

ENVITED Community Meeting

From SDV to AIDV: Enabling Smart Vehicles Through Data-Driven Engineering

3 November 2026 | Munich (House of Communication)

Pre-event to the ASAM International Conference

Software- and AI-defined mobility are reshaping how complex vehicle systems are designed, validated, and operated. Functionality is increasingly software- and AI-driven, continuously evolving over the lifecycle, shifting focus from hardware toward software, data, and continuous validation.

This shift spans multiple domains beyond ADAS and automated driving, including energy management, powertrain optimization, predictive maintenance, and user-facing digital systems.

A key challenge is validating end-to-end and hybrid architectures. Modern systems combine classical modular software with AI-based components, as well as end-to-end approaches where behavior is learned from data. These paradigms challenge traditional validation, as system behavior emerges from complex interactions rather than deterministic logic. New strategies are required, covering component-level verification, system integration, full-system behavior in dynamic scenarios, and continuous validation under OTA updates.

Data-driven Engineering and Simulation play a central role - not only for development and testing, but as a foundation for scalable validation and structured evidence generation across the lifecycle. This increases the importance of simulation credibility, including transparency, traceability, and reliability.

The ENVITED Community Meeting, organized by the Automotive Solution Center for Simulation, brings together industry and academic experts to address these challenges. As a pre-event to the ASAM International Conference, it connects methodological innovation, industrial application, and standardization across software-defined vehicle systems.

envited
Research Cluster
Environment for Virtual Test Drive

sposored by Ansys part of SYNOPSYS®

AGENDA

3 November 2026

10:15 AM CET	***REGISTRATION*** Networking and Warm-up
10:45 AM	**Welcome & Introduction** Alexander F. Walser | Automotive Solution Center for Simulation e.V.
10:50 AM	***KEYNOTE:*** **The Next Function Won't Be Written - Why Validation is Becoming Part of Development** Marc Holzäpfel | Holzäpfel Engineering Advisory
11:20 AM	**Building a Digital Twin - A Terrain Engineer's Perspective** Mirco NIERENZ | TrianGraphics GmbH
11:40 AM	**Quantifying Simulation Quality for Virtual Homologation** Sana KAHN| AAI Innovations GmbH
12:00 noon	**From Requirements to Continuous Homologation:** **Agentic AI for ODD Aware ADAS & Autonomous Driving Validation.** Stephen LERNOUT | Deontic
12:20 PM	***LUNCH BREAK***
1:20 PM	**Re-Engineering Automotive Development with Virtualization and AI** **Shift-left through virtual first - from System to Silicon** Matthias LÖFFLER | ANSYS part of Synopsys
1:40 PM	**Scalable SDV Validation & Data Strategy: Bridging Cloud-Native SiL, Modular HiL and E2E Fleet Testing through an Open AI-Driven Ecosystem** Guillaume OBDAM | AKKODIS Germany Consulting GmbH
2:00 PM	**Code Compliance as the Foundation for Trustworthy AI-De fi ned Vehicle Development** Andreas LAURINGER| Kontrol GmbH
2:20 PM	**Missing Edges First: Deterministic Gap Analysis on a Safety Evidence Graph** Arnd HEKERMANS | viprove
2:40 PM	***COFFEE BREAK***
3:20 PM	**AI-Enabled Validation of Autonomous Driving Functions in Critical Project Phases** Martin HENNE, Jasmin PUCHALOWITZ | EDAG GmbH
3:40 PM	**From Static Approval Evidence to Dynamic Evidence Systems:** **Rethinking Virtual Homologation for Software-De fi ned and AI-Enabled Vehicles** Sabine RÖHRICHT, Udo KAEMPF | MHP Management und IT-Beratung GmbH
4:00 PM	**From Ticket Chaos to Trace Intelligence: Scaling AI-Driven Defect Management** Akhil RAJAGOPAL, Hanno STAGE | MHP Management- und IT-Beratung GmbH
4:20 PM	**World Café**
5:30 PM	***END***

ABSTRACTS

KEYNOTE

The Next Function Won't Be Written - Why Validation is Becoming Part of Development

Marc HOLZÄPFEL | Holzäpfel Engineering Advisory

The software-defined vehicle has laid the groundwork for the AI-defined vehicle. Central computers with access to sensors and actuators, together with connectivity, make it possible to bring learned functions into every domain of the vehicle.

