Vehicle software: who can approve the next change?

Software ownership changes decision rights, but does not remove review or integration. This report examines sourcing, requirements exchange and after-sale activity as separate measures. Public evidence supports a coordination mechanism; it does not measure the causal effect of hand-off counts on release speed.

41% vs 25%

Carmakers building their software stack mainly in-house, China against the US (AlixPartners survey, late 2025)

3 vs 5 years

Time to bring a new vehicle to market, Chinese against established carmakers (AlixPartners 2026 outlook, as reported by Forbes)

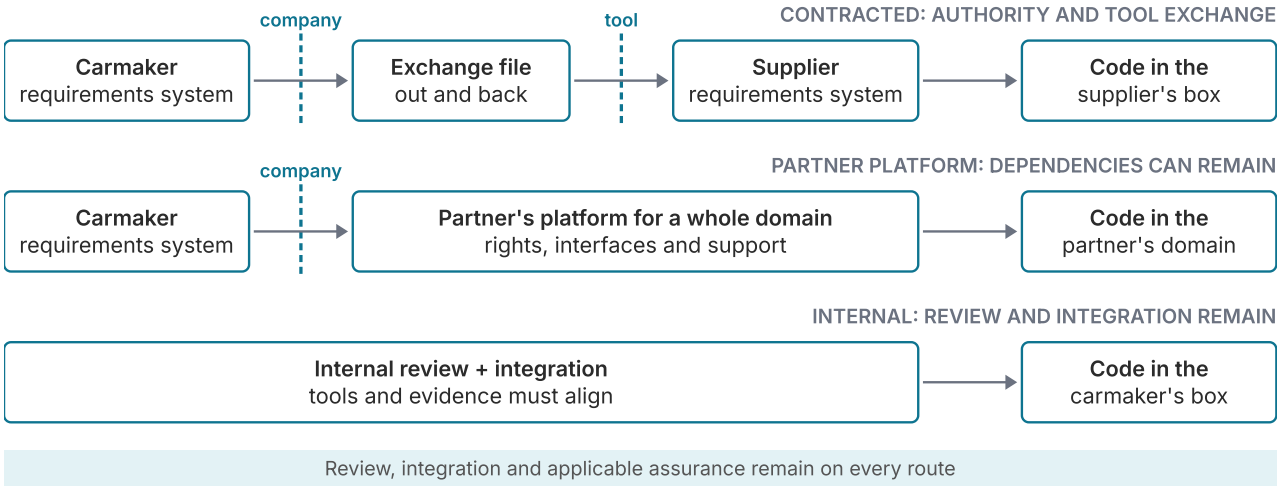
25.7%

Share of China's annual recall vehicle count covered by OTA recall campaigns in 2025 (SAMR); not completed repairs

SUMMARY

Decision rights, traceability and assurance shape the path of a change.

Illustrative responsibility arrangements; no measured border counts or time effects. NRG analysis.



Source: NRG analysis; prostep ivip (2024); Object Management Group; AlixPartners (2026); Forbes (2026); McKinsey (2026); China Association of Automobile Manufacturers; SAMR; Roland Berger (2024)

AlixPartners' late-2025 survey reports sourcing and deployment practices, not measured hand-off counts. Its organisational responses cannot establish how much ownership explains a development-speed gap.

NRG's framework separates commercial authority, exchange interoperability and release assurance. An external interface can require agreement; an internal team can still have tool differences and review queues. These are possible mechanisms, not quantified effects.

Among automaker respondents in the survey's regional groups, 41% in China primarily build SDV software in-house, compared with 25% in the US and 27% in EU/UK. These are not vehicle, controller or code shares.

ReqIF transfers structured XML data. Automation can preserve identifiers across files, while a shared repository can still contain inconsistent revisions. The useful test is whether meaning and evidence survive the exchange.

Architecture forecasts, recall coverage and city-NOA supplier categories describe different populations. None is a release-speed score. Bringing work in-house moves responsibility; internal review, integration tests and applicable regulatory obligations remain.

Whole-vehicle development and software deployment are different clocks.

The three-versus-five-year comparison remains a secondary report of a 2026 outlook. The survey's deployment question concerns the respondent organisation, not necessarily software reaching customer vehicles.

Whole-vehicle years. The AI bar is a hypothetical maximum saving from the 2025 outlook, not an observed cycle.



Source: AlixPartners Global Automotive Outlook (2026), as reported by Forbes

Forbes, reporting AlixPartners' 2026 outlook, puts the gap at two years: about three years to bring a new vehicle to market for Chinese carmakers, five for established ones, which makes the Chinese cycle roughly 40% shorter (NRG calculation). Both figures cover the whole vehicle, not its software alone, but they set the clock the software has to fit into.

AlixPartners' 2025 outlook identified potential AI-assisted savings of up to eight months from a typical five-year product-development cycle. If that maximum were achieved, the result would be about 4.3 years (NRG calculation). It is not an observed average or a software-release benchmark. Its separate recommendation to decouple hardware and software clocks describes a development approach, not a measured treatment effect.

In the November–December 2025 study, 1,002 executives participated, including 495 automaker respondents. The automaker deployment item asks for time from code complete to deployment of a new SDV application at the organisation. It does not identify the timing of rollout to a customer fleet. Bug-fix shares are a separate question.

