

The V-Model, recompiled: Q&A with Aptiv

11-May-2026 15:06 GMT

Matthew Beecham

S&P Global

Supply Chain and Technology, Automotive

How the V-Model has evolved from sequential validation to continuous software-driven verification.

In automotive development, the V-Model has long served as the backbone of the engineering discipline, mapping requirements to verification in a mirrored structure that ensures traceability from concept to production. For decades, it provided the industry with a predictable rhythm: define, design, implement, test — then validate back up the chain.



Source: Getty Images/Jackie Niam

But the model was born in an era of hardware-dominated vehicles, when software changes were infrequent and tightly controlled. Each electronic control unit was developed, validated and frozen largely independently, with integration occurring late in the program cycle. The rise of advanced driver assistance systems began to strain this cadence, introducing software complexity that outpaced traditional gate-based validation cycles.

Today, the shift to software-defined vehicles, centralized compute and continuous deployment is reshaping how the V-Model is applied rather than replacing it outright. Verification is increasingly continuous, simulation-heavy, and distributed across development pipelines rather than concentrated at the end of the program. Standards such as ISO 26262, UNECE R155 and R156, and modern DevSecOps practices have effectively extended the model into the life cycle of the vehicle itself.

The result is not the abandonment of the V-Model, but its transformation from a sequential process into a persistent, continuously executing discipline.

To understand how this shift is being implemented in practice — and what it means for safety, architecture and development workflows — we spoke to Brian Witten, chief technology officer for Aptiv's Intelligent Systems, Sensors & Compute.



The following is an edited transcript of the conversation.

S&P Global Mobility: The V-Model has long underpinned automotive development. Is it still fit for purpose in a world of centralized compute and continuously evolving software-defined vehicles?

Brian Witten: The V-Model still works. What has changed is the pace at which it runs.

ISO 26262 continues to require requirements traceability and verification evidence. Centralized compute does not change those obligations; instead, it changes the assumption that the V is traversed once, gate to gate, before a program ships. At Aptiv, the process is now continuous. Every

commit triggers a verification cycle tied back to requirements, running across software-in-the-loop, hardware-in-the-loop, and vehicle-in-the-loop environments while generating evidence throughout the process. The Aptiv LINC™ Software Platform, together with the underlying Wind River development environment, automates much of this traceability.

Centralized compute can actually make the V-Model easier to execute, not harder. Validation is performed against a stable platform abstraction rather than dozens of ECU variants sourced from multiple suppliers.

Aptiv is heavily invested in zonal architectures and SDV platforms — how is this architectural shift forcing changes to traditional verification and validation workflows?

Validation now happens at the platform level, not just the function level. The Gen 6 ADAS platform decouples hardware and software, supports multiple [system on chip] and architectural configurations, and runs OCI-compliant containers so software can move to the most suitable device. That flexibility only works when verification evidence travels with the components — a fundamentally different model from the one most validation organizations were originally designed around.

Zonal architectures also introduce new validation requirements. Signal integrity and timing across the network must now be validated in ways that traditional point-to-point ECU architectures never required. At the same time, cloud-native infrastructure is changing what is possible at scale. In Aptiv's own program data, modular CI/CD pipelines reduced an Android OS integration cycle from months to just 48 hours.

CI/CD pipelines are increasingly central to software development in other industries. Can they be reconciled with ISO 26262 expectations in safety-critical automotive programs at scale?

Yes. Aptiv is already doing this today. ISO 26262 prescribes rigor and traceability, not cadence. Treating CI/CD as a way to bypass verification, rather than automate and strengthen it, reflects a process failure — not a limitation of the standard itself.

In Aptiv's pipelines, every run executes static analysis, fault injection, requirements-linked test cases, and ASIL-appropriate coverage gates before any merge is allowed. Git, GitHub, Coverity, Wind River Studio and other tools are integrated directly into the build process, while hardware-in-the-loop and software-in-the-loop testing are orchestrated automatically. Work-product evidence is generated as a byproduct of the pipeline itself, rather than through a separate, manual documentation effort.

The more useful question may actually be the reverse: maintaining ISO 26262 evidence manually at software-defined vehicle (SDV) development speed is what becomes impossible. Automation is what makes the standard scalable and practical.

Simulation and virtual validation are gaining momentum across ADAS and automated driving development. At what point — if any — can they meaningfully reduce reliance on physical testing for homologation?

For ADAS perception and prediction, the industry is already operating in a simulation-first environment. The long tail of real-world edge cases is simply too vast to validate through physical driving alone, which means a significant share of verification must happen in simulation. Aptiv already applies simulation across hazard analysis, risk assessment, fault analysis, and high-

frequency design work, supported by tools such as Ansys medini analyze within the functional safety workflow.

