www.exida.eu

14th *exida* Symposium

Functional Safety and Cybersecurity in the Automotive Industry

2026

October 27–28

(Starts with a get-together on October 26 in the evening)

Arabella Alpenhotel
Seeweg 7
83727 Spitzingsee

exida.com GmbH
Prof.-Messerschmitt-Str. 1
85579 Neubiberg

SYMPOSIUM 2026

The Automotive Industry is being transformed by software-defined vehicles, Artificial Intelligence, and increasing connectivity, bringing Functional Safety, Cybersecurity, and Free and Open-Source Software (FOSS) closer together than ever before. As AI is increasingly used in both vehicle functions and engineering processes, new challenges arise in safety assurance, verification and validation, cybersecurity, and lifecycle management. This symposium focuses on practical approaches for developing safe, secure, and trustworthy automotive systems. Topics include AI in safety-related applications, AI-assisted engineering, FOSS governance, the convergence of Functional Safety and Cybersecurity, statistical evaluation methods for SOTIF, and the latest developments in ISO 26262, ISO/SAE 21434, the Cyber Resilience Act. The symposium brings together experts from industry, academia, and standardization to exchange experience, discuss emerging challenges, and shape the future of Automotive Functional Safety and Cybersecurity.

Be part of the discussion, build your network, and share your insights with peers from across industries — all within a stunning natural setting that inspires innovation and reflection.
It will be our great pleasure to welcome you to the *exida* Functional Safety and Cybersecurity Symposium in the Automotive Industry 2026.

In case of questions or need for assistance with the registration, please contact us.

Kerstin Tietel
+49 89 44118232
kerstin.tietel@exida.com

www.exida.eu

Excellence in Dependable Automation

SYMPOSIUM 2026

FROM POSITIVE RISK BALANCE TO DRIVER ASSISTANCE: APPLYING SIFAD TO L2+

apl. Prof. Dr.-Ing. Moritz Werling & Ludwig Drees
BMW Group
Kilian Zwirglmaier
Qualcomm

Traditionally, SAE Level 2 systems have relied on actuation limiters to keep system failures controllable by the driver. However, as Level 2 functions expand into operational design domains with reduced clearances between traffic participants, particularly where hands-off operation is permitted, some failures can no longer be assumed to be fully controllable. In these cases, the probability of uncontrollable failures must be reduced to an acceptable level. SIFAD (Safety Integrity Framework for Automated Driving) provides a rigorous and efficient framework for estimating and quantifying residual risk, not only for SAE Level 3 systems but also for advanced L2 functions. The talk demonstrates the application of Bayesian networks through practical examples and shows how they can support a quantitative safety argument for driver assistance systems.

SYMPOSIUM 2026

FROM EXPLORATORY SIMULATION TO FIELD-RELEVANT SAFETY EVIDENCE: EXPLAINABLE STATISTICAL EVALUATION OF AN AEB FUNCTION

Prof. Dr. Stefan-Alexander Schneider
Hochschule Kempten

Rainer Faller
exida

The presentation introduces the application of Explainable Statistical Evaluation (ESE) to an Automatic Emergency Braking function in the Euro NCAP Car-to-Pedestrian Nearside Child Obstructed 50% scenario. In the first part, the FASim simulation framework is used as an evidence generator to systematically vary scenario parameters, execute Monte-Carlo experiments, derive safety-relevant performance measures, and provide traceable input for an ESE-based Bayesian network. The case study illustrates how simulation can complement a limited number of physical validation tests by identifying critical parameter regions, interactions, and potential performance limitations.

The second part presents methodological extensions intended to simplify the early application of ESE in industrial development. A key distinction is made between exploratory sensitivity analysis and quantitative field-risk estimation. For exploratory screening of the parameter and ODD space, real-world occurrence distributions are not necessarily required; transparent design distributions, such as uniform or beta distributions, can initially be used to reveal influential parameters and critical combinations. Real-world or plausibly justified ODD distributions are introduced subsequently to calibrate and weight the Bayesian network for field-relevant risk statements. The resulting staged approach supports the generation of statistically founded, explainable, and traceable evidence contributions for the safety argumentation of connected and automated vehicle functions.