NRG does not combine the whole-vehicle benchmark and deployment item into a speed ranking. The sources do not match individual sourcing choices, interface counts or review queues with observed release times. They motivate a process investigation; they cannot allocate the regional gap to hand-offs.

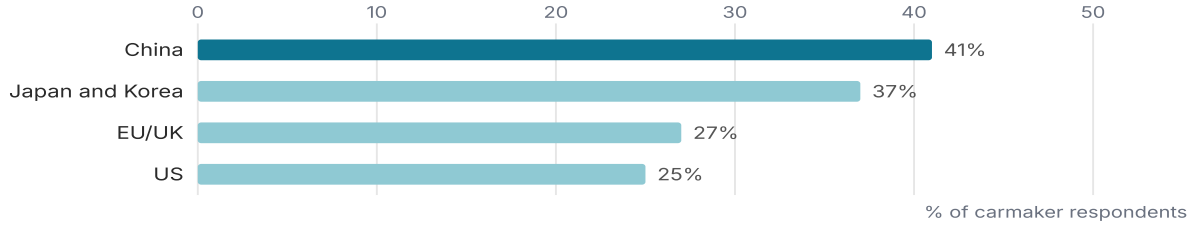
> **67%**
> of automaker respondents report more than eight weeks from code complete to deployment of a new SDV application at their organisation (late-2025 survey).

Sources: AlixPartners (2025, 2026); Forbes (2026); NRG calculation

Sourcing responses are not vehicle ownership shares.

The late-2025 AlixPartners survey included 1,002 executives, 495 at automakers. Regional percentages describe automaker respondents within each group. The chart does not give regional sample sizes or measure code ownership.

Automaker respondents primarily building SDV software in-house; US, EU/UK, China, Japan/Korea groups



Source: AlixPartners survey (November–December 2025, published April 2026)

Most respondents combine building and buying: 48–68% across the four regional groups charted. The survey categories describe sourcing strategies, not pure ownership models. A company can write one layer and license another, own hardware while buying software, or retain a supplier's platform while controlling selected interfaces. The percentages therefore cannot count the commercial decisions on a particular release path.

The survey authors associate in-house sourcing with architecture authority. Separately, 39% of Chinese automaker respondents reported a central/zonal architecture strategy, while 60–67% of Western respondents chose hybrid/in-transition designs. Strategy responses do not establish deployed fleets. Comparable samples and release scope matter.

Capgemini's separate June 2025 fieldwork records organisational strategy and implementation responses. It is not a matched sample of the AlixPartners respondents and cannot turn the regional sourcing comparison into a controlled test of development performance. Survey constructs and populations need to remain separate.

NRG treats architecture authority as decision rights: which interfaces may change without a new agreement, and who can test the result? A high in-house share does not answer either question alone. Internal teams and supplier licences can retain dependencies. A bought component can offer stable interfaces, reuse and explicitly agreed long-term update support with little recurring negotiation.

Other measures among regional automaker respondents

Region	Platform-level software reuse	Over half of R&D on software-defined vehicles
China	48%	36%
US	31%	21%
EU/UK	33%	19%
Japan and Korea	33%	41%

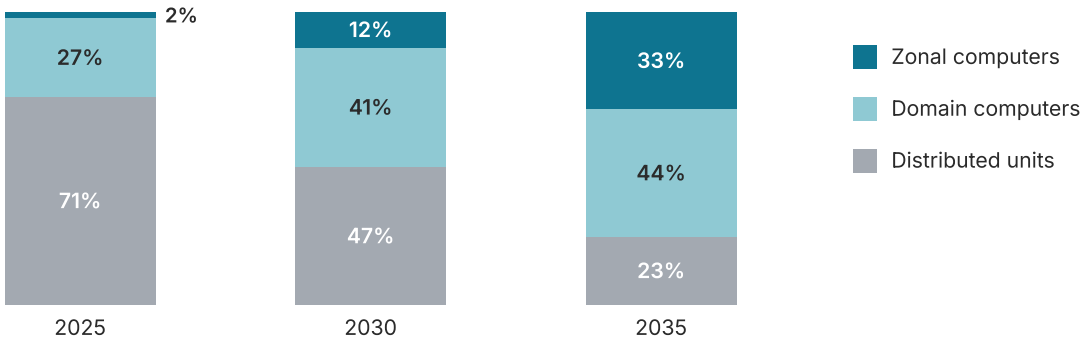
Source: AlixPartners (2026)

Sources: AlixPartners (2026); Capgemini Research Institute (2025)

Distributed designs fall in a model of light-vehicle production.

McKinsey's January 2026 Current Trajectory model covers passenger cars and light commercial vehicles. The 2025 architecture shares are a model base; 2030 and 2035 are conditional forecasts, not 2026 installed-fleet observations.

Light-vehicle production by architecture, %. Distributed is a broad design category, not literally one ECU for every function.



Source: McKinsey (2026)

McKinsey puts distributed architectures at 71% of light-vehicle production in its 2025 model base. In Current Trajectory, the share falls to 47% in 2030 and 23% in 2035; domain architectures reach 44% and zonal designs 33% in 2035. The categories are production designs, not the installed fleet. A distributed design need not assign exactly one controller to every function.

McKinsey expects the control-unit market to grow 4.0% a year to 2035, against 10.2% for software development and 11.9% for software integration, verification and validation services, the fastest-growing line in its forecast.

