TRUSTWORTHY VEHICLE CONTROL

Efficient coexistence of automotive safety, cybersecurity, hard real-time execution and high availability

SAFE | **SECURE** | **REAL-TIME** | **AVAILABLE**

From protected information to controlled physical vehicle effect

Hans-Leo.Ross@Cariad.Technology

SUMMARY The architecture is successful only when trustworthy, timely and sufficient physical control remains available after faults and attacks.

Hans-Leo Ross (TX-SB) | October 2026

Mission: Protect vehicle controllability, not merely individual components

The required vehicle behaviour is the reference for every safety, security and availability decision.

Vehicle and traffic objective

- Controllable lateral and longitudinal motion
- Bounded response time and predictable degradation
- Sufficient physical actuation capability
- Controlled transition to a minimal-risk condition

Closed-loop assurance chain

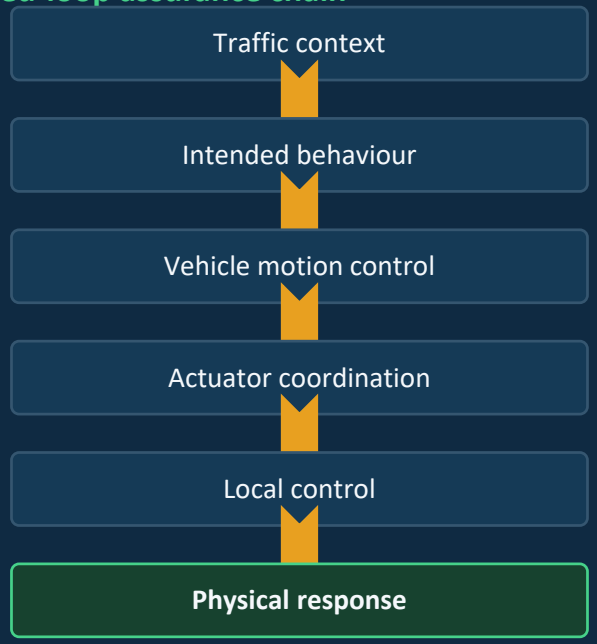

Design consequence

- Security findings become safety-relevant when they cause missing, late, stale or untrusted information.
- A cryptographically valid command may still be physically dangerous.
- The vehicle reaction must depend on operating state, FTTI, redundancy and remaining capability.

SUMMARY Safety and security are different disciplines, but both must preserve the same closed-loop vehicle capability.

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Assign every protection mechanism a precise trust dimension

No single mechanism establishes trustworthy closed-loop control.

Source trust

SecOC, MAC, certificates, authenticated boot

Is the source authorised?

Data integrity

E2E, CRC, ECC, MAC

Was information altered?

Freshness

Counters, timestamps, replay protection

Is it current?

Timing

Deadline, WCET, TSN, time sync., policing

Is it available in time?

Execution

CFM, lockstep, program supervision

Was the intended path followed?

Functional validity

Plausibility, state machine, safety envelope

Is the result permissible?

Effect integrity

Feedback, observers, residuals

Was the physical effect achieved?

Availability

Redundancy, reconfiguration, fallback

Can capability continue?

SUMMARY Use complementary evidence: identity, integrity, freshness, time, execution, plausibility, physical effect and remaining capability.

INTERNAL

Mission: Combine safety and security without confusing their responsibilities

Authentication protects provenance; safety supervision protects behaviour and physical effect.

Safety mechanisms

- ECC, lockstep and hardware diagnostics
- Watchdog, alive, deadline and logical supervision
- AUTOSAR E2E, CRC, counters and Data IDs
- Plausibility, command monitoring and safety envelopes
- Redundancy, voting, fallback and graceful degradation

Security mechanisms

- Secure boot, secure update and anti-rollback
- HSM, protected keys and monotonic counters
- SecOC, MACsec, TLS/DTLS and freshness
- Firewalls, access control and network segmentation
- IDS, event logging and incident response

Required coupling

- Verify source and freshness before use.
- Apply E2E, deadline and semantic checks separately.
- Do not switch dangerous actuator states directly from one authentication failure.
- Translate security failures into analysable system effects and controlled degradation.

SUMMARY E2E and SecOC are complementary: one diagnoses safety communication faults, the other resists intentional manipulation.

Hans-Leo Ross (TX-SB) | October 2026

E/E Architect's Lightboard

Safety

Mission: Prevent dangerous physical behavior.

Mechanisms:

- ECC & Lockstep
- Hardware Diagnostics
- Plausibility Checks
- MPU / MMU

Security

Mission: Protect data, identity, and access.

Mechanisms:

- Secure Boot & HSM
- SecOC & MACsec
- Firewalls
- Intrusion Detection (IDS)

Timing

Mission: Bound execution and delivery times (FTTI).

Mechanisms:

- Time-partitioned RTOS
- Deadline Supervision
- TSN Network Slots
- Bounded WCET

Health & Availability

Mission: Survive faults via redundancy.

Mechanisms:

- Independent Power Paths
- Fallback Modes
- Selective Recovery
- Diverse Sensing

Synthesis Statement: No single mechanism establishes closed-loop control. Use complementary evidence: Identity (Security) + Behavior (Safety) + Deadlines (Timing) + Redundancy (Health).

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Scale protection according to domain consequence, timing and exposure

The same vehicle platform needs different safety, security and real-time profiles by domain.

Domain	Safety criticality	Security exposure	Timing	Primary architecture objective
Chassis	Very high	Very high	Hard	Maintain controlled actuation
ADAS / AD	Very high	Very high	Hard / firm	Trust perception, planning and MRM
Body	Mixed	Medium–high	Mixed	Contain mixed criticality
HMI	High for warnings	High	Firm	Correct, timely driver interaction
Entertainment	Low direct	Very high	Soft	Contain external attack surface
V2X / OTA / Diagnostics	Contextual	Very high	Firm / soft	Trust external data and changes
Cloud	Indirect	Very high	Soft	No direct uncontrolled actuation
Power & COM networks	Enabling	High	Hard–mixed	Preserve energy and data paths

SUMMARY **Mechanism selection follows vehicle consequence, attack exposure, timing budget, fallback capability and physical redundancy.**

INTERNAL

Mission: Keep Chassis control deterministic, local and physically available

Brake, steering and propulsion require hard real-time supervision plus independent actuation paths.

