

Engineering Trust Into Electro-Mechanical Braking

A system-level view of ASIL-D design



Executive Summary

Electro-mechanical braking replaces hydraulic pressure generation and distribution with independently controlled electric actuators at the wheels. In a full four-corner dry architecture, the loss of the hydraulic fallback places greater reliance on electrical power, sensing, communication, computation, and wheel-end actuation.

No single component can provide ASIL D capability by itself; the system must detect and contain relevant faults within the applicable fault-tolerant time interval, transition the affected function to a defined controlled state, and preserve sufficient vehicle-level braking capability through the remaining healthy paths. The sections that follow examine why EMB is arriving now, what that safety principle requires at the component level, and what it means in practice for OEMs and Tier-1 suppliers building the next generation of by-wire vehicles.

Contents

Why EMB, Why Now	4
Architecture Evolution	4
Commercialization and Market Timing	4
Vehicle-Level Benefits	4
The 12V-to-48V Transition	5
Braking Availability for Automated Driving	6
Vehicle-Level Availability and Degraded Braking	6
Differential Braking as a Potential Fallback	6
Functional-Safety Architecture and ASIL Decomposition	7
SOTIF and Command Execution	7
The Core Engineering Problem: Safety Without a Mechanical Backup	7
Architecture-Dependent Outcomes	8
Diagnostic Monitoring Without Sensor Duplication	8
Power as a Fault Domain	9
System Architecture Overview	9
Technical Deep Dive	10
Power Delivery, Redundancy, and MCU Supervision	10
Wheel-Speed Sensor Interface	11
Driving the Motor: Gate Drivers and System Simplification	12
External Current-Sensing Options	13
Knowing Where the Motor Is: Rotor Position Sensing	14
Electronic Parking Brake and Static Hold	15
Translating Driver Intent: Pedal-Travel Sensing	16
What This Means for OEMs and Tier-1s	18
Conclusion	19

Why EMB, Why Now

Architecture Evolution

Braking architectures have evolved from mechanical linkage to hydraulic actuation and electronically controlled electrohydraulic systems. Direct electromechanical actuation at the wheels is now emerging alongside these established architectures. Rather than moving toward one end state on a fixed schedule, the industry is developing several architectures in parallel. Electrohydraulic braking (EHB) remains widely used; semi-dry systems combine electromechanical actuation at selected wheels – often the rear axle in announced programs – with hydraulic actuation elsewhere; and full-dry EMB uses electromechanical actuators at all four wheels, with no hydraulic fluid in the service-brake actuation path. OEMs select among these architectures according to vehicle-segment requirements, platform cost, production timing, packaging, electrical architecture, and safety concept. These architectures are therefore likely to coexist, with full-dry EMB emerging alongside – rather than immediately replacing – hydraulic and electrohydraulic systems.

Commercialization and Market Timing

Full-dry EMB is moving beyond laboratory development and into an early commercialization phase, with initial series-production programs reported for 2026. However, much of the broader market remains in development, validation, supplier selection, or pre-production. Full-dry EMB is therefore an emerging production technology, not yet a broadly adopted market standard. China has emerged as an important early market for EMB development and initial commercialization, while OEMs and Tier-1 suppliers in other regions continue to evaluate and develop semi-dry and full-dry systems. Regulatory frameworks are also evolving to address braking architectures without a conventional hydraulic actuation path.

Vehicle-Level Benefits

Several potential vehicle-level benefits are driving interest in EMB. Direct electromechanical actuation may reduce selected delays associated with hydraulic pressure generation and fluid transmission. One publicly reported OEM comparison cites approximately 80-100 milliseconds to full clamp for EMB versus about 150 milliseconds for an electrohydraulic system, although the measurement boundary and operating conditions must be confirmed before these values are used as a general comparison. Faster command-to-brake response could benefit time-critical functions such as automatic emergency braking, but system performance depends on the complete sensing, decision, communication, and actuation chain. Independent wheel actuators can also enable more direct control of individual wheel braking force and closer coordination with regenerative braking, stability control, and vehicle-motion functions. Hydraulic systems can also provide wheel-selective pressure control, so EMB's potential advantage should be framed in terms of actuation independence, local feedback, and the potential for finer control resolution rather than exclusive capability. Full-dry EMB removes the conventional master cylinder, brake-fluid lines, reservoir, and centralized hydraulic pressure-generation hardware from the service-brake actuation path. This may simplify packaging and assembly while eliminating brake-fluid filling, bleeding, and fluid-related service for that path. Any claimed reduction in total system mass or cost should account for the motors, power electronics, wiring, connectors, and mechanical transmissions added at the wheels. Because target clamp force and actuator response are electronically commanded and controlled, EMB can support software-based calibration and actuator-health monitoring using current, position, voltage, and temperature data. Over-the-air changes to braking behavior would require controlled development, cybersecurity protection, safety analysis, validation, configuration management, and regulatory review.