In the software-defined vehicle, engineers write a function and use testing and simulation to confirm what they wrote. More and more functions will no longer be written line by line. Engineers will describe the target behavior, train for it and find out what the function actually does through testing and simulation.

A new generation of robotics companies shows how fast this can go. These companies work on a similar problem with much shorter product cycles. Their robots learn by demonstration and train in simulation. Development on the left side of the V-model is speeding up. Validation on the right side is not yet built for that pace.

Verification and validation move from the end of development into its center. The object lists that today's tools rely on are disappearing inside end-to-end models. When my teams moved driver assistance to an AI-first architecture, we could train models faster than our tools could analyze them. Engineers use closed-loop simulation and testing to see what a model does and decide what it learns next. This loop does not stop at start of production. It runs again with every update. And the results have to hold up in a safety case.

Automotive engineering already knows how to validate systems it did not build itself. For decades, suppliers delivered them as black boxes. In this keynote I look at what that experience is still worth when functions learn. The shift asks a lot of the tools this community builds. It also asks organizations to decide as fast as their functions learn.



Marc HOLZÄPFEL has spent 18 years bringing automated driving from research into series production. At Porsche, he delivered InnoDrive, a predictive driver assistance system, to market. At CARIAD, he led the Automated Driving Alliance with Bosch: a program of more than 1,200 engineers serving multiple Volkswagen Group brands. There he drove the shift to an AI-first architecture and built the cross-company governance that turned unpredictable releases into a fast, reliable cadence. Today, as an independent Engineering Advisor, he helps technology companies build engineering ecosystems for complex cyber-physical products.

11:20 AM

Building a Digital Twin - A Terrain Engineer's Perspective

Mirco NIERENZ | TrianGraphics GmbH

“What does it cost to model ten kilometres of roads through a small town?” This seemingly simple customer request opens up a range of engineering questions. Which details matter for the intended simulation? What data is available, and how much work is needed to turn it into a usable digital twin?

This presentation offers a practical view of creating simulation environments from the perspective of a terrain engineer. Using the small-town example as a common thread, it follows the process from defining requirements and assessing geospatial data to preparing and enriching road networks, buildings and their surroundings. It highlights how data quality, missing attributes and the required level of detail influence both modelling decisions and project effort.

The talk also examines where procedural methods, automation and AI can help, including data enrichment, texture generation and the creation of 3D assets. It considers the review and rework these methods require, alongside their potential to reduce manual effort.

Finally, the presentation addresses export, validation and ongoing project maintenance. It discusses why standards compliance needs to be complemented by checks of road logic, geometry and behaviour in the target simulator. The central message is practical: the length of a road network alone says little about the effort required. Clear requirements, suitable source data and agreed acceptance criteria provide the basis for planning and delivering a digital twin that serves its intended simulation purpose.

11:40 AM

Quantifying Simulation Quality for Virtual Homologation

Sana KAHN| AAI Innovations GmbH

As vehicles become software and AI defined, physical testing alone can no longer cover the scenario space needed for safe approval. Regulations such as UN R157, the EU ADS Regulation (EU) 2022/1426 and the UNECE NATM framework increasingly accept simulation as evidence. However, a virtual test result is only as trustworthy as the toolchain behind it: the scenario, the environment and map data, the sensor and vehicle models, and the simulation framework itself.

Today, simulation quality is described in many different, mostly non comparable ways. Credibility assessments are often manual, tool specific and hard to exchange between OEMs, suppliers, technical services and authorities.

This talk presents the approach of the ASAM QSQ (Quantifying Simulation Quality) project. QSQ works towards industry consensus on how simulation quality is defined, measured and communicated, and towards a standardized, machine readable format for storing and exchanging quality metrics. Key messages:

1. Virtual homologation needs measurable, not only documented, simulation quality.
2. Quality metrics must cover the whole toolchain, from HD maps and 3D environments to models and frameworks, and must be traceable across the ASAM OpenX ecosystem.
3. A shared, machine readable metrics format makes credibility evidence comparable and reusable, and allows quality checks to run continuously in CI/CD and OTA driven development.
4. Pre-competitive cooperation is required to turn this into accepted practice between industry, technical services and regulators.

