

0.1 Executive Summary

Who Builds the Architects: From Drawing Board to Reasoning Engine

archi-intelligence Research Series · Working Paper 2026-04

DOI: 10.5281/zenodo.21035221

Taking the design tool—an engineering artifact spanning six decades—as its starting point, this paper traces the principal line of evolution from the constraint solver in Sutherland’s 1963 Sketchpad to the contemporary toolchains of MBSE, EDA, and AUTOSAR. It advances a proposition that vendor market narratives routinely obscure: the deepest identity of a design tool is constraint maintenance, not graphical depiction. On this basis, the paper repurposes the AI²-ML (Architecture Intelligence Maturity Levels) framework established across the trilogy—turning it from a measure of the architecture being designed to a measure of the tool that designs it. It scores the core species of the automotive architecture toolchain along five dimensions (D1–D5: constraint-awareness depth, reference-architecture integration, AI-assisted generation, AI reasoning autonomy, and interoperability), and composes these into an overall L0–L5 level under an aggregation rule that takes D1 and D4 as its twin reasoning axes. Every score rests on standards specifications and official vendor documentation as first-tier evidence, within a 2024.1–2026.1 evaluation window.

The central finding is that the category hugs the ground while its summit stands vacant. Of twelve tools, only one—SCADE, now part of Synopsys—reaches L3, and even there its verifiability extends only to the software and model layer; the remaining eleven fall within L1–L2, and not one reaches the whole-system L4. Two further observations follow. The “heat of AI” and the “capacity for reasoning” have decoupled: the AI capabilities announced across 2025–2026 land almost entirely on generation (D3), and lift no tool’s reasoning autonomy (D4). And SysML v2, though finalized by the OMG in 2025, remains thin in practice; absent a conformance test suite, every high interoperability score rests on vendor self-attestation.

On this basis the paper argues that the ultimate form of an architecture-design tool is not autonomous generation but **verifiable constraint satisfaction—rendering provable the correctness of a design and of its every change**. It defines the summit of AI²-ML (L5) as “no unverified change,” establishes constraint-awareness (D1) and reasoning autonomy (D4) as the core dual axes, and shows that the vacant L4–L5 highland—the empty frontier—is the same blank that the industry’s architecture debt, identified in the trilogy, marks out from the supply side. AI²-ML is a capability-maturity ladder; it deliberately excludes market share, and secures through tiered, observable evidence that every score can be independently recomputed.

Keywords: architecture-design tools; constraint-aware reasoning; change-impact analysis; Correct-by-Construction; MBSE; SysML v2; architecture maturity; AI²-ML