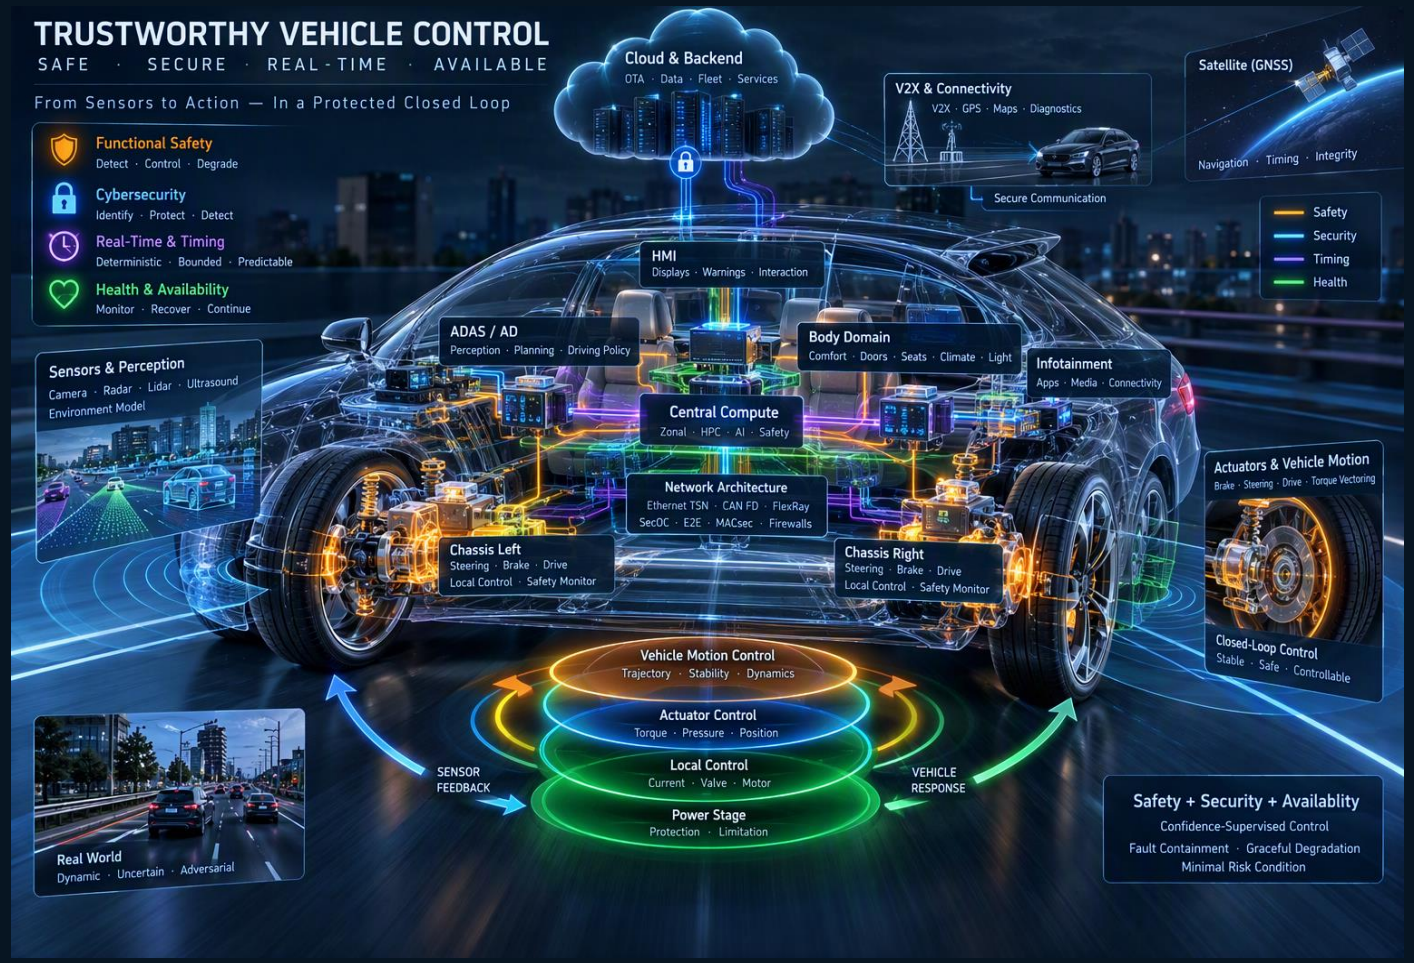

Efficient mechanism combination

- Local current, pressure, torque and position loops
- SecOC or fixed-latency MAC plus E2E and deadline checks
- Command/monitor, gradient limits and state-machine validation
- Redundant sensors, compute, communication, power and power stages
- Local arbiter, effect feedback and degraded-control fallback
- No session-dependent handshake in the fast inner loop

SUMMARY For Chassis, “fail secure” must never cause abrupt loss of steering, braking or stabilising authority.