Homologation, however, is a separate issue. UNECE R157 already allows software-in-the-loop (SIL) and hardware-in-the-loop (HIL) evidence for automated lane keeping systems, and regulators are gradually moving further in that direction. Even so, the highest-assurance functions still require physical correlation for type approval, and that is likely to remain the case for the foreseeable future.

The industry is ultimately moving toward a multipillar validation model. Simulation provides broad scenario coverage and exercises edge cases at scale. Physical testing establishes real-world correlation. Field data then closes the loop after deployment. What is changing is the balance between those pillars: validation evidence increasingly spans all three, rather than being concentrated in a single physical test campaign at the far right side of the V-Model.

As original equipment manufacturers push to own more of the software stack, including middleware and operating systems, where does this leave tier 1 suppliers like Aptiv in the long-term architecture of the vehicle?

Aptiv's strategy has been closely aligned with the major megatrends reshaping mission-critical industries, including automation, electrification and digitalization.

The Gen 6 ADAS platform is available either as a turnkey solution or as a set of modular building blocks. It incorporates standardized middleware and OCI-compliant containers, enabling OEMs to deploy and scale their own software on top. The LINC™ Software Platform abstracts applications from the underlying SoCs and operating systems, allowing the same codebase to run across multiple hardware platforms with minimal rework.

Aptiv's value extends across the full technology stack — from high-performance compute, safety-critical real-time operating systems and zone controllers to sensors and development infrastructure. Delivering vertically integrated ASIL-D real-time systems, while also assuming global safety responsibility, requires the depth of experience, technical insight and operational focus that Aptiv brings.

As the platform layer consolidates around suppliers such as Aptiv, the feature layer is increasingly shifting toward the OEM.

Software reuse is essential to scaling SDV platforms but also introduces systemic risk. How can suppliers ensure safety integrity is preserved when code is reused across multiple platforms and hardware generations?

We think about this in three dimensions.

First, there is rigorous component qualification. Each reusable element is treated as a safety element out of context (SEooC), with explicit ASIL decomposition and clearly documented integration assumptions. The LINC™ Software Platform conforms to ASPICE, ASIL-B and OCI standards, while providing the deterministic environment required to satisfy ISO 26262 requirements.

Second, hardware abstraction layers and standardized APIs help contain the impact of silicon changes. This reduces the amount of redesign and rework required when hardware platforms evolve.

Third, every component still undergoes full re-validation when deployed on new hardware. Gen 6 is designed so the stack can be re-verified against a new SoC variant without rebuilding the entire safety case from scratch.

Supply-chain integrity is another critical dimension, and one that is often underestimated. Reused code creates leverage not only for manufacturers, but also for attackers. Every binary in the platform therefore requires cryptographically signed provenance, a software bill of materials (SBOM), and, increasingly, quantum-resistant signatures. Without those controls, code deployed across a million vehicles effectively becomes a million-vehicle attack surface.

Finally, V-Model evidence must travel with the component wherever it is deployed. If the verification and traceability evidence does not accompany the software across platforms, the entire argument for safe reuse begins to break down.

Do zonal architectures ultimately simplify integration and validation by reducing ECU complexity, or do they introduce new system-level dependencies that make verification harder?

It is both — but ultimately a net positive when implemented correctly. Zonal architectures separate I/O from compute, consolidate large numbers of distributed ECUs, and significantly reduce wiring mass and physical complexity. As a result, manufacturing, end-of-line testing, and repairability all become more efficient.

At the same time, zonal architectures introduce a new set of very concrete engineering and validation requirements. Deterministic timing across the in-vehicle Ethernet backbone becomes critical, as does managing resource contention on centralized compute platforms running mixed-criticality workloads. Architecturally, it is equally important to ensure that a failure within one part of a zone controller cannot propagate into unrelated vehicle functions. Addressing these concerns adds new verification demands on top of the requirements for network determinism and freedom from interference.

This is ultimately the same trade-off the broader computing industry made decades ago: moving from many small, independently programmed components to a smaller number of highly concurrent processors that provide a far more flexible compute fabric. The production and scalability benefits are substantial but fully realizing them requires restructuring and accelerating the V-Model to support continuous integration, continuous testing and ongoing verification.

The tooling landscape is becoming increasingly fragmented, particularly with the rise of AI-driven development and simulation tools. Is the industry moving toward consolidation or further divergence?

More divergence is expected in the near term, followed by gradual consolidation around standards rather than individual vendors. The industry is still in an early phase of AI-assisted development, scenario generation, neural simulation and synthetic data pipelines. Many of these tools are genuinely effective, but a significant portion will likely not persist into production use by the next major architecture cycle. That pattern is typical of a rapidly expanding tool category.