Integration includes more than checking one supplier against another. Central and zonal computers combine functions, require assurance across domains and can retain externally supplied code. NRG treats consolidation as a change in where integration takes place, not as evidence that internal ownership reduces its net cost.

Architecture and contract scope interact. A platform owner may specify interfaces while licensing major components; a supplier can offer stable interfaces and reusable software. Internal teams can also retain separate tools and review queues. The practical comparison is responsibility for a particular change and its test evidence, rather than a count of boxes or a presumed frictionless integrated model. That question needs release-level evidence.

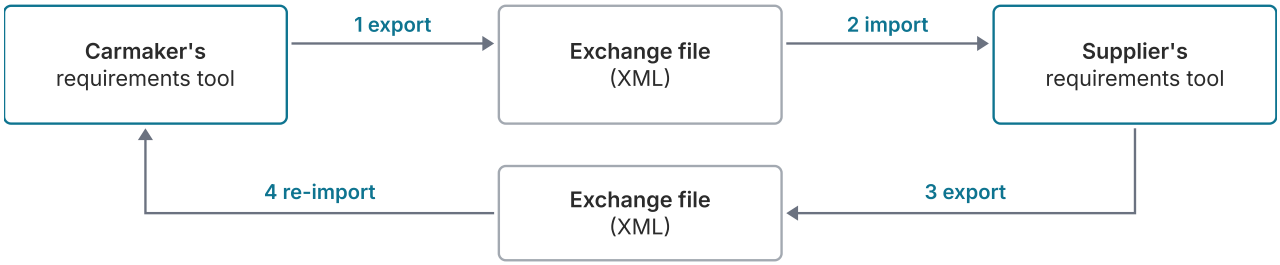
Sources: McKinsey (2026); NRG analysis

ReqIF is structured data; the hard part is preserving meaning.

OMG's ReqIF format exchanges structured requirements across tools. The January 2024 implementation guide remains a valid dated interoperability reference; an exchange file does not imply manual re-entry on every change.

The guide's system-attribute round-trip example; exchange may be automated and is not a measured workflow frequency.

SYSTEM-ATTRIBUTE EXAMPLE: STRUCTURED XML EXCHANGE



This example exchanges attributes; persistent identifiers must survive

Source: prostep ivip ReqIF Implementation Guide (2024)

OMG's Requirements Interchange Format uses XML to move structured requirements between tools. Sharing a repository is not a prerequisite. Export and import can be automated, with persistent identifiers preserving object identity. A file-based exchange can therefore be reliable; it is not the opposite of a data-based workflow.

The prostep ivip guide, version 1.10 dated January 26, 2024, gives a system-attribute example: export, import, return export and re-import. Its purpose is implementation interoperability, not measuring elapsed time or proving that every change repeats manual copying. Recommendations in that guide are non-normative.

NRG separates transport from semantics. Identifiers can survive while an attribute is interpreted differently, a revision is stale or a test link points to superseded evidence. Conversely, putting requirements in one database does not by itself align internal teams. Useful process checks ask what the current requirement means, who may change it and whether impact analysis remains linked to verification.

Automotive SPICE 4.1, dated August 24, 2026, supplies a current assessment model. Its existence does not establish universal assessment or make evidence automatically comparable. Review and integration remain inside a company as well as across a supply contract. The mechanism is possible reconciliation work; neither this guide nor the model supplies a per-exchange delay estimate.

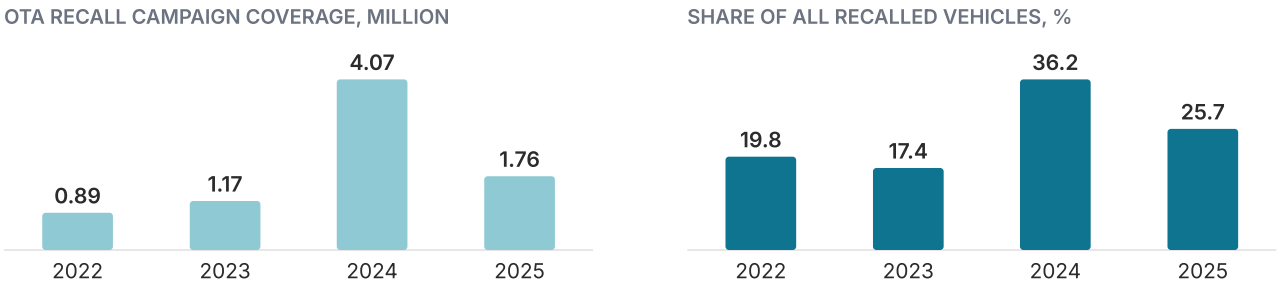
Persistent IDs

Preserve identity across exchanges, then check revision state, attribute meaning and test links. A shared repository alone does not establish traceability.

OTA recall campaigns covered a quarter of China's recall vehicle count in 2025.

SAMR's latest completed annual series counts vehicles within announced campaigns. Coverage does not establish completed repairs, unique OTA users or software quality. Annual campaign composition can produce large swings.

Vehicles covered by China OTA recall campaigns, and share of annual recall vehicle count; not repair completion



Source: State Administration for Market Regulation, annual recall notices; NRG calculation

OTA recall campaigns covered 0.887 million vehicles in 2022, 1.173 million in 2023, 4.068 million in 2024 and 1.756 million in 2025. Dividing by the annual recall vehicle totals gives 19.8%, 17.4%, 36.2% and 25.7% (NRG calculation). These are campaign-covered counts, not verified completed fixes. The non-monotonic series reflects the size and composition of campaigns; it is not a steady adoption curve.