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

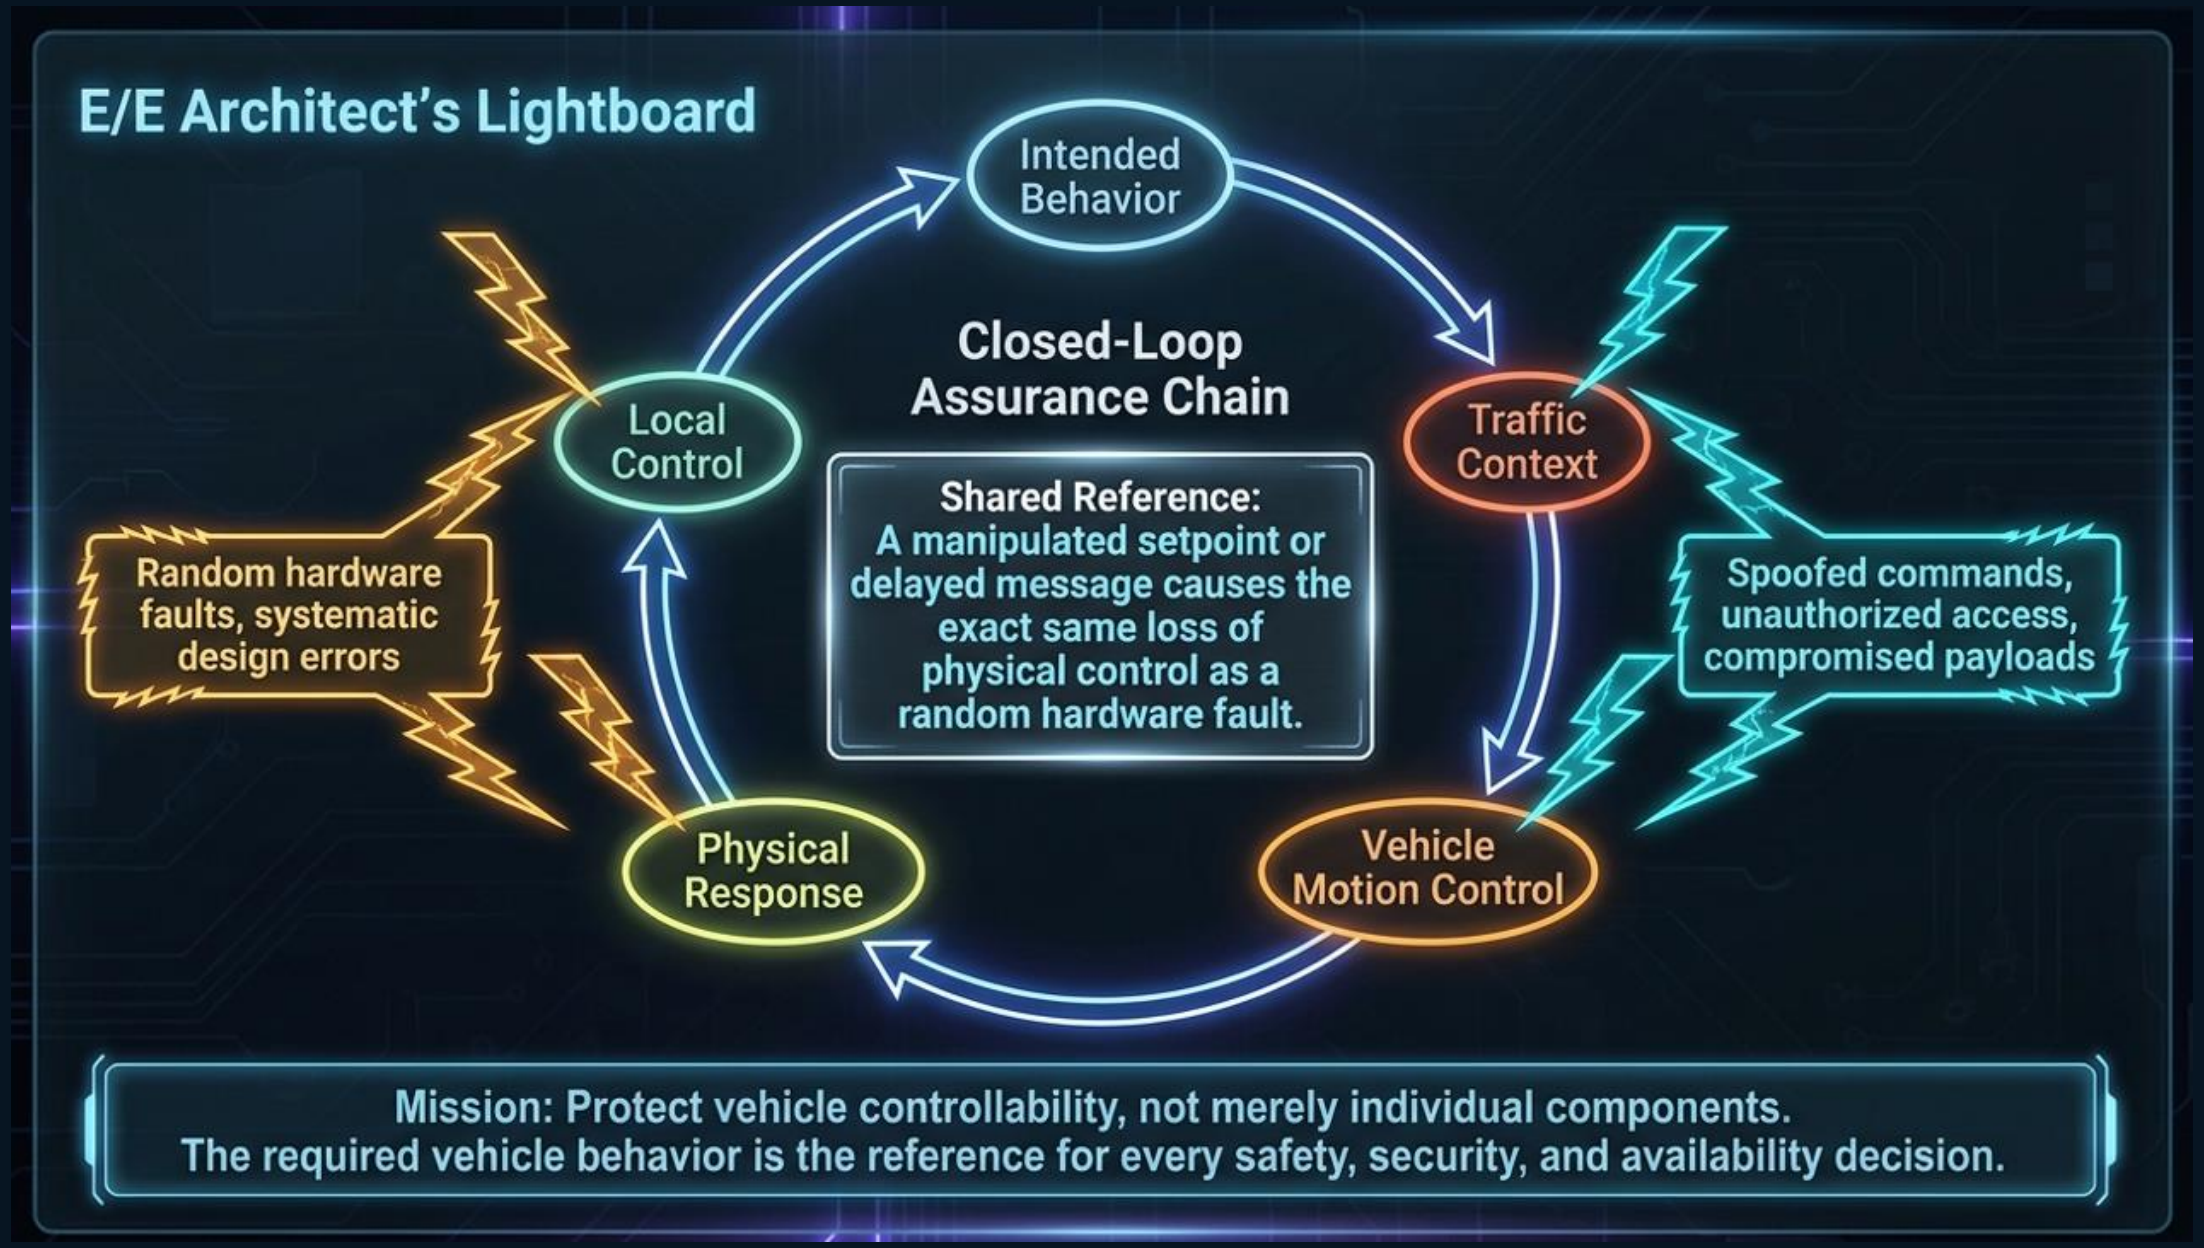

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

E/E Architect's Lightboard

ADAS / AD Domain

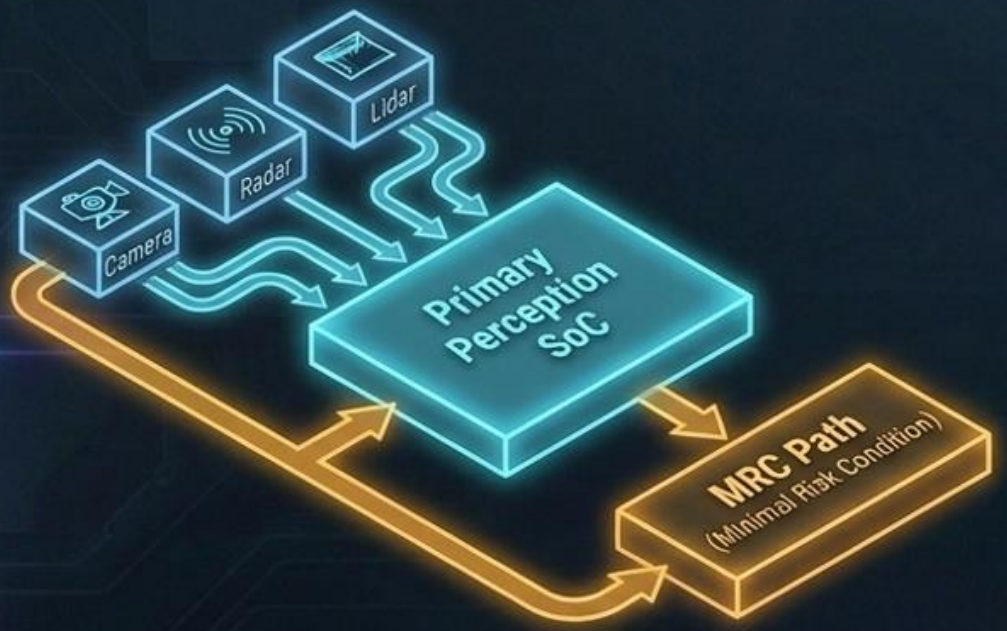

Body Domain

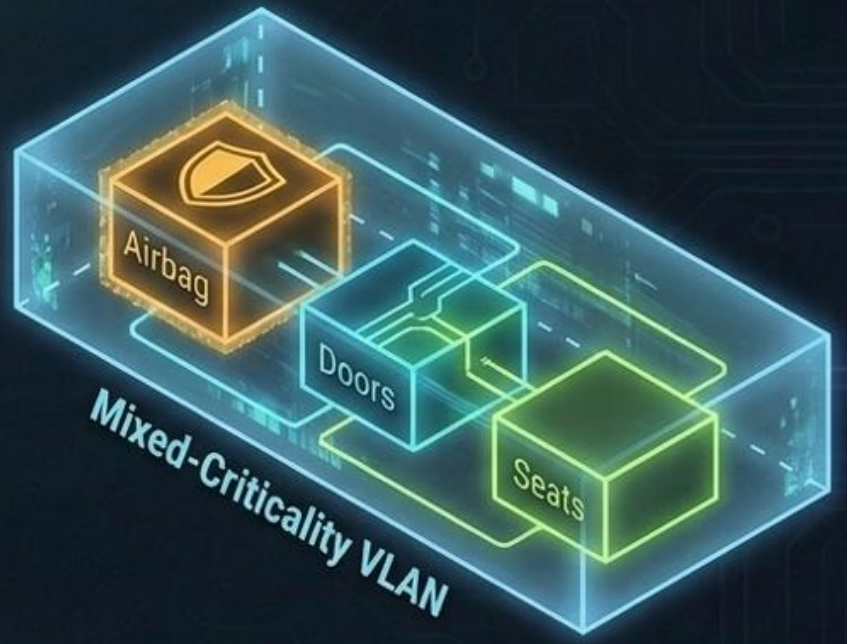

Independent Minimal Risk Condition (MRC)

The MRC path must NOT share a common-cause failure domain with the primary perception pipeline. SOTIF (ISO 21448) governs functional insufficiencies.

The Anti-Lockout Directive

Security lockouts or firewall rules must NEVER prevent emergency safety functions (airbag deployment, emergency door release). Emergency roles must be locally pre-provisioned.

INTERNAL

Mission: Make ADAS/AD trustworthy from sensing to minimal-risk manoeuvre

Perception uncertainty, cyber manipulation and compute faults converge in vehicle-control decisions.