The talk closes with the current status of QSQ, open questions, and how the community can contribute.

12:00 noon

From Requirements to Continuous Homologation:
Agentic AI for ODD Aware ADAS & Autonomous Driving Validation.

Stephen LERNOUT | Deontic

ADAS and autonomous driving systems must be validated across an exploding number of requirements, Operational Design Domains, geographic regions, and edge cases. Yet much of today's scenario-based validation workflow remains manual. Engineers translate requirements into test scenarios, construct road networks, adapt scenarios to new ODDs and repeat the process whenever regulations, standards or system requirements change. As autonomy scales, this approach becomes increasingly difficult to sustain. This presentation introduces an agentic AI approach to turning requirements and ODD definitions directly into standards compliant simulation scenarios, creating the foundation for a more automated and ultimately continuous homologation workflow.

Deontic uses generative AI agents to interpret natural language requirements and automatically generate and validate the corresponding road networks, scenes, OpenDRIVE and OpenSCENARIO assets. ASAM standards provide the interoperability layer connecting generated scenarios with existing simulation and validation environments. The approach extends beyond individual scenario generation by making the ODD a first class component of validation. Using structured ODD models based on ISO 34503 and ASAM OpenODD, scenarios can be generated within an ODD, at its boundaries and across relevant combinations of road infrastructure, speed, environmental conditions and geographic constraints. This enables engineering teams to systematically improve scenario diversity and ODD coverage rather than simply increasing scenario volume.

The presentation also outlines the path toward continuous homologation. Regulatory and standards changes can be monitored automatically, allowing an AI agent to identify which validation requirements, scenarios or scenario libraries need to be updated. In the next stage, results from simulation environments can flow back into the validation process, creating traceable links between requirements, ODDs, scenarios, simulation outcomes and homologation evidence.

Finally, we explore how this validation intelligence layer can connect with emerging world models and Vision Language Action systems, using structured scenarios not only to generate simulation assets but also to evaluate increasingly complex autonomous driving AI.

The result is a shift from manual scenario engineering toward an ODD aware, standards based and continuously evolving validation pipeline, enabling autonomy teams to scale validation as quickly as their systems, markets and regulatory environments evolve.

1:20 PM

Re-Engineering Automotive Development with Virtualization and AI

Shift-left through virtual first - from System to Silicon

Matthias LÖFFLER | ANSYS part of Synopsys

Automotive development is entering a new era where increasing system complexity, software-defined architectures, and faster innovation cycles demand a fundamental rethink of traditional engineering. This presentation explores how a virtual-first, AI-enabled development approach can shift engineering left, connecting system-level decisions with software and silicon much earlier in the development lifecycle. By combining virtualization, simulation, and AI across the System-to-Silicon continuum, engineering teams can evaluate architectures earlier, accelerate development and validation, and reduce dependency on physical prototypes. The result is a more integrated, agile development model built for the complexity of next-generation vehicles.

1:40 PM

Scalable SDV Validation & Data Strategy: Bridging Cloud-Native SiL, Modular HiL and E2E Fleet Testing through an Open AI-Driven Ecosystem

Guillaume OBDAM | AKKODIS Germany Consulting GmbH

Summary

The transition to Software-Defined Vehicles (SDV) has rendered traditional, sequential V-Model validation cycles as a bottleneck for DevOps-driven CI/CD. Currently, the automotive industry faces a fragmented landscape where development, validation, production, and aftersales operate in technical silos.

This is further complicated by complex standards that often function as vague frameworks rather than concrete bases for validation. Consequently, validation tools are frequently retroactively adapted to specific implementations rather than addressing deviations from core requirements. As OEMs move away from universal standards, the industry is left with unclear requirements and a proliferation of customer-specific, error-prone tool adaptations. These inconsistent IT architectures result in data being collected but only superficially evaluated, while valuable insights from vehicle fleets remain disconnected from early-stage development. Furthermore, varying global regulations regarding security and data protection necessitate rigorous anonymization, adding another layer of complexity.