Figure 1. EHB vs. Semi-dry vs. Full-dry EMB Comparison

This paper focuses primarily on the full four-corner dry EMB architecture shown at right, where the fail-operational safety requirements discussed in the remainder of this paper are most acute; semi-dry designs, which retain hydraulic actuation on the front axle, are referenced only where the contrast helps clarify why those requirements exist.

The 12V-to-48V Transition

The development of EMB is closely linked to the evolution of vehicle electrical architectures. Current EMB development includes 12V implementations that leverage established vehicle power infrastructure. As actuator power requirements increase, however, the higher current required at 12V places greater demands on conductors, connectors, voltage-drop control, transient performance, and thermal management. At equal delivered power, a 48V rail carries approximately one quarter of the current of a 12V rail. Depending on conductor sizing and system constraints, this can reduce resistive loss and voltage drop or enable a lighter wiring harness. The final benefit depends on cable length, conductor cross-section, connector resistance, duty cycle, allowable voltage drop, temperature, and cost targets; voltage alone does not determine the system outcome. Electrical-power demand can be particularly challenging at the front axle during higher-deceleration events because dynamic load transfer increases front-axle braking demand. Several announced programs use rear-axle electromechanical actuation with hydraulic braking retained at the front as an intermediate architecture. This choice may limit the initial electrical-power and integration scope while preserving a hydraulic front-brake path, but it should not be presented as a universal development sequence. A 48V architecture may offer greater design margin for high-power or four-corner EMB, but the timing and extent of market adoption remain uncertain. The relative adoption of 12V, 48V, semi-dry, and full-dry systems will depend on vehicle segment, actuator design, electrical infrastructure, component cost, safety concept, and production strategy. In the near term, 12V and 48V EMB architectures are likely to coexist rather than one replacing the other on a fixed schedule.

Same Power. Less Copper.

480 W load / 1 m copper conductor




Figure 2. Illustrative 12 V/48 V conductor trade-off for a 480 W nominal load.

Per 1 m conductor; return path and connectors excluded. Actual results depend on temperature, duty cycle, bundling, voltage-drop limits, and OEM derating.

Braking Availability for Automated Driving

Vehicle-Level Availability and Degraded Braking

As automated-driving responsibility increases, the vehicle cannot rely solely on timely human intervention to mitigate a braking fault. For functions whose safety concept requires the system to achieve a minimal risk condition, sufficient braking capability may need to remain available after specified faults. The required availability, degraded performance, and operating duration must be derived from the item definition, hazard analysis, operational design domain, fault assumptions, fault-tolerant time intervals, and applicable regulations.

A four-corner EMB architecture can support this objective through independently controlled wheel actuators and vehicle-level redistribution of braking demand. Whether this degraded mode satisfies the safety goal depends on the braking capacity remaining at the available wheels, stopping-distance and stability requirements, and the timing of the coordinated response. Physical separation alone does not establish independence; the safety case must also address shared power, communication, computation, sensing, software, thermal, environmental, and mechanical dependencies.

Differential Braking as a Potential Fallback

Independent wheel-brake control can generate a yaw moment through differential braking and may contribute to a minimal risk maneuver following degradation or loss of the primary steering function. This capability is not exclusive to EMB: electronic stability control (ESC) and wheel-selective hydraulic brake-by-wire systems can also generate differential-braking yaw. EMB may improve the independence and controllability of this function by avoiding a shared

hydraulic pressure-generation path, provided that shared electrical, communication, and control dependencies are adequately addressed.

Differential braking is not equivalent to conventional steering. The available path-control authority depends on vehicle speed, tire-road friction, remaining brake authority, actuator availability, brake thermal condition, vehicle-stability constraints, and road geometry. Before this capability - sometimes referred to as "steer-by-brake" - can be credited as part of a vehicle's fallback strategy, its performance must be validated against the specified steering-failure modes, fault-detection and actuation-response timing, and relevant operating conditions.

Functional-Safety Architecture and ASIL Decomposition

ASIL decomposition is a general technique defined in ISO 26262 and is not specific to EMB. Where permitted and justified, an ASIL D safety requirement may be decomposed into redundant requirements allocated according to an applicable decomposition scheme - for example, ASIL B(D) and ASIL B(D). The original safety goal remains ASIL D, and the decomposition does not remove the rigor required to satisfy it.