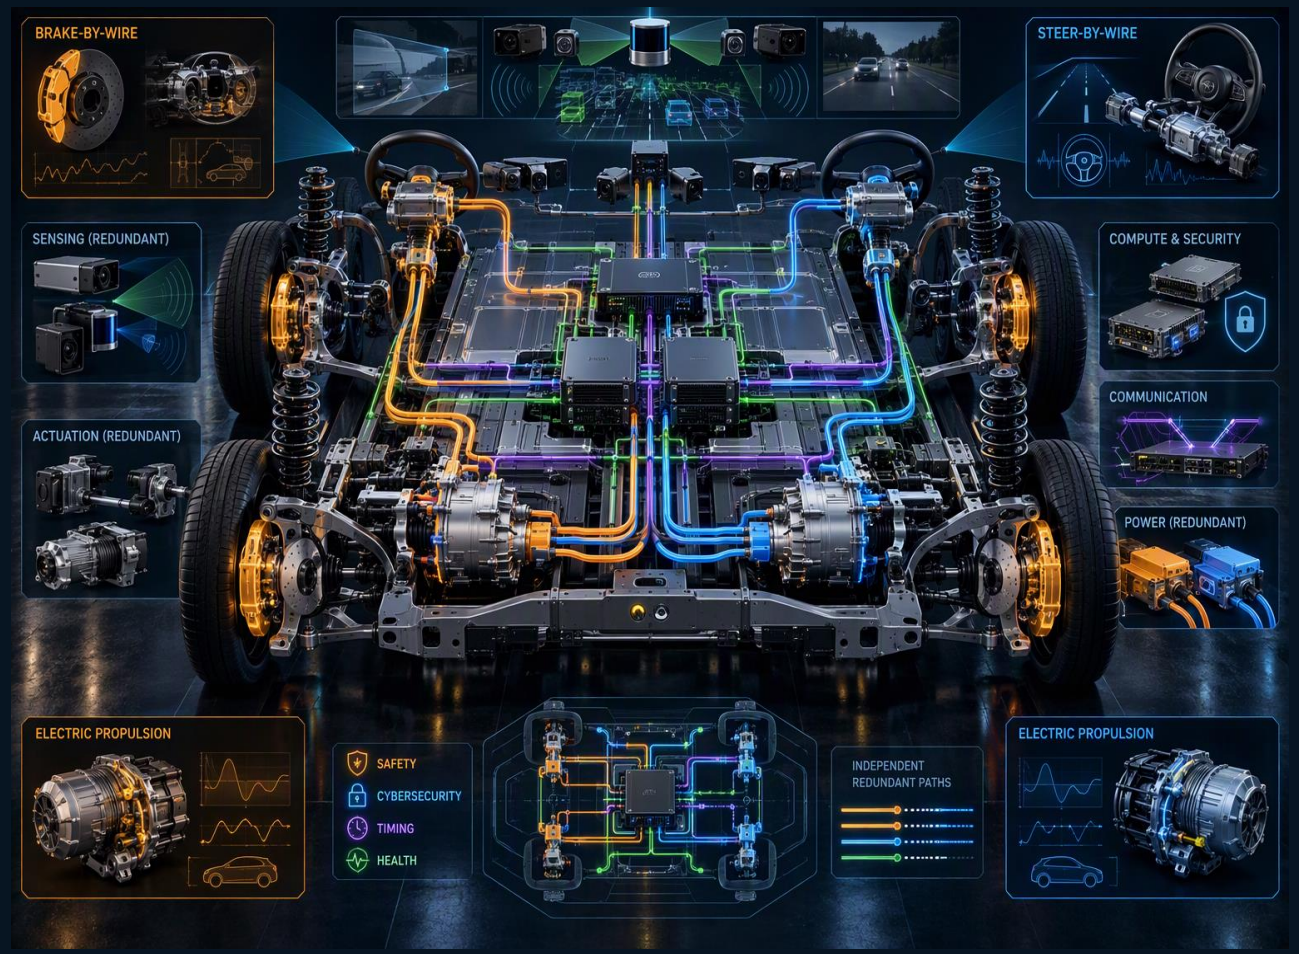

ADAS / AD assurance pattern

- Diverse sensing and independent state estimation
- Source authentication, secure sensor paths and protected calibration
- ODD and capability monitoring
- Trajectory, rule and safety-envelope verification
- Time synchronisation with quality status
- Independent fallback trajectory or MRM path
- Vehicle effect feedback through motion states and actuator residuals

SUMMARY **ADAS/AD needs both source trust and independent physical plausibility; neither cryptography nor sensor fusion alone is sufficient.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Separate mixed-criticality Body functions while preserving essential vehicle services

Airbag, doors, lighting, seats and climate require differentiated safe states and access rights.

Airbag	**Very high safety integrity**	Independent crash sensing, firing-path diagnostics, protected configuration, strict access control
Doors / access	**Safety + security**	State plausibility, anti-trap, authenticated access, mechanical emergency release
Exterior lighting	**Roadworthiness**	Lamp and supply diagnostics, command plausibility, degraded fallback patterns
Seats / restraint	**Occupant protection**	Position limits, occupancy consistency, motor current monitoring, protected coding
Climate / thermal	**Availability / asset protection**	Temperature and pressure monitoring, limited fallback, compartment isolation

SUMMARY **Body efficiency comes from shared platform services with strict partitioning and function-specific safe or degraded states.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Protect driver interaction without allowing infotainment to control the vehicle

HMI is a safety interface; entertainment is an exposed but non-authoritative trust zone.

HMI and driver information

- Timely and unambiguous warnings
- Independent safety tell-tales or protected display path
- State-consistent takeover and degradation messages
- Input validation, debouncing and authority checks
- Supervision of display availability and stale information

Entertainment and user connectivity

- Sandboxing, least privilege and process isolation
- Secure boot, signed applications and controlled updates
- TLS, certificate validation and privacy protection
- Firewall and gateway mediation to vehicle domains
- No direct write authority to Chassis or safety-critical state
- Rate-limited, schema-validated data exchange

SUMMARY **The cockpit may present vehicle state, but external applications must never become an uncontrolled command path.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Treat V2X, positioning, maps, diagnostics and OTA as untrusted external inputs

Connectivity must enrich the vehicle without becoming a single point of control or fleet-wide failure.

GNSS / positioning

Multi-source plausibility, inertial cross-check, time-quality monitoring, spoofing detection

Maps / navigation

Signed content, provenance, versioning, consistency with perception and road context

V2X

PKI, message signing, freshness, misbehaviour detection, strict application filtering

Diagnostics

Authenticated roles, least privilege, session limits, protected coding and write access

OTA / ORU

TLS transport, signed artefacts, anti-rollback, A/B image, staged release and rollback

Gateway to vehicle

Segmentation, firewall, IDS, rate limiting, protocol break and schema validation

SUMMARY **External data may support decisions, but critical control must remain locally verifiable and operationally degradable.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Make Cloud services lifecycle-safe without coupling them to the fast control loop

Cloud trust is essential for updates, fleet monitoring and services, but not for immediate actuator stability.

Cloud security baseline

- Mutual authentication and TLS
- IAM, certificate and key lifecycle
- Tenant and service isolation
- Secure APIs, secrets and event logs
- Vulnerability and incident management

Safety boundary

- No direct uncontrolled actuator access
- Local validation of cloud-derived data
- Defined loss-of-connectivity behaviour
- Version compatibility and rollback
- Field evidence linked to vehicle capability

Efficient operating model

- Perform heavy analytics and prognosis offboard.
- Keep time-critical control and fallback onboard.
- Transmit health evidence and immutable event context.
- Use staged deployment, health gates and fleet rollback.

SUMMARY **Cloud services may influence planning and lifecycle decisions, while local vehicle architecture retains time-bounded control authority.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Engineer the Power network as a physical prerequisite of digital safety

Redundant software cannot preserve vehicle control after a common loss of energy, grounding or actuation power.

Power-network hazards

- Common supply, ground or wake-up loss
- Voltage transients and brownout
- Thermal overload and load shedding
- Shared fuse, connector or distribution failure
- Incorrect power-state transition

Required mechanisms

- Independent or sufficiently decoupled supplies
- Voltage, current and temperature supervision
- ORing and selective isolation
- Prioritised load shedding
- Independent emergency energy reserve
- Separate actuator power paths

Architecture rule

- Do not claim fail-operational control from compute redundancy alone.
- Analyse internal common causes and external effects such as impact, water, heat and harness damage.
- Preserve the smallest physically effective control path.

SUMMARY **Physical redundancy is necessary because power, wiring, packaging and environment can defeat purely logical redundancy.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Provide deterministic communication without allowing security traffic to consume control deadlines

The COM network must bound latency, isolate traffic and retain trustworthy alternative paths.