SYMPOSIUM 2026

AI ORCHESTRATION FOR SAFETY ENGINEERING: TRANSFERRING LEGACY-BMS LESSONS TO AUTOMOTIVE E/E SYSTEMS

Dr. Timo Burggraf
Varta Storage GmbH

Large embedded codebases, incomplete safety boundaries, and shared resources challenge both legacy BMS and automotive E/E development. Drawing on the presenter's background in automotive functional safety, including EGAS and ISO 26262-based ABS/ESC development, this talk transfers lessons from a real legacy-BMS analysis to automotive practice.

It shows how an AI orchestrator, specialized tool-grounded agents, and independent QA can identify traceable potential findings in watchdog, CAN, SPI, memory, and shutdown paths, including missing FTTI evidence and common-cause dependencies. The automotive transfer covers safety/QM partitioning, freedom from interference, dependent-failure analysis, and bidirectional traceability. AI identifies and structures potential findings; engineering teams verify the evidence; the Safety Manager retains final authority.

SYMPOSIUM 2026

DELEGATION, TRUST, AND AI: REFLECTIONS FROM THE FUNCTIONAL SAFETY FRONTIER

Dr. Christopher Temple

The integration of Artificial Intelligence (AI) into automotive systems represents a transformative step for the industry, enabling perception, decision-making, and predictive maintenance at a scale previously unattainable. Yet this evolution introduces deep challenges for system architects and safety engineers: How do we justify trust in non-deterministic behavior? How do we assure systems that learn, adapt, or operate on opaque internal logic? And how can safety cases extend to components whose behavior may never be fully specified?

This presentation reflects on these questions through the lens of system integration, delegation of responsibility, and assumption management—core themes familiar from traditional functional safety. Drawing parallels to established supply chain practices, such as the use of safety elements out of context (SEooCs), we explore whether existing patterns of argumentation, partitioning, and validation are sufficient when applied to AI-based elements.

As AI becomes both a design tool and a runtime decision-maker, the session challenges attendees to think about where trust originates—and whether current assurance practices are prepared for what's coming.

SYMPOSIUM 2026

WOULD YOU SIGN IT? WHEN AI BECOMES THE SAFETY ENGINEER

Dr. Robert Maier
exida

AI assistants are becoming part of automotive safety development: whether our processes, standards, and organizations are ready or not. They promise to strengthen and accelerate engineering work, yet quietly challenge established assumptions about expertise, review, independence, and responsibility. As their contribution grows, the uncomfortable question is no longer whether we will use AI, but whether we are prepared to stand behind the resulting work products. This interactive session brings the elephant in the room into the open and invites you to examine where confidence ends and accountability begins. Through active participation and thought-provoking themes that cut to the heart of safety engineering, you will confront your own boundaries for using AI in safety-critical development. The session will challenge assumptions, surface tensions, and set the stage for the practical exchanges that follow.

SYMPOSIUM 2026

AI-MODERATED SAFETY ANALYSES – SCALING EXPERTISE WITHOUT COMPROMISING SAFETY

Stefan Bläsius
BMW Group
Dr. Michael Blum
exida

The increasing complexity of modern vehicle systems and the shortage of experienced safety experts challenge traditional safety engineering practices. AI technologies offer new possibilities to capture, reuse, and scale expert knowledge across projects and organizations.

This presentation demonstrates how AI can support the moderation of safety analyses by ensuring methodological rigor, improving documentation quality, and reducing preparation effort. The concept enables organizations to leverage accumulated safety knowledge more effectively while keeping accountability and decision-making firmly in human hands. The session highlights benefits, implementation approaches, and lessons learned from introducing AI into established safety engineering processes.