Total recall vehicle count fell 39.1% in 2025 from 2024 (NRG calculation). That denominator change matters when reading OTA shares. A recall resolved through a software remedy can still require controlled testing and execution. The annual notice is not a measurement of the time between code completion and repair.

Wider OTA evidence uses different units. ResearchInChina reports installation in China passenger cars and a 2030 forecast. WAYS monitoring reports release versions and functions under its own tracked scope. The August 2026 observation is a separate period, not an annualised comparison with 2025.

NRG treats this activity as a reason to examine update responsibilities after sale. It does not show that ownership makes fixes quick or that every update waits for a supplier. Code rights, agreed support and validation obligations decide the path of a specific change. Campaigns, versions, functions and repeated update events must remain distinct measures.

Separate OTA populations and units; not directly comparable

Measure	Figure	Period
China passenger cars with OTA installed	76.2%; over 95% forecast	Jan–Oct 2025; 2030
Versions / functions in WAYS monitoring	46 / 576	August 2026
Regulator cumulative reported updates	4,047 OTA activity reports / 486m vehicle-update events	to end-2024; repeat events

Source: ResearchInChina (2026); China Passenger Car Association and WAYS (2026); SAMR

Sources: SAMR; China Passenger Car Association and WAYS; ResearchInChina; NRG calculation

Rules attach to covered products and approvals, not ownership models.

Update-management requirements and government filings are different obligations. The selected dates below are scoped milestones; applicable categories, characteristics and transition clauses require their own legal and technical assessment.

Selected scoped milestones; future dates shown separately. Consult the applicable categories and normative text.

Jul 2022	VCA / EU	R155 selected new-type scope
Jul 2024	VCA / EU	R156 new types; R155 broader scope
Feb 2025	China	Covered intelligent-connected product / OTA notice
Jan 2026	China	GB 44496 effective; Amendment No.1 on Jan 28
Jul 2026	VCA / EU	R156 new complete vehicles, with scope qualifiers
Jan 2027	China	Future: GB 47955 effective date
Jul 2029	VCA / EU	Future: R156 completed multi-stage vehicles

Source: VCA (updated May 2026); MIIT/SAMR (February 2025); official GB 44496 and GB 47955 status listings

VCA's May 22, 2026 guidance describes staged UN R155 cybersecurity and R156 software-update requirements for specified vehicle and approval categories. Its schedule includes R155 new-type phases in July 2022 and broader phases in July 2024, plus an R156 new-type milestone in July 2024. It identifies July 2026 for new complete vehicles and July 2029 for completed multi-stage vehicles under the relevant scope.

R156 application includes qualifiers concerning affected type-approved characteristics. An assessed software-update management system is not a universal duty to file every software change with a European government. This report does not apply one schedule to all vehicle categories or substitute the timeline for the applicable approval text.

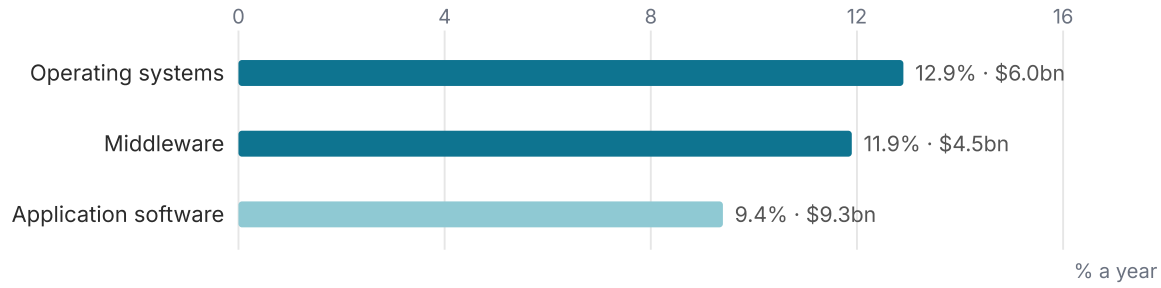
China's February 25, 2025 MIIT/SAMR notice covers specified intelligent-connected products and OTA activity, including recording and filing requirements. It should not be restated as every carmaker selling in China filing every possible update. Later production-consistency action adds context without establishing replacement of that notice.

GB 44496-2024 took effect January 1, 2026; Amendment No.1 was issued and effective January 28, 2026. Its substantive effect was not established here. GB 47955-2026 was issued June 27, 2026 and takes effect January 1, 2027, still future at this check. Detailed transition clauses were not verified and are omitted. Internal ownership does not remove applicable assurance obligations.

Layer-growth rankings depend on the publisher's classification.

GMI forecasts faster growth for operating systems and middleware than applications. Other definitions and horizons yield different rankings; the August 2026 MarketsandMarkets edition calls applications its fastest-growing layer.

Expected annual growth by software layer to 2035, with 2025 market size



Source: Global Market Insights (2026)

Global Market Insights estimates automotive software at USD 19.8bn in 2025: USD 6.0bn operating systems, USD 4.5bn middleware and USD 9.3bn applications. Its projected layer growth rates are 12.9%, 11.9% and 9.4% respectively to 2035. These values belong to one publisher's classification and are estimates, not audited software sales.

