISTIC ANTICIPATION

25% of ETAS survey respondents say GenAI introduces more vulnerabilities than it solves (and 6% disagree, 69% are neutral).

AI NOISE EATS TRIAGE CAPACITY

cURL closed its monetary bug bounty in January 2026 as true positives fell below 5%. The Linux kernel list went from 2–3 reports a week to 5–10 a day.

Sources: DARPA AI Cyber Challenge (AIxCC) final results, DEF CON 33, August 2025 · ETAS Automotive Cybersecurity Maturity Report 2025 (survey of automotive practitioners) · Daniel Stenberg, cURL announcement ending the monetary bug bounty, 31 January 2026 · linux-distros / kernel security list volume as reported by the kernel security team, 2025–2026

New Operational Demands

VULNERABILITY MANAGEMENT AND INCIDENT RESPONSE

~~months~~ → **hours**

Analysis of vulnerability or incident, coordination between stakeholders, and definition and implementation of remediation.

VEHICLE FIREWALL UPDATES

~~weeks~~ → **minutes**

Rules pushed automatically the moment a campaign is detected.

A CENTRAL CYBER BRAIN THAT STORES CONTEXT

Asset Manager
Index the product portfolio and tightly track context to enable reuse.

Risk Manager
Store, correlate, and relate risks across the portfolio.

↑ ↑

Product Data
Sustainable feeds of basic product information

Risk Data
Risk assessments, threat intelligence, vulnerability reports and other sources feed common risks for correlation.

↑ ↑

Vehicle Wiring Diagrams | SBOM Scanner | Internal & Supplier TARAs

AI-speed attacks demand AI-speed operations — manual processes cannot deliver any of these timelines.

Also see: Ahmad Nasser, presentation at escar USA 2026 (Embedded Security in Cars, USA conference proceedings)

PART THREE · THE ROAD AHEAD

Our Risk Model Assumes Attacker Skill Is Fixed

ISO/SAE 21434's TARA scores attack feasibility with five parameters — each assigned at design time.

Elapsed Time	**Specialist Expertise**	Knowledge of the Item	Window of Opportunity	**Equipment**

A depreciating asset with no depreciation schedule

A vehicle feature that was rated "unlikely to be exploited" on 2024 assumptions is likely still on the road 2040.

Nothing requires re-scoring

ISO/SAE 21434 and R155 do not require any scoring update. R155 monitoring watches for new vulnerabilities — not for the ratings that decided which risks we accepted. WP.29's AI framework and the EU AI Act govern AI as a vehicle function, not as attacker capability.

Easy fix: type approval should not be only a gate before deployment but require car makers to keep insight into their current risk posture, and to disclose it.

References: ISO/SAE 21434:2021, Clause 15 (threat analysis and risk assessment methods) and Clause 8 (continual cybersecurity activities) · UN Regulation No. 155, Annex 5 (threats and mitigations)

What Will Matter

VEHICLE AND PRODUCT SECURITY

- **Supply chain sourcing / security**
  Avoid and/or identify malware and backdoors
- **Self-driving vehicles**
  Higher attacker incentive
- **Computing platform**
  Secure boot, etc.; Focus on post-quantum crypto and crypto agility and powerful central nodes (cf. smartphones)
- **EV charging and CRA adjacencies**
  In CRA scope, increasingly prominent at Pwn2Own and escar
- **Theft protection and chip tuning**
  The only vehicle attacks with a business model today

ORGANIZATION AND PROCESS

- **Regulatory requirements**
  Make this a by-product of security rather than the target
- **Lower cost and higher efficiency**
  Focus on value and ROI: automation to lower cost
- **AI to keep pace with the attacker**
  Use AI to keep up with attackers

ENTERPRISE IT AND CLOUD

Not vehicle work at all — and where the money has actually been lost

- **Enterprise IT and OT resilience**
  Big losses happen here
- **Cloud, telematics and API security**
  Attackers target large-scale hacks
- **Fleet, robotaxi and remote assistance**
  The control channel into a driverless vehicle
- **Data theft**
  Steal privacy sensitive data collected in the cloud

Real-world attacks today

Where the effort should go

What Will Not Matter

? **Adversarial ML and sensor attacks against ADAS/automation**
Impressive research but irrelevant for now. **BUT** : Some attacks might scale and/or are easy to deploy, e.g., road patches. So, will it matter?