Four physically separate EMB wheel actuators may support fault containment and provide architectural resources for degraded vehicle-level braking, but they do not by themselves establish a valid decomposition. Independence must be demonstrated for the elements to which the decomposed requirements are allocated. The safety case must address freedom from interference and dependent failures, including common-cause and cascading failures involving shared power, communications, computation, software, sensing, thermal exposure, and mechanical resources.

SOTIF and Command Execution

SOTIF (ISO 21448) addresses hazards that can arise without a component malfunction, including limitations in the perception, decision-making, or control functions used for automated driving - for example, misclassification of an object or scene under challenging conditions. SOTIF is primarily a vehicle- and function-level consideration that extends well beyond the brake actuator, and its validation requires scenario-based testing across the operational design domain.

Within that broader validation, the EMB actuator has a narrower but important role: executing valid and authorized braking commands predictably within defined timing and performance bounds, reporting its available capability and diagnostic status, and responding in a defined manner to implausible, conflicting, delayed, corrupted, or out-of-range commands. Lower latency can improve braking performance, but safety depends on correct command handling and validated end-to-end timing - not latency alone. Vehicle-level perception and decision functions remain responsible for determining when and how much braking to request.

The Core Engineering Problem: Safety Without a Mechanical Backup

Earlier braking architectures could retain some braking capability through a hydraulic or mechanical path after selected electronic faults. A full four-corner EMB system removes the conventional hydraulic pressure path from the service-brake actuation chain, placing greater reliance on electrical power, sensing, communication, computation, motor control, and mechanical actuation. Its safety concept must therefore define how relevant faults are detected and contained, what controlled response each affected wheel actuator executes, and how the vehicle preserves the braking and stability capability needed to reach a safe or minimal risk condition. In functional-safety terms the braking function shifts from fail-safe, where removing electronic actuation still left a hydraulic path intact, to fail-operational, where the required capability must survive the fault.

This is where it is important to distinguish among different forms of redundancy. Steer-by-wire architectures may require redundant torque-generation, sensing, control, communication, and power paths because loss of lateral-control authority can be hazardous. The implementation may use duplicated motors, windings, controllers, power channels, or other independent actuation resources rather than a second complete steering mechanism.

Full four-corner EMB can use a different fault-tolerance strategy. No individual wheel node necessarily has to remain operational after every fault. Instead, fail-operational or degraded braking capability may be achieved primarily at vehicle level through local fault detection and containment, followed by coordinated redistribution of braking demand among the remaining available actuators.

Wheel-Level Responsibility

- Detect locally observable faults within the applicable fault-tolerant time interval
- Contain the effects of a local fault within the affected actuator, where required by the safety concept
- Execute the defined local fault reaction and transition to the controlled state specified for the fault and operating condition
- Report diagnostic and availability status to the vehicle controller, where communication remains available

Vehicle-Level Responsibility

- Redistribute braking demand among the remaining available actuators
- Manage braking asymmetry and maintain vehicle stability within defined degraded operating limits
- Provide the degraded stopping performance required for the fault and operating condition
- Support transition to the vehicle-level safe state or minimal risk condition defined by the safety concept

Architecture-Dependent Outcomes

Whether vehicle-level redistribution provides the required degraded braking performance depends on the platform architecture and operating condition:

- **A full four-corner EMB architecture** with vehicle-level control capable of redistributing braking demand may preserve the degraded stopping performance and stability control required by the safety concept using the remaining available actuators. Depending on operating conditions and available brake authority, differential braking may also contribute to directional control, but it is not equivalent to maintaining conventional steering capability.
- **Semi-dry architectures** combine wheel-level electromechanical fault handling with the redundancy and fallback characteristics of the hydraulic path serving the remaining wheels. Their degraded braking strategy may therefore use both redistribution among available EMB actuators and continued hydraulic braking, depending on the system design.
- **Architectures without vehicle-level brake-force redistribution capability** cannot rely on this mechanism to mitigate the loss of a wheel actuator. The wheel node's local fault-detection and controlled-response requirements remain, but the vehicle safety case must determine whether the resulting loss of braking contribution provides acceptable stopping and stability performance.
- **Power architecture** directly affects wheel-node availability. With a single input, loss of that supply may make the actuator unavailable and trigger its defined local fault response. Where the safety concept requires input-level redundancy, monitored dual inputs may allow transfer to a healthy source and continued operation. This does not eliminate faults in shared downstream conversion, grounding, control, or distribution resources.