Mordor's current summary puts operating-system growth at 9.82% through 2031 and applications at 48.53% of its 2025 estimate. MarketsandMarkets' August 2026 edition instead calls applications fastest-growing, at 10.0% through 2033. Those taxonomies are not reconciled. A larger market base or a different CAGR alone does not reveal which inclusion causes the gap.

Platform control can give a carmaker authority over interfaces, but does not establish ownership of every component or the right to change licensed code. NRG's mechanism is conditional: explicit interfaces and support obligations can reduce unresolved decisions, while a proprietary platform can add dependence on its owner.

Capgemini's 2025 organisational survey asks distinct strategy and implementation questions. A unified-platform priority and successful scaling of an SDM use case are not the same construct or a measured conversion funnel. Neither percentage is a share of vehicles in service. The comparison motivates questions about implementation; it does not prove that owning a platform caused faster deployment.

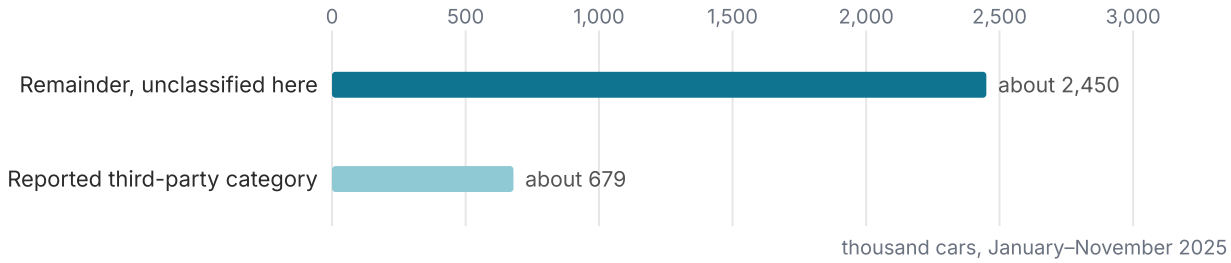
> **83%**
>
> call a unified platform key to strategy; separately, 14% report successfully scaling an SDM use case. Different constructs among automotive organisations (Capgemini, June 2025 fieldwork).

Sources: GMI (September 2026); Mordor (current 2026 summary); MarketsandMarkets (August 2026); Capgemini (2025)

A reported supplier category is not the complement of in-house ownership.

A January 2026 account of CAAM's city-NOA report supplies aggregate and supplier figures for January–November 2025. The full issuer taxonomy was not obtained; subtracting a category does not identify ownership of the remainder.

City-NOA reported category and residual, January–November 2025. Residual is unclassified here, not wholly in-house software.



Source: CAAM report via Xinhua (January 2026); NRG calculation; full classification unresolved

Xinhua's account reports 3.129 million passenger cars with city navigation-assisted driving in January–November 2025 and penetration of 15.1% using passenger-car insurance registrations as the denominator. The account calls the volume cumulative sales. The source's mixed labels should not be flattened into a share of all vehicle sales, and the observation is not a fresh 2026 adoption estimate.

A reported leading supplier's 414,400 units and 61.06% category share imply about 678,677 units in the reported third-party category. Dividing by 3.129 million gives about 21.7%; subtracting leaves roughly 2.45 million (NRG calculation). That arithmetic depends on the source's category boundary. It does not verify the full methodology.

The remainder may contain different development and licensing arrangements. It cannot be labelled wholly developed by carmakers, owned code or absence of suppliers. Unit volumes also do not establish the number of firms choosing each model. No conclusion about a majority of firms or the ownership of the biggest sellers follows.

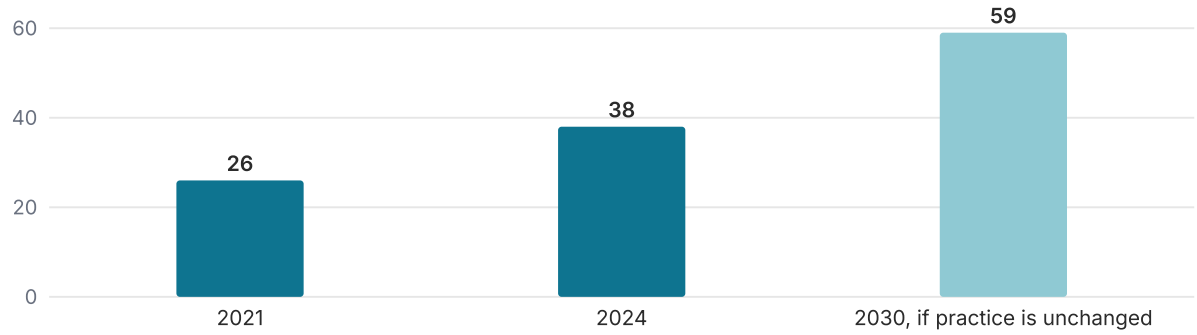
A lead partner can simplify a carmaker's immediate contracting interface while retaining sub-suppliers, software licences and tool exchanges behind it. NRG treats this as a possible responsibility arrangement, not a measured reduction to one hand-off. Interface rights, support scope and integration evidence matter more than the number of names on the first contract.