CAN / CAN FD

Priority arbitration, bus fault confinement, E2E, SecOC or CAN-MAC, closed bus-load analysis

CAN XL

More payload for protection metadata, but strict frame-size and burst limits remain necessary

Ethernet TSN

Time-aware scheduling, policing, time sync, reserved bandwidth, replicated paths and duplicate elimination

MACsec / SecOC

Link authenticity and application/PDU authenticity address different termination boundaries

Diagnostics / OTA

Separated queues, rate limits and lower priority than control traffic

Network redundancy

Independent switches, links, routes and time domains where common-cause risk matters

SUMMARY **Determinism requires closed configuration, policing and reserved resources; encryption or authentication overhead belongs in WCET and FTTI budgets.**

INTERNAL

Mission: Build trust zones from vehicle authority and fault-containment boundaries

Trust zones must state who may access which resources, at what time and with what actuation authority.

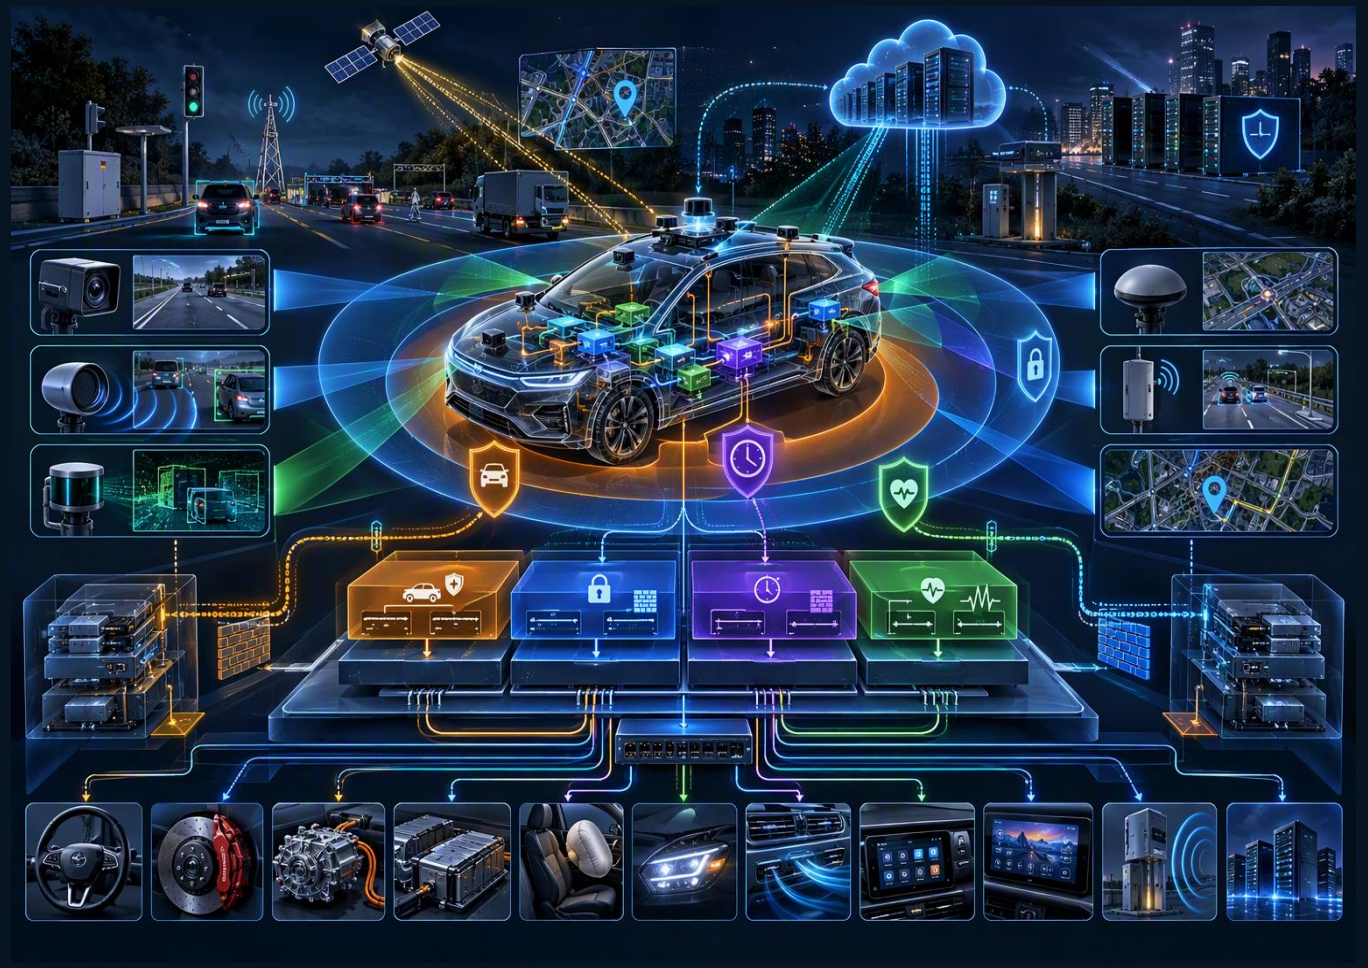

Recommended zoning logic

- Connectivity zone: externally exposed, strongly authenticated, no direct actuation
- Infotainment/HMI zone: user-facing, sandboxed, read-mostly vehicle access
- ADAS/AD zone: high-compute, time-synchronised, safety-envelope constrained
- Body zone: mixed criticality with peripheral and power containment
- Chassis zone: smallest attack surface, deterministic services, local fallback
- Independent management zone: protected health, security and recovery authority

SUMMARY **A trust zone is not merely a VLAN; it combines identity, memory, time, I/O, power and actuation boundaries.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Combine spatial, temporal and physical separation into usable fault containment

Isolation is effective only when it protects memory, time, shared resources and the physical effect path.

SPATIAL

MPU/MMU, IOMMU, process isolation, protected memory, peripheral and DMA rights

Where may the function act?

TEMPORAL

CPU and ISR budgets, TSN slots, HSM quotas, memory-bandwidth and interrupt policing

When and for how long may it act?

LOGICAL

Alive, deadline and logical supervision; control-flow checkpoints and mode consistency

Did the required protection path execute?

PHYSICAL

Separate power, sensing, communication, compute and actuation paths; packaging and environmental independence

Can a real vehicle effect still be produced?

SUMMARY **Fault containment needs four coordinated dimensions: spatial, temporal, logical and physical independence.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Recover at the smallest effective fault-containment region

High availability is achieved by preserving healthy control paths, not by resetting the largest subsystem.

Decision evidence

- Fault scope
- Timing margin
- Safety confidence
- Security confidence
- Remaining sensors
- Remaining compute
- Remaining communication
- Physical actuation authority

SUMMARY **Detection without precise containment can reduce availability and create a larger vehicle-level hazard.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Structure the Chassis MCU around deterministic execution, independent evidence and local recovery

A high-availability MCU must offer separation, observation, synchronisation, containment and recovery as platform capabilities.