SYMPOSIUM 2026

STRENGTHENING THE LINUX OS FOR AI AGENT GOVERNANCE IN THE AGENTIC ERA

Michael Engel
Red Hat

As AI agents transition from cloud-hosted assistants to autonomous workloads running on embedded Linux systems—in vehicles, robots, and industrial controllers—a new governance challenge emerges that existing standards do not yet fully address. Unlike traditional software with deterministic control flow, agentic AI workloads autonomously invoke tools, access files, and issue commands based on probabilistic inference, combining the broad syscall surface of an interactive shell with the autonomous execution of a daemon. This paper proposes a governance extension to the Red Hat In-Vehicle Operating System (RHIVOS) that composes existing Linux kernel primitives, namespaces, and cgroups, into a layered enforcement architecture specifically designed for agentic workloads, requiring zero kernel modifications. The framework distinguishes four complementary layers: containment (kernel-enforced isolation), governance (policy-based evaluation of agent behavior), validation (architectural separation of plan creation from execution with formal verification before action), and functional safety (compliance with standards like ISO 26262 and IEC 61508, requiring independently certified safety controllers, hazard analysis, and formal verification). For safety-critical deployments, the governance and validation layers complement—but do not replace—the functional safety layer, providing defense-in-depth within the Linux application processor while preserving clear scope boundaries. Built on RHIVOS's ISO 26262 ASIL-B certification foundation, this capability demonstrates that the Linux kernel already contains the primitives necessary to govern autonomous AI agents—what is needed is a principled framework to compose them.

SYMPOSIUM 2026

OPEN DISCUSSION: AI TOOLS IN SAFETY-CRITICAL AUTOMOTIVE DEVELOPMENT – OPPORTUNITIES, CHALLENGES, AND PRACTICAL EXPERIENCE

Dr. Rafael Zalman
Infineon

Artificial Intelligence is increasingly finding its way into safety-critical automotive development, from requirements engineering and software implementation to verification, validation, and documentation. While AI tools offer significant potential to improve efficiency and quality, they also raise important questions regarding trustworthiness, compliance with standards, traceability, and functional safety.

This interactive discussion session invites participants to share practical experiences, discuss different approaches, and exchange opinions on the use of AI in safety-related development. The focus is on real-world applications, lessons learned, challenges, and best practices rather than formal presentations. The objective is to foster an open dialogue, identify common concerns, and explore how AI can be integrated responsibly into automotive safety engineering.

SYMPOSIUM 2026

DISCUSS THE CURRENT STATE OF OPEN SOURCE SOFTWARE CERTIFICATION METHODS

Fabrizio Tronci
Red Hat

Red Hat continues to develop a tailored certification methodology that is suitable for open source software. We will hold an open discussion about current approaches for mitigating identified misalignments with traditional safety-critical development methodology.

SYMPOSIUM 2026

DEMONSTRATING FREEDOM FROM INTERFERENCE (FFI) THROUGH XEN HYPERVISOR

(THE JOURNEY OF AN OPEN-SOURCE HYPERVISOR TO MEET THE REQUIREMENTS OF ISO 26262 AND ISO/SAE 21434)

Ayan Kumar Halder & Senthil Rajagopal
AMD

Software-defined vehicles are consolidating the digital cockpit onto single-SoC multicore platforms that combine the safety-critical instrument cluster and ADAS display, the infotainment system, and cloud connectivity. This cuts cost, weight, and wiring complexity, but sharing hardware between mixed-criticality domains makes freedom from interference (FFI) the central challenge: a fault in the QM infotainment domain must not corrupt the memory, timing, or communication of the ASIL-rated cluster. On a connected cockpit, however, FFI is not a safety problem alone—an infotainment or connectivity domain exposed to the outside world becomes an attack surface, and there is no safety if security is compromised. ISO/SAE 21434 therefore sits alongside ISO 26262 as a first-class concern: the isolation the safety case depends on must itself be resilient against malicious interference, not just accidental faults. This presentation shows how the open-source based Xen hypervisor meets that dual challenge, demonstrating FFI through core hypervisor isolation features—hardware-assisted memory partitioning, deterministic scheduling and CPU pinning, and controlled inter-domain communication—that keep the ASIL-B cluster fully separated from the QM infotainment and connectivity VMs under an ISO 26262-aligned safety lifecycle and an ISO/SAE 21434 cybersecurity process.

Since Xen is open source, unlike proprietary hypervisors, an open-source safe hypervisor gives OEMs and Tier-1s full visibility into the code underpinning their safety and security argument, freedom from vendor lock-in and per-unit licensing costs, and the flexibility to adapt, audit, and maintain the platform across long vehicle lifecycles. The talk argues that open source is not only a cost advantage but a strategic enabler of transparency, control, and long-term ownership for the digital cockpit.