Diagnostic Monitoring Without Sensor Duplication

The baseline EMB wheel module considered in this paper does not duplicate every sensing or actuation path. Instead, local sensing, power-management, and motor-control functions use device diagnostics and system-level plausibility checks to detect faults that could compromise brake-force control. If reliable operation can no longer be maintained,

the actuator executes its defined local fault response and reports its status to the vehicle controller. These diagnostics improve fault detection but do not constitute independent redundant sensing channels.

Fault tolerance is therefore achieved primarily through wheel-level fault detection and containment combined with vehicle-level redistribution of braking demand. Redundant or diverse sensing may be used in other programs where explicitly required by the allocated safety requirements, but it is not assumed for the baseline wheel module described here. This approach is known as system-level module redundancy: the Central ECU cross-compares the wheel-speed signals from all four corners, so four ASIL B channels, one per module, support an ASIL D safety goal at vehicle level. Its validity rests on the independence of those four channels and on the integrity of the ECU performing the comparison.

Power as a Fault Domain

The same principle applies to power delivery. A power-path fault can disable the sensing, control, communication, and motor-drive functions it supplies, so input power, regulated rails, and MCU supervision must be included in the wheel node's fault-detection and containment strategy. The following section examines how that requirement can be addressed in single- and optional dual-input power architectures.

System Architecture Overview

At the vehicle level, a 12V/48V Power Distribution Box and a Central ECU (the platform's central MCU) fan power and control out to each wheel-corner EMB actuator independently, keeping the power-delivery and communication/control paths architecturally distinct all the way to the wheel node.




Figure 3. EMB Vehicle-Level Power and Communication Architecture




Figure 4. EMB Wheel-Node Actuator Block Diagram

Wheel-Node EMB Actuator: 1. Power Management & Wheel-Speed Interface | 2. Motor Gate Driver & Phase-Current Sensing | 3. Motor Rotor-Position Sensor | 4. EPB / Static-Hold Actuation **Vehicle-Level Input—Shown for Context:** 5. Pedal-Travel Sensor

At the wheel-node level, several subsystems work together to support local fault detection, containment, and defined fault response. These include a power-management stage that supplies and supervises the node and can be preceded by an optional dual-input disconnect stage where required by the vehicle safety concept; a wheel-speed sensor interface that powers, decodes, and reports local speed information; a motor gate driver that controls the power stage and, in the implementations described here, integrates in-line phase-current sensing; a motor rotor-position sensor that supports commutation and actuator-travel estimation, with interference rejection appropriate to the selected sensing technology; and, in the architecture described here, an EPB actuation path that uses dedicated outputs of the wheel-node driver to control the static-hold mechanism.

The pedal-travel sensor shown in Figure 4 is outside the wheel-node boundary. It is a vehicle-level driver-command input to the Central ECU (or vehicle-level brake controller), not a wheel-node subsystem. The following sections first examine the wheel-node functions and then address pedal-demand sensing separately as a vehicle-level interface.

Technical Deep Dive

Power Delivery, Redundancy, and MCU Supervision

Power delivery is subject to the same discipline as sensing and motor control: because there's no hydraulic fallback, a fault in the power path can cascade into the loss of the very components it feeds — sensors, motor drivers, and the MCU — so the wheel node's power architecture applies the same detect-and-contain discipline, acting within the applicable fault-tolerant time interval. That same principle extends to the power input itself, for programs that require it: dual 48V input isn't mandatory for every EMB power stage — many designs run a single input — but where a program's safety case calls for input-level redundancy, real redundancy means switchover, not just a second wire. Where required by the vehicle safety concept, two battery inputs can be combined through monitored disconnect switches ahead of one shared downstream conversion path — shown in the System Architecture Overview diagram

above. Redundancy in this configuration applies to the battery inputs and the disconnect switch paths; the 48V-to-12V pre-regulator (APM81815, 48V systems only) and the 12V PMIC (A81415) remain a single, shared downstream path rather than being duplicated. If one input degrades, the switches transfer the load to the healthy input to support continued wheel-node operation — the same fault-detection-and-safe-state principle applied to the power domain. Because OEMs run 12V and 48V platforms side by side today, a single component portfolio rated across both voltage classes carries a design through that transition without a full re-spin per platform generation.