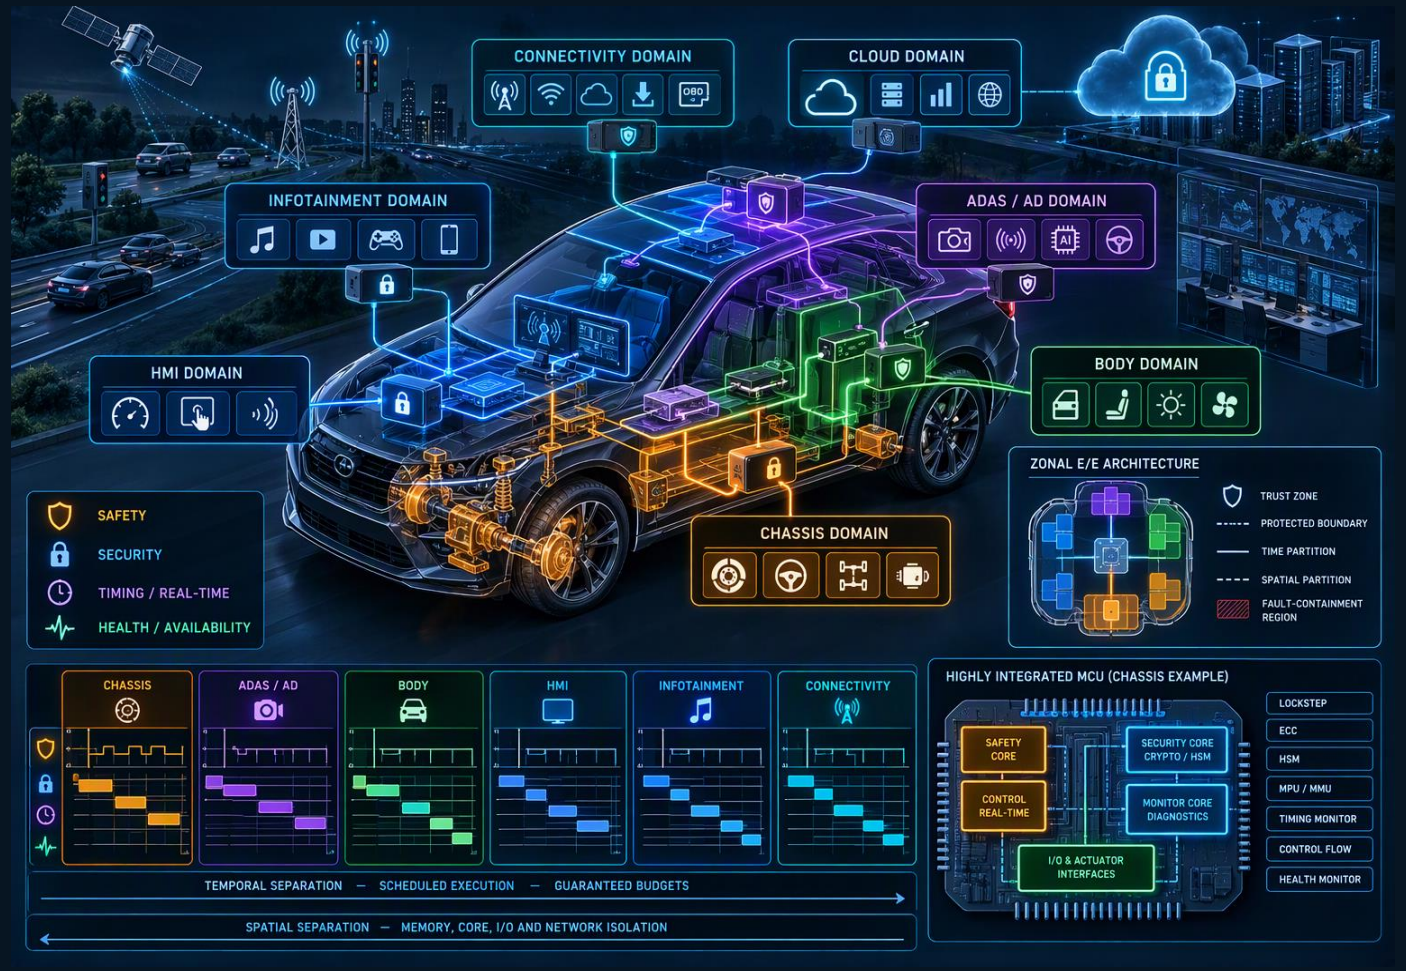

Target MCU structure

- Diverse Core A: full-performance control
- Diverse Core B: delayed safety-oriented control or independent monitor
- Independent meta-monitor with protected time base
- TCM / scratchpad and ECC-protected memories
- MPU plus DMA/IOMMU access protection
- Deterministic on-chip interconnect and bandwidth budgets
- Partitioned HSM with priority crypto queues
- Local resets, clock domains and alarm routing
- Qualified watchdog as final escalation only

SUMMARY **Lockstep detects selected failures; fail-operational control additionally requires independent execution, arbitration and recoverable fault-containment regions.**

INTERNAL

Mission: Preserve temporal control through physically independent Chassis actuation

Internal and external effects require physical redundancy beyond software, cores and logical network separation.