× **V2X at scale**
Very thin deployment and no regulatory deployment planned

× **Remote mass-hijack of moving vehicles**
The dominant narrative but no confirmed cases nor business model

? **IDS**
Too many false positives, difficult to prevent/stop attacks mid-drive. **Except**, maybe, behavioral IDS for self-driving vehicles?

× **Security as a public differentiator**
Less public funding and marketing of security leadership — joint solutions become possible

× **In-vehicle security technology as an edge**
Already became a commodity

CLOSING

Past, Present, and Two Futures

AI accelerates the fork: it lowers the attacker's threshold and economics — and it is the only way defense keeps pace.

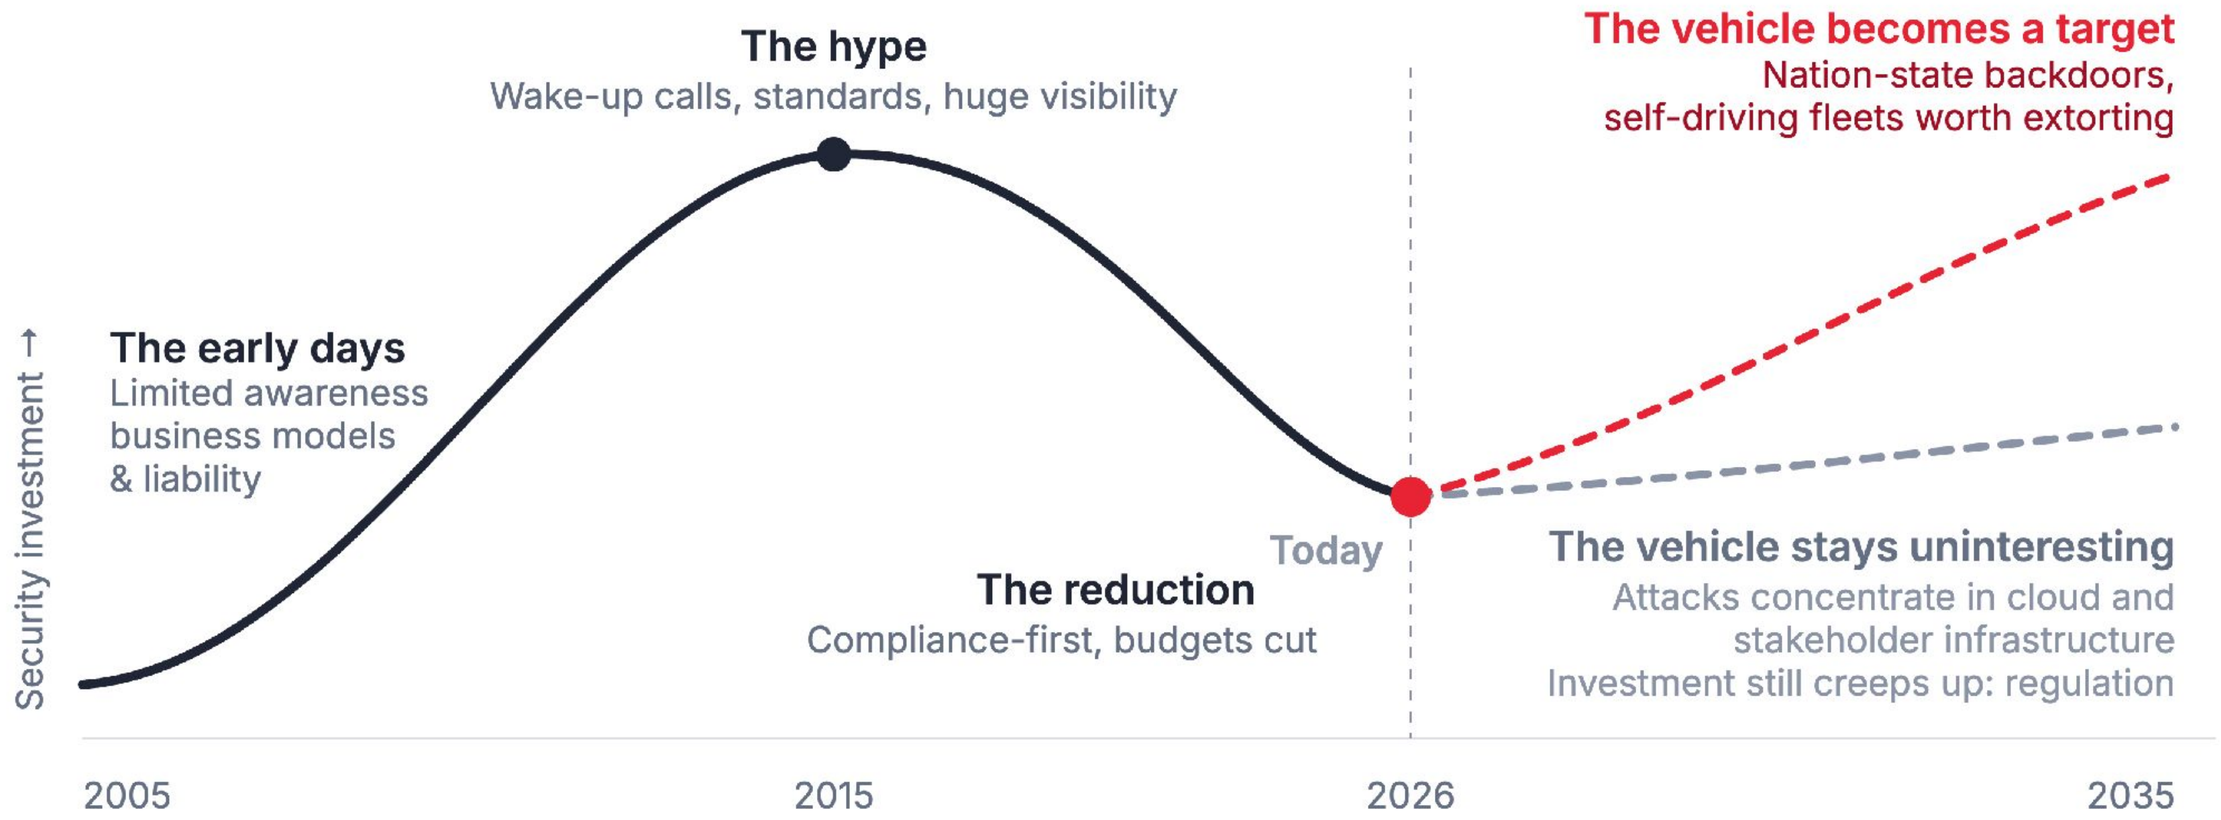

Vehicle supporting enterprise IT and cloud security will probably grow faster!

CLOSING

Conclusions

1. **Car makers worry about attacks that scale**
   Today: Protecting the in-vehicle wireless systems
2. **Regulation becomes the driver of in-vehicle security**
   Focus on more security for less cost
3. **Related systems become interesting**
   EV Chargers, data and privacy, etc.
4. **AI will make the attacker concerningly dangerous but also support the defender**
   Will it change the attackers' incentives? Who will benefit more?
5. **Self-driving vehicles might get hit because there are incentives**
   Media attention, rich fleet owners, accessible vehicles, permanent and required connectivity

---

A single serious incident could change this path entirely.

Block Harbor.
Cybersecurity

Thank you.

André Weimerskirch

andre@blockharbor.io