The wheel-node electronics package is, in effect, a small high-performance computer bolted to a brake caliper — a safety MCU, gate drivers, sensors, and communication transceivers, all needing stable regulated power even while the vehicle battery sags under the transient current draw of an emergency stop. Allegro's A81415, a purpose-built EMB PMIC, addresses this on two fronts simultaneously. As a power source, its 2.2MHz buck-boost pre-regulator holds an intermediate rail stable through battery transients and feeds five internal LDOs (including a $\pm 1\%$ -accuracy rail dedicated to the MCU's ADC). As a safety supervisor, the A81415 runs three independent watchdogs (pulse-width, window, and question-and-answer/challenge-response) that continuously verify the MCU is executing correctly, forcing a safe state within milliseconds if it isn't.




Figure 5. EMB Power Delivery and MCU Supervision Architecture

Wheel-Speed Sensor Interface

In a conventional hydraulic or electro-hydraulic braking system, each wheel-speed sensor (WSS) requires its own dedicated, long-haul wiring harness routed all the way back to a central braking ECU. In a decentralized EMB architecture, this topology is highly inefficient. Instead, the sensor output can be routed directly into the EMB actuator module already residing at that wheel corner. Here, the module's local MCU decodes and processes the signal, eliminating redundant wiring and allowing wheel-speed sensing to operate natively within the vehicle's zonal architecture.

The A81415 EMB PMIC is purpose-built to execute this consolidation. Alongside the power delivery and MCU supervision functions described earlier, it integrates a fully protected WSS interface directly into the power management IC. This eliminates the need for discrete circuitry or a stand-alone interface chip. The interface drives the wheel-speed sensor through high-side and low-side switches—protected against short-to-battery and short-to-ground faults—utilizing low-side current sensing to decode the signal. It then reports both the speed data and diagnostic status over the same SPI bus already used for power-management telemetry.




Figure 6. A81415 Wheel-Speed Sensor Interface Block Diagram

From a safety and compatibility standpoint, the integrated decoder covers the full spectrum of industry-standard protocols: standard two-level (speed-only), PWM two-level with tooth pulses (including standstill), and three-level AK protocol at both standard and high (14 mA/28 mA interpolated-pulse) resolutions. More importantly for the system's ASIL-D safety case, the interface actively supervises the physical health of the sensor connection. It features onboard thermal protection, leakage-current compensation up to 2 k Ω on the sensor lines, and dedicated fault detection for short-to-ground/open-circuit, short-to-battery, out-of-range pulse timing, low supply voltage, and signal grey-zone conditions. Because that protocol support is configurable rather than fixed in hardware, a design built around the A81415 can accommodate whichever wheel-speed sensor protocol an OEM specifies without board-level rework - a flexibility a discrete, protocol-specific interface implementation would not offer.

The result is a system with significantly higher integration and diagnostic coverage than what a generic safety PMIC or traditional braking ASIC can provide. By folding the WSS interface into the PMIC, the design achieves one less discrete component per wheel node, eliminates a set of vehicle-level harness runs to route and diagnose, and embeds a critical sensing function directly inside the local architecture of the EMB actuator.

Driving the Motor: Gate Drivers and System Simplification

Once rotor position is known, the controller has to energize the correct motor phases through a three-phase gate driver and external power MOSFETs — and in a compact, vibration-heavy caliper, the wiring required to do that reliably becomes its own engineering problem. A conventional approach needs six dedicated PWM traces between the MCU and the gate driver, plus additional lines for current sensing, diagnostics, and enable signals — 15 or more traces in an assembly where every millimeter of board area and every EMI path matters. Allegro's automotive 3-phase gate drivers address this directly across both voltage classes: the AMD89106 and AMD89108 for 12V systems, and the AMD89186 and AMD89188 for 48V systems, all four sharing the same functional architecture.

The AMD89108 and AMD89188 support PWM-over-SPI, letting the MCU send duty-cycle and timing commands digitally over the existing SPI bus rather than driving six physical PWM lines.

All four parts also integrate two in-line current-sense amplifiers, reading phase current back over SPI directly from the gate driver — when paired with one of these parts, no separate discrete current sensor is required. This in-phase sensing approach — shared across all four parts — is arguably their most important feature: because current is measured directly in the motor phase leads rather than in the low-side return path, it places no restriction on when in the PWM cycle current can be sampled, and it supports higher accuracy through simple averaging rather than precisely timed single-shot sampling, while also detecting a wider range of fault conditions than low-side sensing can. And because only two in-phase sensors are needed per motor — the third phase current can be reconstructed from the other two under defined sampling and operating conditions, with that derived channel continuously checked by the gate driver's own offset and gain diagnostics rather than a physical third sensor — it delivers that accuracy and fault coverage with less resistive power loss than a three-shunt implementation.