SYMPOSIUM 2026

DRIVE35: UK INNOVATION SCHEME FOR MIXED-CRITICALITY SDV

Kaspar Matas & Gavin McCall
Codethink

DRIVE35 is a collaboration between Jaguar Land Rover (JLR), Arm, and Codethink for architecting and implementing the concept of a centralised compute module capable of hosting all vehicle software, regardless of criticality or function. This entails:

- Allocation of JLR safety application between main compute running the Codethink Trustable Reproducible Linux (CTRL OS) and safety domains in accordance with in-context platform constraints
- Early elimination of systematic faults by building the CTRL OS for the Arm Fixed Virtual Platforms (FVP), due to its virtualised functional parity with future silicon
- Engineering a continuous collaborative in-context safety case with the exida assessed Eclipse Trustable Software Framework, a methodology validated as suitable for software development up to ASIL D/SIL3
- Construction of the centralised compute module with continuous fault induction, stress and soak testing quality controls and evidence gathering via the Codethink Software Integration Factory (SIF)

This presentation presents how the combined innovative approaches from JLR, Arm, and Codethink are aiming to reduce the number of processing nodes in JLR's next generation vehicle programme by more than ~80%, to achieve a zonal architecture suitable for mixed-criticality SDV use cases.

SYMPOSIUM 2026

GENERAL PURPOSE FOSS IN SAFETY-RELATED SYSTEMS – QUALIFICATION & INTEGRATION STRATEGIES

Jan Toennemann
Vector Informatik

Free and open-source software (FOSS) constitutes the backbone of nearly every technical device and digital infrastructure we come in contact with. With some libraries being present in billions of devices worldwide, there should be no question on the robustness of the actual implementation — but does that mean they are suited to the use in safety-related systems?

Based on actual, widely employed FOSS projects, we are going to examine how qualification and integration strategies could look like. The explicit documentation of design assumptions highlights how traditional desktop, server, and IoT use cases can show significant gaps when compared to the requirements of safety-related systems. With upstream projects rarely interested in merging hardening activities targeted at safety-related systems, e.g., due to decline in performance, architectural changes typically require some divergence from the original project.

Let's examine the technical challenges and business cases of FOSS qualification and integration in safety-related systems compared to targeted (re)engineering of required functionality.

SYMPOSIUM 2026

OPEN SOURCE BEST PRACTICES AS A FOUNDATION FOR QUALITY MANAGEMENT

Philipp Ahmann
ETAS GmbH

Modern software development has become highly iterative and code-centric, moving beyond traditional waterfall and V-model approaches. Nevertheless, open source software forms the foundation of many security- and safety-critical systems and has repeatedly demonstrated its ability to deliver high-quality software at scale. This success is enabled by practices such as peer review, automated testing, continuous integration, traceable and transparent contribution processes, and collaborative governance.

This session explores what research and industry experience reveal about effective open source development practices and how these practices relate to established quality management methodologies, such as ISO 9001. It presents examples of open source projects, including Eclipse S-CORE, that have adopted development processes aligned with industry and regulatory standards, including ISO 26262, ISO/SAE 21434, and Automotive SPICE.

The presentation also examines project health and maturity assessment frameworks, including LFX Insights and OpenSSF Scorecards, developed by open source foundations, communities, and ecosystem initiatives dedicated to measuring project health, sustainability, and maturity. By mapping proven open source practices to quality management principles, the session highlights how organizations can leverage open development models while maintaining the rigor required for safety- and quality-critical applications.

The session is informed by ongoing work within the ELISA (Enabling Linux in Safety Applications) Project, where open source best practices, project maturity models, and quality management principles are currently being analyzed and mapped to the needs of safety-critical domains.

SYMPOSIUM 2026

SHORT RANGE WIRELESS TECHNOLOGY SPARKLINK TARGETING CYBER SECURITY AND FUNCTIONAL SAFETY FEATURES

Dr.-Ing. Francesc Fons & François Fischer
Huawei

In this talk, the authors will introduce the short-range wireless communications technology SparkLink as well as the International SparkLink Wireless Short-Range Communications Alliance (iSLA, in short). Founded in 2020, the iSLA community is committed to a globalized industry alliance with the goal of promoting the innovation and industrial ecosystem of new wireless short-range communication technologies, namely SparkLink (a.k.a. NearLink), to support the fast development of new applications such as smart vehicle, smart homes, smart terminals and smart manufacturing.