This contribution focuses on bridging the gap between development, validation, and pre-launch engineering with real-world performance by focusing on two key pillars:

- Development & Validation Phase: The primary objective is to refine system reliability and performance before market entry. By utilizing advanced simulation and testing frameworks, this approach ensures that products meet stringent safety and functional requirements during the design freeze and SOP (Start of Production) lead-up.
- Early Aftersales Integration: A distinctive "Early Warning" capability integrates data from the first months of field operation to validate development assumptions against real-world stressors. This allows for Rapid Iteration of validation test cases and Data-Driven Calibration of digital twins using field telemetry.

The proposed "kit-of-parts" system integrates commercial validation and data logger tools with open-source components. By strictly adhering to established standards (e.g., Eclipse openMDM®, ASAM ODS, ASAM MDF/OSI, Modelica FMI/FMU, and OSRF ROS2), the system achieves unprecedented environmental parity between virtual SiL and real-time HiL. To address exponential test-space complexity, the framework implements a dual-engine execution strategy: Rule-based Testing to ensure functional safety (ISO 26262) and AI-driven Testing to employ generative scenario modeling for "edge case" discovery and automated anomaly detection.Ultimately, this approach establishes a coordinated, cross-platform data foundation spanning virtual cloud instances, real-time laboratory setups, and end-to-end vehicle environments - essential for reducing time-to-market while ensuring comprehensive vehicle safety.

Degree of Innovation

The core innovation is a unified data warehouse that collapses traditional silos between development, production, and aftersales. By establishing a cross-platform data foundation - synchronizing with the cloud, the lab, the test vehicle and vehicles using shadow test mode - this framework enables OEMs to accelerate time-to-market while maintaining rigorous, real-time safety validation.

2:00 PM

Code Compliance as the Foundation for Trustworthy AI-Defined Vehicle Development

Andreas LAURINGER| Kontrol GmbH

As software-defined and AI-defined vehicles continue to increase in complexity, ensuring compliance, traceability, and validation across the complete development lifecycle becomes a critical challenge. At Kontrol, we address this challenge through a code-centric compliance approach that integrates software quality, regulatory requirements, safety expectations, and engineering evidence into a continuous validation framework.

The presentation will demonstrate how code compliance can serve as the backbone for modern vehicle development and validation by connecting development artifacts, simulation results, test evidence, and governance requirements in a single traceable workflow.

The session will cover:

- Simulation for continuous validation and OTA-driven development through automated compliance monitoring and evidence generation.
- Validation of hybrid and end-to-end system architectures with traceable links between requirements, code, models, tests, and deployed software.
- Virtual homologation approaches supported by auditable compliance evidence and digital engineering artifacts.
- Integration of simulation and compliance verification into modern CI/CD pipelines.
- AI-enabled simulation and data-driven testing approaches with transparent governance and validation records.
- Data quality, traceability, and trustworthy simulation supported by end-to-end lineage and compliance controls.
- Interoperability and standardization through alignment with industry standards and ecosystems such as ASAM OpenX.
- Opportunities for pre-competitive collaboration around common compliance, validation, and traceability challenges facing the automotive industry.
- Using practical examples from software-intensive engineering environments, the presentation will show how organizations can reduce compliance effort, improve audit readiness, accelerate development cycles, and build confidence in increasingly autonomous and AI-enabled vehicle systems.

2:20 PM

Missing Edges First: Deterministic Gap Analysis on a Safety Evidence Graph

Arnd HEKERMANS | viprove

Software- and AI-defined vehicles produce requirements, scenarios, simulation assets, test runs and ALM links faster than anyone can decide whether the resulting evidence holds together. What stays open is the assessor's question: which required links are missing, which present links do not prove what they claim - and will the same check give the same answer tomorrow? Homologation rarely fails on an artifact somebody produced. It fails on an edge nobody drew - or on an edge that is there, green, and proves nothing.

This talk presents viprove as a typed safety evidence graph with Visualize and Analyze as its core stages - with report and seal as Analyze outputs - plus an owned Execute path only where we run the stack ourselves. Visualize builds the graph: project artifacts - requirements, OpenSCENARIO scenarios, ALM items, tests, runs - ingested as typed nodes and edges against a frozen domain layer. The graph is what makes a missing edge nameable; which edges are required is decided by the grammar, in the next step. Analyze runs a versioned, executable grammar pack over that graph and returns a scoped verdict: findings by type (norm, SYS, test, over-fill, integrity including ASIL, timing by reason, grammar), a report, and a seal. New run evidence — whether from that owned SiL/IT Execute path or from external HiL/CI - is re-ingested and re-analyzed the same way. Report and seal are always outputs of Analyze, not of the run source.