The result is a smaller, cheaper MCU footprint and fewer PCB traces to route around a hot, high-vibration mechanism, plus a shrinking physical footprint for the gate driver package itself as this integration trend continues. That same digital interface also lightens the MCU's software burden: commutation timing and diagnostics that would otherwise require dedicated timer peripherals and interrupt-driven PWM management become simple SPI register reads and writes, freeing firmware complexity and MCU headroom for the over-the-air tuning and higher-level control that the software-defined vehicle model depends on.




Figure 7. 3-Phase Gate Driver Application Circuit

External Current-Sensing Options

Precise braking force depends on precise motor torque control, and torque is directly proportional to phase current — so the system needs to measure current flowing into the motor phases quickly and accurately, often at peaks between 50A and 150A. Where the gate driver already integrates in-line current sensing, as described above, this requirement is already met with no additional component. For designs paired with a gate driver that lacks that built-in capability, current measurement typically falls to a separate shunt-based sensor, which carries real costs: shunt-based sensing generates significant thermal losses at these currents, difficult to manage inside a sealed, already-hot caliper, and shunts typically require separate isolation amplifiers to protect the low-voltage MCU from motor-side transients. Allegro's isolated magnetic current sensors (such as the ACS37220) solve this by integrating a Hall-effect element and a primary copper conductor into a single surface-mount package with well under 100 microohms of resistance — substantially lower power dissipation than a typical discrete shunt under comparable operating conditions, with built-in galvanic isolation and a dedicated overcurrent fault output. For programs using a gate driver without inline sensing built in, that reduction in thermal load — combined with in-phase sensing's ability to support simple, high-accuracy averaging and detect more fault conditions than low-side sensing alone — is what makes it possible to close a torque-control loop supporting an ASIL D safety goal without adding a thermal management problem to the caliper. A two-sensor implementation can reconstruct the third current in a three-wire motor, but its suitability depends on sampling conditions, diagnostic coverage, and channel fault behavior — a role that per-channel offset and gain diagnostics built into the current-sensing IC help fulfill. A three-sensor implementation instead measures all three phases directly, allowing the sum $I_u + I_v + I_w = 0$ to be evaluated as a plausibility check; a phase short to ground, or a drifting or failed sensing channel, then presents as a non-zero sum without relying on a reconstruction that presumes the very fault-free behavior being checked. In-line sensing measures only current reaching the motor windings; detecting a fault that bypasses them — a phase-to-chassis short, or a shoot-through event — requires a DC-link measurement upstream of the bridge.





Figure 8. 3-Phase In-Phase Current Sensing Architecture

Knowing Where the Motor Is: Rotor Position Sensing

An EMB actuator uses a brushless DC motor driving a ball-screw or roller-screw mechanism to apply clamping force, which means the controller needs continuous, high-resolution knowledge of rotor position both to commutate the motor and to track pad travel precisely. The central challenge is magnetic interference: high motor-phase currents generate stray fields strong enough to corrupt a conventional magnetic angle sensor's reading, and shielding against that interference adds weight, size, and cost that a wheel-hub caliper assembly can't easily absorb. This isn't a 48V-specific problem — the same stray-field challenge exists in today's 12V EMB systems, and in other by-wire actuators such as electric power steering, since any motor drawing significant current near a position sensor faces it regardless of voltage class or application.

Resolvers have long been the default answer to this problem in demanding motor-control applications — as electromechanical transformers, they are highly robust against external magnetic interference, which is exactly why they remain common in traction-motor drives for electric vehicles today. But a resolver is fundamentally an analog device: it requires a dedicated excitation signal, analog-to-digital demodulation circuitry, and a comparatively large, heavy rotating transformer assembly — none of which fits easily into the tight space and cost envelope of a wheel-hub caliper. What EMB needs is robust interference rejection delivered by a compact semiconductor position sensor with a digital output that a caliper-mounted PCB can accommodate.

Allegro addresses this with two sensing architectures that reject stray fields by physical design rather than by shielding. On the Hall-effect side, Allegro offers the A33023, a differential-Hall angle-sensor SoC with motor-commutation outputs and single- or dual-die package options. On the inductive side, the A17802 is a single-die inductive position sensor that rejects interference through frequency-domain filtering on its coil-based sensing element rather than through a magnetic core at all. Both deliver latency in the tens-of-microseconds range with sub-degree accuracy. Where a program's safety case calls for diverse-physics redundancy, a Hall-based sensor from this family can be paired with the inductive A17802 at the system level; hardware-level dual-die redundancy within one package is currently a Hall-sensor option.