Sources: CAAM January 2026 report as described by Xinhua; primary taxonomy unavailable; NRG calculation

A dated spending model separates ownership cost from market revenue.

Roland Berger's October 2024 study, based on 30 experts, modelled software spending of USD 26bn in 2021 and USD 38bn in 2024. Its 2030 spending and saving figures are conditional scenarios, not observed results.

Automotive industry software development spending, USD bn a year



Source: Roland Berger (2024)

The study's USD 38bn estimate for 2024 and USD 59bn projection for 2030 describe industry software spending under its model. They are not interchangeable with software-market revenue estimates from other publishers. The 2021-to-2024 growth is about 13.5% using rounded endpoints (NRG calculation); the study describes roughly 14% using its own underlying model.

Its USD 17bn annual saving by 2030 is a potential outsourcing outcome, not a demonstrated cost reduction. Reuse, licensing, product differentiation and supplier economics determine what can be bought efficiently. The age of the study does not invalidate the dated scenario, but it must not be presented as a 2026 spending observation.

Roland Berger's May 2026 unbundling analysis provides newer procurement context: software scope and monetisation can be separated from hardware. It does not publish a directly comparable replacement total for the old model. NRG therefore uses it to explain contract mechanisms, rather than inventing an updated spending figure.

Internal development carries capability and maintenance costs; external software carries licence, integration and support obligations. Neither choice removes assurance or guarantees a short change path. The useful contract review asks who can alter an interface, which updates are included and how long support continues. These obligations can be clear under either sourcing model, and ambiguous under either one.

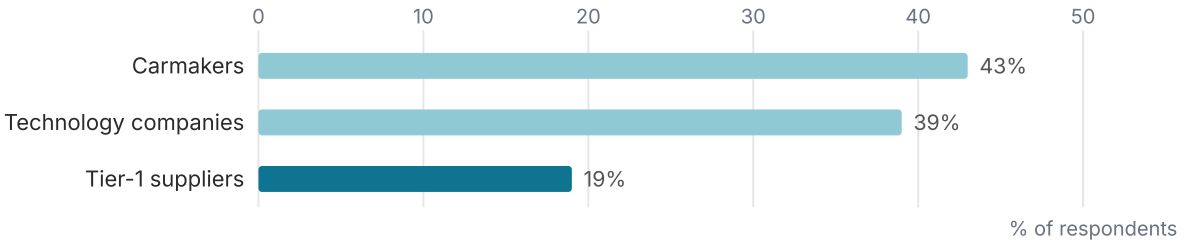
> **USD 17bn**
>
> potential annual saving by 2030 in the October 2024 study's outsourcing scenario; not an achieved result or a 2026 estimate.

Sources: Roland Berger (October 2024); Roland Berger unbundling analysis (May 2026); NRG calculation

Tracked platform reuse is a distinct survey measure.

AlixPartners' late-2025 survey asks about reuse enabled and tracked across model series and platform generations. Shares describe respondents, not the percentage of code reused.

Respondents achieving platform-level software reuse, by type



Source: AlixPartners survey (2026)

In the survey, 19% of Tier-1 supplier respondents reported tracked platform-level reuse, against 43% of automakers and 39% of technology companies. The automaker–supplier gap is 24 percentage points (NRG calculation). The full sample contains 495 automaker, 149 Tier-1 and 358 technology respondents. It is a self-report, not an audit of code, installed platforms or customer-specific engineering work.

Low tracked platform reuse does not prove suppliers rebuild all software for each customer. Program-level reuse, untracked reuse and architecture divergence are separate categories. A standard product may be reusable without meeting the highest maturity category. Reported reuse capability does not show whether a specific interface change is cheap.

NRG uses reuse as a question about amortisation and maintenance, not a verdict on every supplier's business model. Shared components can spread development cost, but integration and assurance may still differ by programme. A buyer needs to know which evidence travels with a reused component and which must be regenerated.

Separately priced software and support can make obligations visible without removing negotiation or lowering total cost. Partner preferences below are distinct questions with different organisational populations. They cannot show that platform providers deliver better releases or control fewer decisions. Release scope, evidence portability and maintenance rights need direct assessment for each component.

Separate organisational questions; not vehicle shares

Measure	Share	Basis
Hyperscalers / AI as critical partners	71–82%	carmakers, range across regions
Platform / embedded-software partners	51–76%	carmakers, range across regions
Exploring new supplier markets	up to 84%	automotive organisations

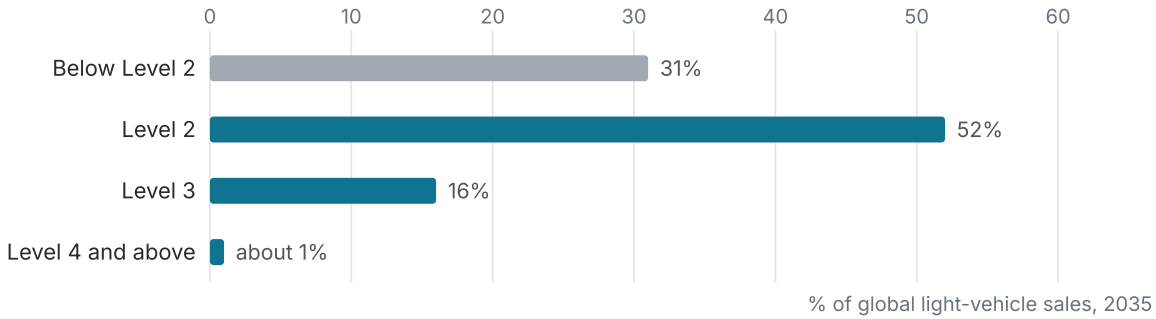
Source: AlixPartners (2026); Capgemini Research Institute (2025)

Sources: AlixPartners (2026); Capgemini (2025); NRG calculation

A 2035 driving-level forecast has a light-vehicle denominator.