Domain mapping and grammar pack are frozen together, pinned by a domain hash and a grammar version; the seal is a SHA-256 over the canonical report body, excluding only timestamp and seal - so the same graph under the same freeze reproduces the same hash.

Determinism is deliberately scoped. It holds for Analyze against a given freeze: same graph, same grammar version, same counting policy, same seal. LLM assistance is used only to map normative documents into the domain layer, under human review, and never in the verdict. A green result means structural consistency under one grammar pack, not a release and not a homologation statement; an empty domain layer yields zero norm gaps, which is a warning and not compliance; and a seal is a replayable record, not a certificate.

We demonstrate live on an L2 ADAS evidence graph - lane keeping and ACC, hazard to safety goal to SYS.2 to SWE.1 to evidence - where a single sealed Analyze run returns a typed finding inventory (system gaps, integrity, timing by reason, grammar, and related checks). The inventory stays on screen; we walk two of its patterns in depth. First, the ACC longitudinal command path under an ASIL C goal: a system requirement with no software requirement beneath it, and no independent monitor on the same goal. Second, a timing verdict where a HiL PASS and a green verifies link cannot serve its safety goal because the claim target (200 ms) does not match that goal's FTTI (150 ms). Other timing findings remain visible in the same seal - for example a pending run record, or goals with no FTTI reaction evidence yet. We then re-run Analyze on stage and get the identical seal - the "same answer tomorrow" question, answered live. This is a graph over artifacts, not a live simulation run. The demo project was authored with known defects; the grammar pack was not told where they are.A second, complete Analyze run on a robotics safety stack - different domain pack and integrity model (PL/SIL), using the same Visualize · Analyze path with a sealed result - shows the engine is not automotive-only.

The ENVITED-X integration point - Services and Checks as the process category, a validation report as the output - is presented as a documented design, not a connected dock. We close with two open questions for discussion: how to express confidence in PDF-based norm extraction, and how far gap closing can be automated without leaving the deterministic core.

3:20 PM

AI-Enabled Validation of Autonomous Driving Functions in Critical Project Phases

Martin HENNE, Jasmin PUCHALOWITZ | EDAG GmbH

The validation and release of autonomous driving functions remain among the most challenging activities in automotive development, particularly during late project stages when homologation, release decisions, and quality targets converge. While modern development environments increasingly leverage simulation and automation, final validation activities still rely heavily on manual testing at vehicle level or on full-system integration test benches. This approach is constrained by limited vehicle availability, labor-intensive test execution, and software tool chains that are often expensive, fragmented, and insufficiently aligned with the speed requirements of contemporary development processes.

At the same time, recent initiatives across automotive OEMs, suppliers, and engineering service providers clearly demonstrate a growing expectation that Artificial Intelligence will become an integral part of the entire validation process. From automated test generation and intelligent scenario selection to AI-supported result analysis and risk-based validation strategies, numerous approaches are emerging. However, despite significant momentum, the industry still lacks broadly established methods, standards, and tool ecosystems that enable scalable and trustworthy AI-driven validation for safety-critical systems.

This presentation provides a practical overview of the current state of AI adoption in the validation of autonomous driving functions, focusing on the challenges and opportunities encountered during critical project phases. Using selected real-world examples, it highlights how AI technologies can increase testing efficiency, improve coverage, accelerate release activities, and reduce dependency on scarce physical test resources. Furthermore, the presentation discusses the benefits already achieved in industrial projects, outlines remaining technical and organizational challenges, and explores the future potential of AI-driven validation workflows for homologation-relevant testing and safety assurance.

By sharing lessons learned, emerging best practices, and concrete application examples, the presentation aims to contribute to the ongoing discussion on how AI can transform the validation landscape for autonomous and software-defined vehicles.

3:40 PM

From Static Approval Evidence to Dynamic Evidence Systems:
Rethinking Virtual Homologation for Software-Defined and AI-Enabled Vehicles

Sabine RÖHRICHT, Udo KAEMPF | MHP Management und IT-Beratung GmbH

Software-defined and AI-enabled vehicles are changing how vehicles are developed, validated and regulatorily assured. While type-approval frameworks still focus on vehicle types, systems and components, approval-relevant behaviour is increasingly being shaped by software, AI models, data and configurations that evolve through controlled releases.