- 6 Hall plates form 3 differential channels:
CH1 = HP1-HP4, CH2 = HP2-HP5, CH3 = HP3-HP6
- Differential sensing cancels common-mode stray fields
- The main angle is calculated by combining all three channels
- Comparing composite angles supports ASIL-D safety goals at system level

Figure 9. A33023 Hall-Plate Sensing Arrangement

Electronic Parking Brake and Static Hold

Service braking and parking hold impose different requirements. Service braking requires dynamic, repeatable control of wheel-brake force, whereas parking hold must retain sufficient clamp force for an extended period—including on a grade—without continuous motor current.

In a representative electrohydraulic braking architecture, parking hold is provided by a separate electric parking-brake actuator at the rear caliper, with its own motor, drive electronics, harness, and mechanically self-retaining transmission. A full-dry EMB architecture already includes an electromechanical service-brake actuator at the wheel, creating an opportunity to integrate the static-hold function rather than add a second complete actuator. Figure 10 illustrates these two implementation approaches. The comparison is representative rather than universal; the actual implementation depends on the vehicle and actuator architecture.

	EHB (Today)	EMB (Built-In)
ACTUATOR	Separate motor per rear caliper	No separate EPB unit — solenoid + pawl added to rear caliper
CIRCUIT	Dedicated EPB harness & driver	Dedicated half- and full-bridge, same driver IC
HOLD	Self-locking screw, own motor	Pawl locks a ratchet gear — motor can't self-lock
REDUNDANCY	Separate from hydraulic path	Preserved — pawl path is independent of the motor/inverter



Figure 10. EPB Actuator - EHB Add-On vs. EMB Built-In Comparison

One possible EMB implementation reuses the service-brake actuator and adds a mechanically retaining pawl or latch to the wheel-end mechanism. Once engaged, the latch can maintain holding force without continuous motor torque. The required driver topology depends on the mechanism's actuation and release requirements and may use either a half-bridge or full-bridge implementation.

Allegro offers half-bridge and full-bridge driver options for different static-hold actuator architectures. For 48 V actuator power stages, the A89500 is a 100 V half-bridge gate driver that controls external MOSFETs, allowing the power stage to be scaled to the actuator's current and thermal requirements. Its 3 mm x 3 mm package helps minimize control-stage PCB area in space-constrained wheel-end electronics. The device also includes an integrated bootstrap diode and protection functions such as supply-undervoltage lockout, overtemperature protection, and cross-conduction protection.

For lower-voltage applications that prioritize integration, the A4950 combines a full-bridge motor driver, integrated power MOSFETs, and current sensing in a single device. Its 7 V-to-40 V operating range and 3.5 A peak-current capability can support compact EPB or static-hold actuators within the electrical and thermal limits specified in the product documentation. The integrated power stage reduces external component count and PCB complexity, although its fixed current and thermal capability provides less output-stage scalability than an external-FET architecture.

Product selection depends on supply voltage, actuator topology, required direction of active control, peak and continuous current, duty cycle, thermal environment, protection and diagnostic requirements, available PCB area, and the defined response to electrical or mechanical faults. The A89500 and A4950 should therefore be presented as architecture-dependent driver options rather than interchangeable solutions.

The mechanical latch can provide a holding path that remains independent of motor torque after engagement. However, mechanical independence after engagement does not establish independence of the complete parking function. The safety analysis must also address the command, electrical power, driver, actuator, feedback, mechanical linkage, engagement-confirmation, and release paths. Relevant fault conditions include failure to engage, partial engagement, unintended release, unintended application while driving, loss of supply during actuation, driver overtemperature or shutdown, feedback disagreement, and failure to release under emergency or low-voltage conditions.

Translating Driver Intent: Pedal-Travel Sensing

In a full-dry braking system without a hydraulic pedal-to-wheel command path, pedal-position sensing is a primary input to the vehicle-level driver-command chain. Its output must provide the accuracy, availability, and diagnostics required for the brake controller to determine the driver's requested braking demand. Unlike sensors within an individual wheel actuator, the pedal interface is a shared vehicle-level function whose loss or corruption can affect commands to all four wheels. It is a single point of failure for driver intent, with no fourth corner to cross-check it against, so the system-level module redundancy used for wheel-end sensing does not apply here. The redundancy must instead be localized within the pedal-sensing function itself — dual-die, dual-package, or a triple-sensor arrangement resolved by two-out-of-three voting.