McKinsey's January 2026 outlook covers passenger cars and light commercial vehicles. Its Current Trajectory figures are a longer structural scenario, separate from the 2026–2031 planning lens.

Global light-vehicle sales by driving-assistance level, 2035 forecast; below-L2 share is a rounded complement



Source: McKinsey (2026); NRG calculation

McKinsey models Level 2 at 52% of global light-vehicle sales in 2035, Level 3 at 16% and Level 4 or above at about 1%. The rounded complement is about 31% below Level 2 (NRG calculation). These are future sales shares, not shares of the installed fleet or all vehicle categories.

At Level 2 the driver supervises assistance. Level 3 describes automated driving within defined conditions with a takeover role for the driver. Those technical levels should not be conflated with one legal product classification or a blanket approval to deploy. China's forthcoming GB 47955 standard has its own scope; metadata alone does not settle detailed transition obligations.

In the same model, electronics move toward domain, central and zonal designs. The forecast suggests more combined functions and associated assurance work, not a measured count of external decisions. Software growth does not identify which company owns the code or who bears integration cost.

NRG's interpretation is to retain traceability through architecture changes. A requirement should still connect to its current interface assumptions and verification evidence after functions move between computers or teams. Reuse can preserve some evidence; changed operating conditions may require new testing. Neither ownership nor a driving-level label resolves those engineering questions. The scenario remains conditional on technology, regulation and market adoption.

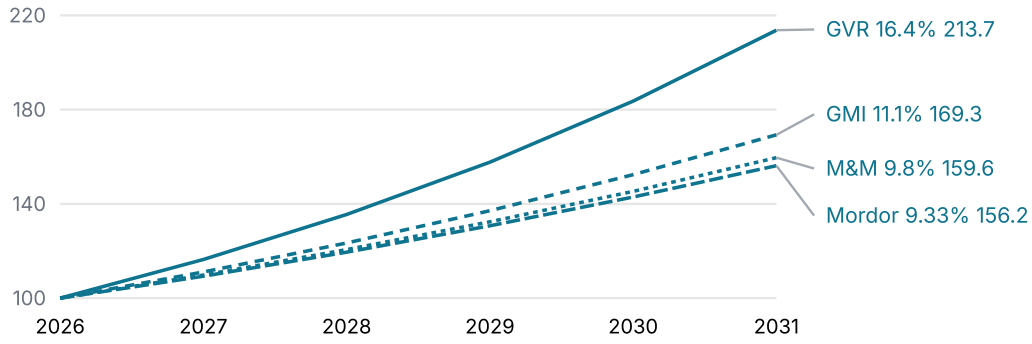
> **Nearly 20%**
>
> annual potential growth for driver-assistance / automated-driving software in McKinsey's 2025–2035 light-vehicle model; not measured 2026 growth.

Sources: McKinsey, The Automotive Software and Electronics Market Through 2035 (January 2026); NRG calculation

A five-year rate illustration is not a common market forecast.

Current publisher CAGRs applied to 2026 = 100 give about 156–214 by 2031 (NRG calculation). Different definitions remain different after normalisation. The lines are neither ownership scenarios nor probability bounds.

2026 = 100; five annual compounding intervals. Published rates, NRG illustration; scopes and official horizons differ.



Source: Mordor (2026–2031); M&M / GVR (2026–2033); GMI (2026–2035); current 2026 summaries; NRG calculation

DECISION RIGHTS	Contract-specific	Inspect code rights, interface approval and included support. A bought component can have clear update rights; internal ownership can retain dependencies.
EXCHANGE QUALITY	Tool-specific	Check persistent IDs, revision state and evidence links. XML is already structured data; manual reconciliation is not inevitable.
RESIDUAL WORK	Release-specific	Measure internal review and integration after code completion. The sources do not estimate a causal time saving from removing hand-offs.

SIGNPOSTS

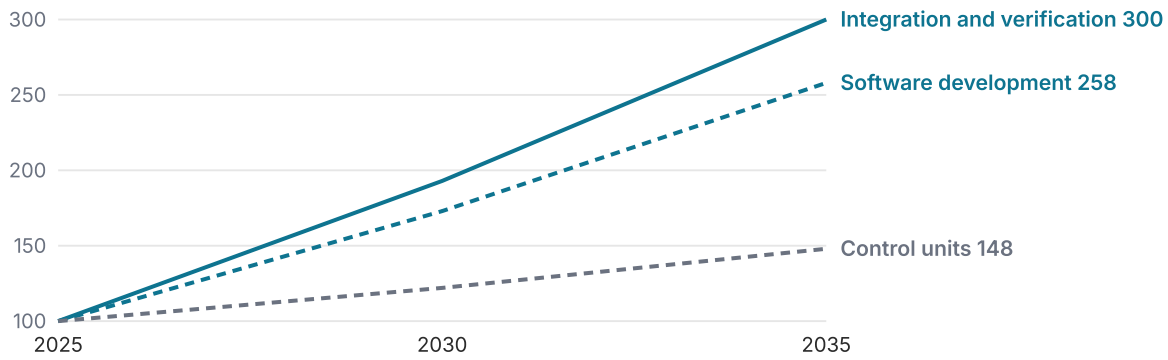
Repeat sourcing questions on comparable respondent populations, against the late-2025 baseline	Share of exchanges retaining identifiers, revision meaning and correct evidence links
Elapsed internal review and integration time for comparable release scopes	Observed deployment and repair completion, separately from recall campaign coverage

Sources: Current publisher summaries checked November 5, 2026; NRG calculation and analytical framework

The model's integration line grows; ownership does not identify the cause.

