The Past Shapes the Future

89th Meeting of IFIP WG 10.4

May 4-7, 2026 – Kaunas, Lithuania

The 89th meeting of IFIP Working Group 10.4 was held at Vytautas Magnus University in Kaunas, Lithuania, under the theme "*The Past Shapes the Future*". This event brought together pioneers of dependable computing and researchers addressing new challenges. The opening session honored Al Avižienis, Jean-Claude Laprie, Hermann Kopetz, Brian Randell, John Meyer and Bill Carter for their foundational contributions, while the technical sessions used their legacy to reflect on the field's current state and future direction.

Artificial Intelligence (AI) is now the central challenge and cannot be ignored: AI is emerging not as a peripheral topic but as the defining challenge reshaping dependable computing. The consensus is clear: agentic AI systems are already being deployed in critical domains, and the risks of disengagement outweigh those of engagement.

AI systems, particularly LLMs and agentic architectures, introduce failure modes that fall outside the scope of traditional dependability frameworks. Their stochastic behavior, sensitivity to distributional shifts, and opaque alignment between intended and actual behavior complicate assurance efforts. Iterative agentic loops can degrade performance, sequential tool chaining can transform benign actions into harmful ones, and accuracy-focused evaluation is insufficient for assessing consistency, robustness, predictability and safety.

Rather than abandoning established approaches, core dependability principles should be adapted for AI systems, including fault tolerance, redundancy, formal specification and assurance argumentation. For safety-critical domains such as aviation, automotive and infrastructure, AI should be introduced with the same rigor applied to earlier disruptive technologies through independent assessment, staged validation and assurance levels proportional to the consequences of failure.

Understandability must become a first-class property: Systems are increasingly deployed without anyone, including developers, operators, or regulators, fully understanding their behavior. Although this is not a new issue, AI-generated code is exacerbating the problem by enabling the creation of systems that no human has formally or conceptually reasoned about. For this reason, understandability must become a first-class research and engineering objective rather than a secondary concern, in particular, when considering code reuse. Systems that clearly express their designers' intent are easier to reason about, verify, and audit when failures occur. This perspective favors architectures that preserve decomposability and auditability over monolithic end-to-end designs, even when the latter offer short-term performance benefits. The same principle extends to assurance cases, which must evolve from static documents into living frameworks that remain valid as systems change.

Hardware is no longer a dependable foundation: Hardware can no longer be assumed to provide a fully reliable foundation for software dependability. At modern transistor scales, crosstalk, environmental noise, and side channel vulnerabilities are inherent properties of hardware rather than defects that can simply be engineered away. Increasingly complex predictive execution mechanisms and opaque supply chains for trusted execution environments further weaken the assumption of a secure and dependable hardware boundary. As a result, dependability must be reconsidered from the ground up, treating hardware as a distributed system in its own right. Chiplet-based architectures, which organize processors as networks of small, independently programmable functional units, were identified as a promising approach for improving fault tolerance and physical isolation. An important open question is the level at which fault and intrusion tolerance should be implemented, whether at the node, chip, chiplet, or functional unit level.

The dependable computing community must engage with policy and legislation: The technical knowledge needed to build more dependable systems already largely exists. What is missing is a set of regulatory and market incentives to apply it. Two well-documented cases illustrate that even catastrophic, high-profile failures have not produced effective legislative responses: forensic software whose false positive error rate was claimed in court to be one in a million but found through principled analysis to be closer to one in eight, and the British Post Office Horizon scandal in which hundreds of innocent people were wrongly prosecuted on the basis of faulty software. This is a structural failure that the dependability field must address directly. Software vendors face little accountability, programmers are generally uncertified, and courts often lack the expertise to rigorously evaluate technical evidence. Greater engagement with policymakers, stronger standards for software testing and validation, and regulatory frameworks that create meaningful incentives for dependability are needed. Experience from building codes, automotive safety, and aviation certification suggests that lasting change requires both incentives and penalties, often triggered by a major failure.

Intelligent vehicles define a near-term research agenda: Self-driving vehicles are entering a scale-up phase, but participants agreed that they have not yet demonstrated reliable autonomous operation, even within constrained operational domains. The gap between claimed and actual autonomy remains significant. Several opportunities exist for IFIP WG 10.4 to make important contributions during the 2026–2030 period, including defining quantitative acceptance criteria for AI and machine learning systems, improving the collection of field data on incidents and tele-assistance interventions, developing comprehensive system assurance cases, and recommending technologies that simplify assurance activities.

Looking toward 2030–2040, attention will shift from individual autonomous vehicles to Systems of Autonomous Systems that integrate self-driving cars, autonomous transport and service vehicles, unmanned aerial vehicles, and infrastructure machines. Ensuring the dependability of these interconnected systems will require new approaches to acceptance criteria, assurance case construction, and system architecture, creating a substantial long-term research agenda.