Allegro offers several contactless sensing options for this function. The A31315 uses 3D Hall-effect sensing, while the A33022 uses a differential-Hall architecture designed for stray-field rejection. Both are Hall-effect-based technologies. The A17802 provides an inductive alternative where a different physical sensing principle is desired.

Figure 11 shows an illustrative architecture using two separately packaged A17802 inductive sensors and a dual-die A33022 Hall sensor. The two inductive devices provide same-principle redundancy with package-level separation, while the dual-die A33022 provides two Hall sensing channels in one package. Combining the inductive and Hall channels introduces physical-principle diversity.




Figure 11. Illustrative redundant pedal-position sensing architecture combining separately packaged inductive sensors with a dual-die Hall sensor.

Independence, diagnostic coverage, and suitability for an ASIL D system design must be demonstrated at system level.

Depending on the selected device and configuration, pedal-position data can be linearized on chip and transmitted over the SENT protocol. The complete driver-command path—not sensor latency alone—must provide the accuracy, update rate, data integrity, diagnostics, and fault response required by the vehicle safety concept. These requirements apply both to manually requested braking and to the arbitration of pedal input with driver-assistance or automated-braking commands.

The number of sensors or dies does not by itself establish independence or compliance with an ASIL D safety goal. The system safety analysis must also address shared targets, mechanics, PCB, power, ground, communication, controller inputs, diagnostics, and environmental exposure. Pedal feel remains a function of the mechanical pedal simulator, any active-feedback mechanism, and the associated control software—not the position sensors alone.

What This Means for OEMs and Tier-1s

System simplification and functional-safety engineering should be considered together. Integrating power supervision, wheel-speed interfacing, gate drive, current measurement, and position sensing can reduce external components and interconnects, but it may also introduce shared-device dependencies. Component selection should therefore begin with the vehicle safety concept and allocated technical requirements—not with maximum integration or minimum component count.

Design Decision	Key Question
Vehicle architecture	Full-dry or semi-dry? What degraded braking performance must remain after specified faults?
Power architecture	12 V, 48 V, or both? Is a single input sufficient, or is input-level redundancy required?
Control and communication	Which functions are centralized or local, and what happens after controller or network faults?
Wheel-speed interface	Local wheel-node interface or centralized ECU connection? What sensor protocols, protection, diagnostics, and vehicle-level plausibility checks are required?
Rotor-position sensing	Which sensing technology provides the required accuracy, interference rejection, and diagnostics? Does the safety analysis require a single channel, dual-die or dual-device redundancy, or different physical sensing principles?
Current sensing	Integrated or external? Are two measured phases sufficient, or are three required?
EPB / static hold	Which actuator and driver topology provides verified zero-power holding and controlled release?
Pedal-demand sensing	What channel availability, plausibility monitoring, and controller independence are required?
Safety evidence	Which safety manuals, FMEDA data, SEooC assumptions, diagnostics, and dependent-failure analyses are available?

Integration can simplify the wheel-node design, but fewer components do not automatically mean fewer safety concerns. OEMs and Tier-1 suppliers must evaluate both the integration benefits and the resulting shared power, package, communication, and diagnostic dependencies. Production EMB programs require clearly defined diagnostic coverage, fault-response behavior, and supporting safety evidence.

Conclusion

The shift from hydraulic fluid to electric wires is, at its core, a shift in where trust lives—from hydraulic pressure transmission to the coordinated behavior of semiconductors, software, electrical power, and mechanical actuators. Earning that trust in EMB does not come from any single high-performance component; it comes from an architecture in which every subsystem follows the same discipline: at the wheel node, detect and contain locally observable faults within the applicable fault-tolerant time interval and execute the controlled response defined for the fault and operating condition; at the vehicle level, coordinate the remaining available wheel actuators to provide the degraded braking and stability capability required to reach a safe state or, where applicable, a minimal risk condition. That is the engineering challenge EMB must solve and the standard against which power, sensing, communication, computation and software, motor control, and mechanical actuation should be evaluated.

That discipline appears throughout the system: position-sensing technologies designed to reject external interference without relying solely on shielding; gate drivers that integrate in-line phase-current sensing; and, where required by the vehicle safety concept, an optional dual-input power architecture that can transfer the load to a healthy source. These approaches place sensing, control, and supervision close to the points where relevant faults can be detected and contained. Realizing the benefits of this integration still requires evaluating shared power, package, communication, computation, and diagnostic dependencies—and the dependent failures they may introduce.