McKinsey's light-vehicle model projects integration, verification and validation from USD 14bn in 2025 to USD 27bn in 2030. The index compares spending lines, not external-hand-off counts or measured internalisation savings.

McKinsey light-vehicle software/electronics market lines, 2025 = 100; 2030/2035 forecasts; NRG calculation



Source: McKinsey (2026); NRG calculation

INTEGRATION AND ASSURANCE	+93% by 2030	The model includes complex central and zonal integration. Its growth does not establish that bought software causes the bill.
SOFTWARE DEVELOPMENT	+73% by 2030	Consolidation changes where work is performed and who pays. More internal code does not by itself shrink net integration effort.
CONTROL UNITS	+22% by 2030	Hardware grows more slowly in this model. Actual workload depends on architecture, reuse, interface dependencies and assurance scope.

SIGNPOSTS

- Integration and assurance spending under comparable scope and accounting
- Internal versus purchased work without treating a transfer as a saving
- Reusable verification evidence and the conditions under which it remains valid
- Actual review queues and change outcomes before attributing speed to ownership

Sources: McKinsey (2026); NRG calculation; NRG analysis

METHOD

How this report was built.

Public evidence was checked on November 5, 2026. Publication dates, observation periods and forecasts are separate. Current accessible summaries were used; inaccessible full methodologies and normative details are not treated as verified.

AlixPartners/Berylls: SDV Survey 2026, April publication, November–December 2025 fieldwork, especially sourcing, deployment and reuse charts. Capgemini: The Era of Software-Driven Mobility, September 2025 publication, June fieldwork. These are organisational self-reports, not vehicle shares or causal tests.

McKinsey: The Automotive Software and Electronics Market Through 2035, January 2026; light vehicles and Current Trajectory. MarketsandMarkets: August 2026 edition, 2026–2033. Mordor: current 2026 summary, 2026–2031. GMI: September 2026, 2026–2035. Grand View: June 2026, 2026–2033. Full paid methods were not reviewed; normalisation does not reconcile scope.

Engineering references: OMG ReqIF 1.2, July 2016; prostep ivip ReqIF guide 1.10, January 2024, system-attribute example; VDA Automotive SPICE 4.1, August 2026. Regulatory references: VCA May 2026 guidance, MIIT/SAMR February 2025 notice, official GB 44496 amendment and GB 47955 status listings. Detailed transition and amendment effects remain unresolved.

SAMR annual recall notices cover 2022–2025 campaigns; ResearchInChina's 2026 abstract observes January–October 2025 passenger cars. WAYS August 2026 monitoring is a separate scope. CAAM January 2026 data are reported by Xinhua; full taxonomy unavailable. The three-versus-five-year clock is Forbes reporting of the 2026 AlixPartners outlook. Roland Berger spending/saving figures are October 2024 model/scenarios; May 2026 unbundling is separate context.

An NRG estimate is an internal-research range; this report contains none. An NRG calculation uses cited public inputs: recall shares, implied category volumes, index paths and growth differences. Coordination and decision rights are bounded NRG synthesis, not measured causal effects. Companies are described rather than named. This is business and market research, not investment advice.

WORKING WITH NATHAN RESEARCH GROUP

Partner With Nathan Research Group

Architecture charts and sourcing percentages leave practical questions open: who may approve a change, how requirements remain traceable and where integration work waits. Industry practitioners can help a team investigate those questions within clear compliance boundaries.

WHO WE PUT IN THE ROOM

Vehicle architecture specialists	Central computing and electrical-electronic architecture, and how carmakers decide which functions to consolidate onto fewer computers.
Embedded software engineers	Software for vehicle controllers and central computers.
Supplier-management specialists	How carmakers work with software and electronics suppliers, across the industry.
Systems-engineering and requirements specialists	Requirements management and functional-safety practice.
Automotive software market specialists	Market sizing, the supplier landscape and how value in the vehicle is shifting from hardware to software.
Former R&D executives	Automotive software and electronics programs.

If your team is evaluating vehicle software development, the software-defined vehicle transition, or the broader automotive software market, tell us the decision you're trying to make.

START THE CONVERSATION

Web	nathanrg.com
Email	team@nathanrg.com General enquiries nathan@nathanrg.com Nathan Choi · CEO
Phone	+82 10 3059 3630 Nathan Choi · CEO +82 2 939 3630 Office · Seoul

Based on public sources and NRG's own research. Figures marked NRG estimate are rounded ranges checked against public data. Nathan Research Group does not request or facilitate material non-public information, and runs every engagement through a documented compliance protocol.