This raises a central question: if the legal approval object remains largely unchanged, but the underlying vehicle configuration becomes more dynamic, how must the assurance system supporting approval evolve?

Virtual homologation must therefore evolve from primarily point-in-time evidence generation towards a lifecycle-oriented evidence system. Approval-relevant characteristics need to remain demonstrable across software releases, over-the-air updates, configuration variants and complex system interactions. For AI-enabled functions, this challenge is amplified because behaviour may depend not only on software versions, but also on models, data and operational conditions.

Simulation and virtual testing can provide scalable evidence across variants, scenarios and releases. From a regulatory assurance perspective, however, virtual evidence must not only be credible when generated, but also remain traceable to the relevant requirement, vehicle configuration and approval decision. Its applicability and validity must therefore be assessable when software, AI models or system configurations change.

The lifecycle perspective further intensifies these requirements. Software and AI-enabled functionality can evolve long after start of production. While not every update requires renewed approval, approval-relevant changes require systematic impact assessment, validation and evidence generation. Frameworks such as UN Regulations Nos. 155 and 156, ISO 24089 and ISO/PAS 8800 increasingly reflect this shift towards lifecycle-oriented assurance.

The contribution examines how approval-relevant changes can be identified efficiently, how virtual evidence can be structured and maintained throughout the vehicle lifecycle, and how trust can be established across increasingly digital assurance processes. It outlines key capabilities for future assurance architectures: traceability, change-impact assessment, digital evidence management, integration of simulation credibility information and cross-functional governance across development, validation, homologation, cybersecurity and AI.

Together, these capabilities form the basis for scalable virtual homologation in software-defined and AI-enabled vehicles, balancing innovation speed, safety and regulatory confidence throughout the vehicle lifecycle.

4:00 PM

From Ticket Chaos to Trace Intelligence: Scaling AI-Driven Defect Management

Akhil RAJAGOPAL, Hanno STAGE | MHP Management- und IT-Beratung GmbH

Modern software-defined vehicle programmes generate vast numbers of defect tickets and trace logs. Critical engineering decisions are often slowed by manual pre-analysis, fragmented toolchains, and reliance on individual experts. In a typical OEM setting, error management teams handle ~300 new tickets per week, yet up to 40% cannot be analysed immediately due to missing information, causing delays and rework. Engineers manually locate trace files, convert binary logs, search for patterns, and copy findings into downstream systems.

This presentation introduces a data-driven AI approach that automates large parts of this workflow. A scalable pipeline ingests ticket and trace data to verify ticket quality, detect duplicates, generate concise summaries, and perform trace-centric root-cause analysis. The key innovation is the tight integration of ticket management and trace analysis in a single pipeline, enabling consistent, repeatable investigations beyond isolated manual efforts. For trace investigations, the system identifies relevant files from large repositories, converts binary traces into machine-readable form, validates data integrity, applies pattern matching, and produces human-readable diagnostics using a large language model (LLM). It also visualises error occurrences and system state transitions to pinpoint fault contexts. Beyond pattern matching, an LLM-based agent autonomously orchestrates investigations by querying signals, cross-referencing ticket context, and iterating towards root-cause hypotheses without predefined rules.

Impact: The approach improves efficiency and decision-making in SDV development. It streamlines ticket creation and analysis for testers and function owners and enables project engineers to perform deep-dive investigations without expert intervention. In production use, initial results indicate up to an 80% reduction in trace analysis effort per ticket, with improved consistency and accessibility for non-experts. Overall, AI-driven automation turns large-scale validation traces into faster root-cause hypotheses, quicker duplicate detection, and more robust ticket quality management, without relying on proprietary tools or data.

Keypoints:

- AI- and data-driven framework for automating ticket and trace analysis in SDV validation.
- Unified, scalable pipeline for ticket quality checks, duplicate detection, summarisation, and trace-centric investigation.
- End-to-end automated trace processing, including file selection, conversion, validation, and pattern-based analysis.
- LLM-supported summaries and visual analytics to accelerate root-cause hypothesis generation and improve accessibility for non-experts
- Measured operational impact, with improved consistency and up to 80% reduction in trace analysis effort per ticket.