When consolidation does occur, it is expected to happen at the standards layer. Frameworks such as ASAM OpenX, OSI, OCI containers and AUTOSAR Adaptive are what make a heterogeneous toolchain manageable in production environments. This assumption is embedded in Aptiv's development infrastructure: the LINC™ Software Platform conforms to ASPICE, ASIL-B and OCI standards, while the underlying Wind River environment is OS-agnostic and designed to be

extensible to third-party tools by default.

Standards are also critical from a qualification standpoint. Under ISO 26262, tool qualification does not scale effectively if the toolchain surface is constantly changing beneath active programs. A stable and standardized interface layer is what enables V-Model evidence to remain valid across tool generations and toolchain evolution.

Over-the-air (OTA) updates effectively extend validation into the post-production phase. Does this mark the end of the traditional V-Model as a pre-deployment framework?

It marks the end of the V-Model as a one-pass exercise, not the end of the V-Model as a discipline. The V now operates across two timeframes. A pre-SOP V establishes the baseline safety case before the vehicle is released, while a per-update V validates each OTA delta against that baseline before it reaches a customer vehicle. The same requirement gates apply: fault injection, regression testing, system integration and vehicle-in-the-loop validation. The same evidence trail is maintained throughout. DevSecOps pipelines that produced the original software are re-executed for every update, with end-to-end traceability preserved and a continuous feedback loop extending back into the cloud.

The regulatory environment has already evolved in this direction. UNECE R155 and UNECE R156 require a cybersecurity management system and a software update management system that span the full vehicle life cycle. Aptiv operates according to these principles regardless of jurisdiction, as they represent the appropriate model for managing a connected, safety-critical system.

In this sense, the V-Model is not disappearing; it is extending. It becomes longer in duration, continuously active across the life cycle of the vehicle, with a closed-loop connection back to the cloud. These are not departures from the model, but rather structural improvements that adapt it to software-defined vehicles.

CONTACTS

The Americas

+1 877 863 1306

Europe, Middle East & Africa

+44 20 7176 1234

Asia-Pacific

+852 2533 3565

www.spglobal.com/mobility

Copyright © 2025 S&P Global Inc. All rights reserved.

These materials, including any software, data, processing technology, index data, ratings, credit-related analysis, research, model, software or other application or output described herein, or any part thereof (collectively the “Property”) constitute the proprietary and confidential information of S&P Global Inc its affiliates (each and together “S&P Global”) and/or its third party provider licensors. S&P Global on behalf of itself and its third-party licensors reserves all rights in and to the Property. These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable.

Any copying, reproduction, reverse-engineering, modification, distribution, transmission or disclosure of the Property, in any form or by any means, is strictly prohibited without the prior written consent of S&P Global. The Property shall not be used for any unauthorized or unlawful purposes. S&P Global’s opinions, statements, estimates, projections, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security, and there is no obligation on S&P Global to update the foregoing or any other element of the Property. S&P Global may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. The Property and its composition and content are subject to change without notice.

THE PROPERTY IS PROVIDED ON AN “AS IS” BASIS. NEITHER S&P GLOBAL NOR ANY THIRD PARTY PROVIDERS (TOGETHER, “S&P GLOBAL PARTIES”) MAKE ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE PROPERTY’S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE PROPERTY WILL OPERATE IN ANY SOFTWARE OR HARDWARE CONFIGURATION, NOR ANY WARRANTIES, EXPRESS OR IMPLIED, AS TO ITS ACCURACY, AVAILABILITY, COMPLETENESS OR TIMELINESS, OR TO THE RESULTS TO BE OBTAINED FROM THE USE OF THE PROPERTY. S&P GLOBAL PARTIES SHALL NOT IN ANY WAY BE LIABLE TO ANY RECIPIENT FOR ANY INACCURACIES, ERRORS OR OMISSIONS REGARDLESS OF THE CAUSE. Without limiting the foregoing, S&P Global Parties shall have no liability whatsoever to any recipient, whether in contract, in tort (including negligence), under warranty, under statute or otherwise, in respect of any loss or damage suffered by any recipient as a result of or in connection with the Property, or any course of action determined, by it or any third party, whether or not based on or relating to the Property. In no event shall S&P Global be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees or losses (including without limitation lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Property even if advised of the possibility of such damages. The Property should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions.

The S&P Global logo is a registered trademark of S&P Global, and the trademarks of S&P Global used within this document or materials are protected by international laws. Any other names may be trademarks of their respective owners.

The inclusion of a link to an external website by S&P Global should not be understood to be an endorsement of that website or the website’s owners (or their products/services). S&P Global is not responsible for either the content or output of external websites. S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process. S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global Ratings’ public ratings and analyses are made available on its sites, www.spglobal.com/ratings (free of charge) and www.capitaliq.com (subscription), and may be distributed through other means, including via S&P Global publications and third party redistributors.