After a first introduction of the technology and its alliance, the authors will walk through a practical use case. To be exact, they will showcase the adoption of this technology in the automotive industry, detailing the roll out of SparkLink in commercial cars through the digital car key function as the first proven success case. They will deep dive in technical aspects concerning functional safety and cyber security features and will present a brief overview of the Intelligent Car Connectivity Industry Ecosystem Alliance (ICCE), comparing it with the current regulation on digital key (DK) from UNECE.

SYMPOSIUM 2026

AGENDA

SECURITY ANALYSIS OF A CONFIGURABLE INTERCONNECT IP IN CONTEXT OF ITS USAGE IN AN ISO/SAE 21434 CONTEXT

Stefano Lorenzini
Arteris
Frauke Blossey & Clemens Roettgermann
exida

The application of ISO/SAE 21434 to semiconductor IP components - especially highly configurable, lower-tier components - poses a structural challenge. While the standard is oriented toward vehicle-level systems, component suppliers (Tier-2/3, IP providers) are still required to contribute evidence toward vehicle cybersecurity arguments. exida and Arteris are going to present an approach, which has been applied for the Arteris FlexGen IP interconnect.

SYMPOSIUM 2026

SYMPOSIUM LOCATION

Arabella Alpenhotel
Seeweg 7
83727 Spitzingsee, Germany

www.exida.eu

Excellence in Dependable Automation

Registration Form ONSITE

I register for the:
***exida* Automotive Symposium 2026**

Date: October 26 - 28, 2026

Location: Arabella Alpenhotel am Spitzingsee
Seeweg 7
83727 Schliersee-Spitzingsee
Germany
www.arabella-alpenhotel.com

Price: € 1,895. -- + tax
The price includes accommodation, food and beverages.[1]

For registration until August 31st, 2026, we will grant a special discount of 25% (€ 1,421.25 + tax).

Please enter the participant's billing address:

Company:* _____

VAT-ID:* _____

Participant Name:* _____

Department: _____

Street*: _____

Post code, city, country:* _____

Participant Email:* _____

Phone number:* _____

Fields marked with an asterisk (*) are required.

Please send the filled page via email to kerstin.tietel@exida.com.

Booking conditions: The symposium will be held in English, and the presentation slides will be in English. In case the registered participant sends a written cancellation 50 days before the start of the symposium the cancellation will be free of charge. Until 21 days before the start of the symposium a cancellation fee of 50% of the fee will be charged. For later cancellations done by registered participants the complete symposium costs will be charged. A replacement of the registered participant with another person is possible at any time. The acceptance of the conditions is part of the registration. *exida.com* GmbH reserves the right to cancel the symposium at short notice and in writing. In this case only the symposium fees will be refunded.

Data protection: The collected personal data is only stored and used for internal purposes related to the management of the training. This data is protected by limited access rights. The duration of the archiving depends on the legal requirements.

_____
Date Signature

[1] Meals or beverages consumed outside of the planned dining will be billed separately on your own expenses.

Registration Form ONLINE

I register for the:
***exida* Automotive Symposium 2026**

Date: October 27 and 28, 2026

Location: Online

Price: € 990. -- + tax

For registration until August 31st, 2026, we will grant a special discount of 25% (€ 742.50 + tax).

Please enter the participant's billing address:

Company:* _____

VAT-ID:* _____

Participant Name:* _____

Department: _____

Street:* _____

Post code, city, country:* _____

Participant Email:* _____

Phone number:* _____

Fields marked with an asterisk (*) are required.

Please send the filled page via email to kerstin.tietel@exida.com.

Booking conditions: The symposium will be held in English, and the presentation slides will be in English. In case the registered participant sends a written cancellation 14 days before the start of the symposium the cancellation will be free of charge. Until 7 days before the start of the symposium a cancellation fee of 50% of the fee will be charged. A replacement of the registered participant with another person is possible at any time. The acceptance of the conditions is part of the registration. *exida.com* GmbH reserves the right to cancel the symposium at short notice and in writing. In this case only the symposium fees will be refunded.

Data protection: The collected personal data is only stored and used for internal purposes related to the management of the training. This data is protected by limited access rights. The duration of the archiving depends on the legal requirements.

_____
Date Signature