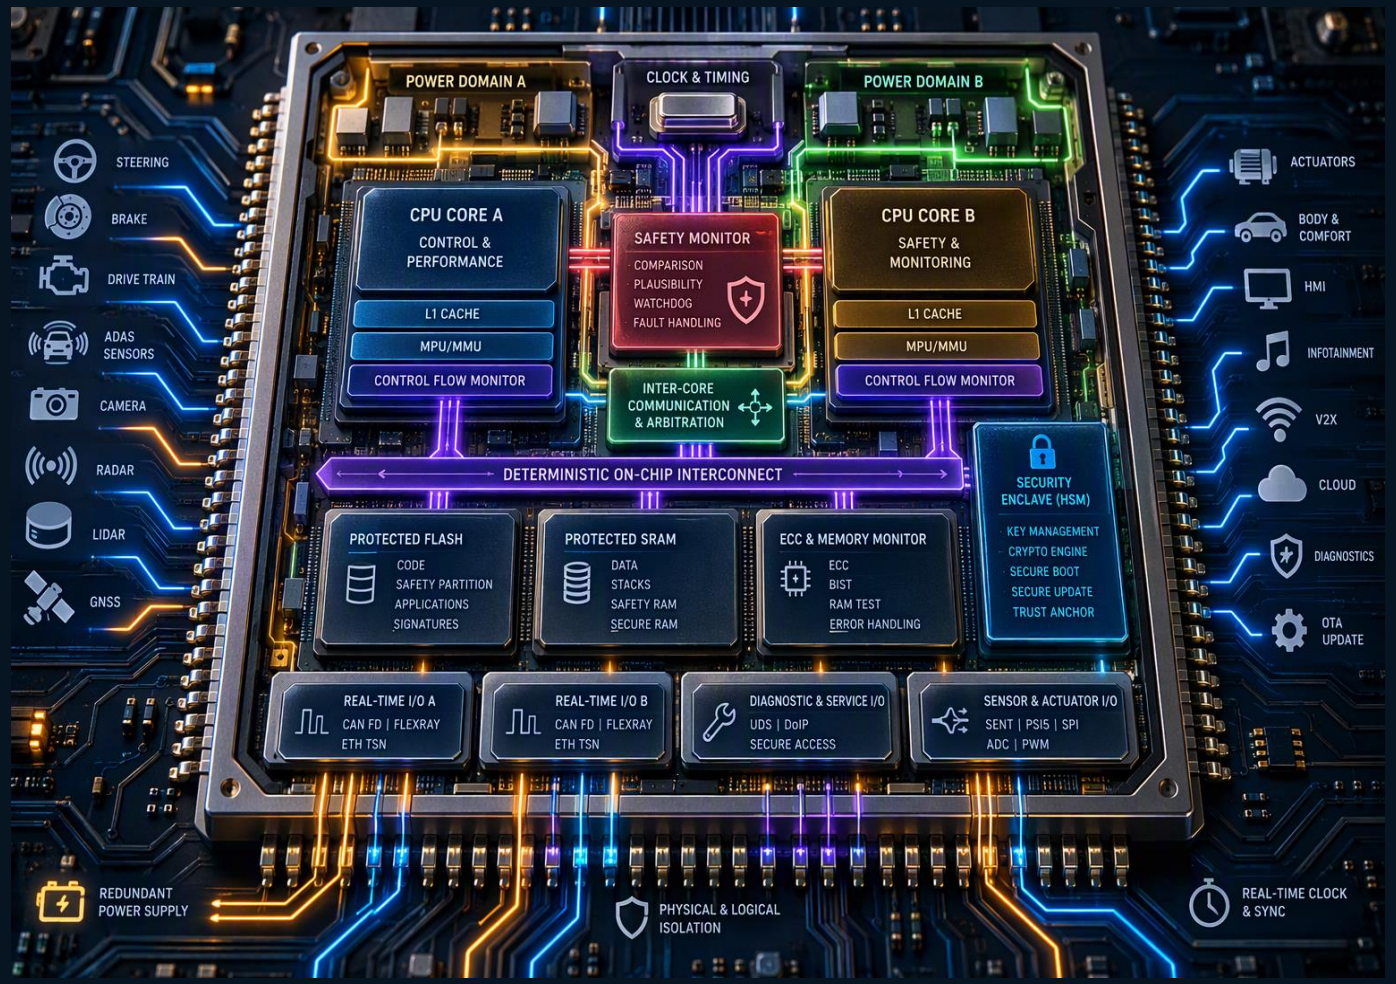

Necessary independence

- Independent sensors and reference supplies
- Independent power and grounding paths
- Independent compute and trust domains
- Independent communication paths or replicated transport
- Independent gate drivers, inverters and motor windings
- Mechanical isolation or controlled interaction
- Packaging against common impact, heat, water and harness damage
- Independent effect feedback from vehicle motion

Physical redundancy converts detection into continued, time-bounded control.

SUMMARY **The decisive metric is remaining physical steering, braking and propulsion authority after an internal fault or external event.**

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

The Limits of Logical Partitioning

The Illusion of Safety

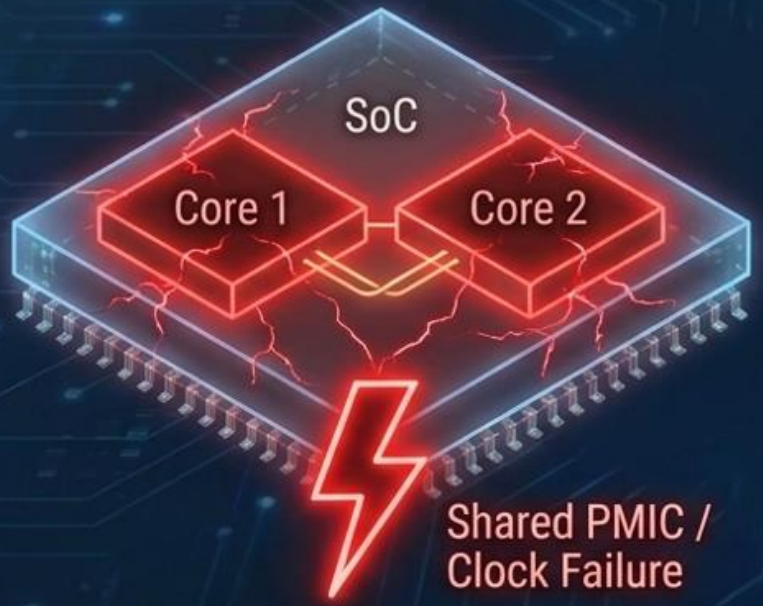

Lockstep detects.
It does **NOT** preserve control.

True Availability

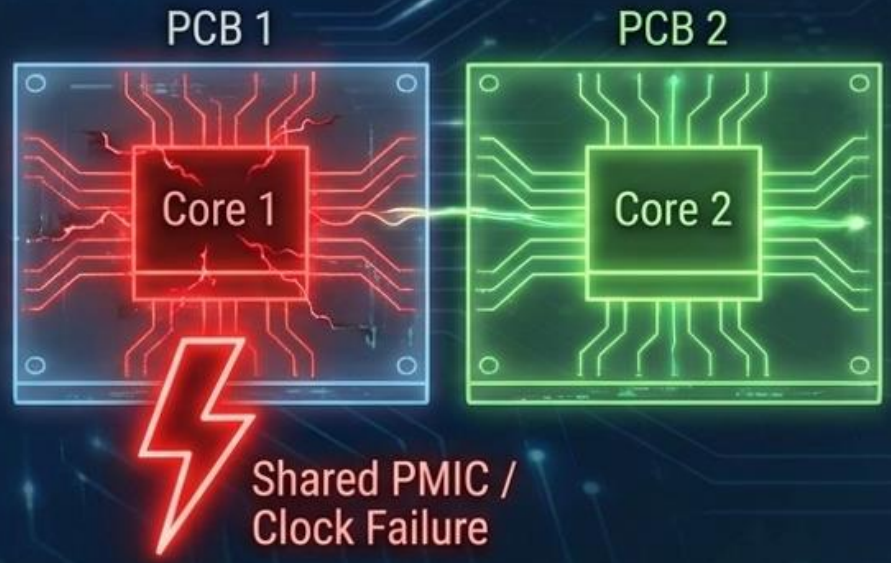

A fail-operational chassis cannot survive common-cause faults (CCF) from shared power, thermal events, or electromagnetic interference within a single SoC.

INTERNAL

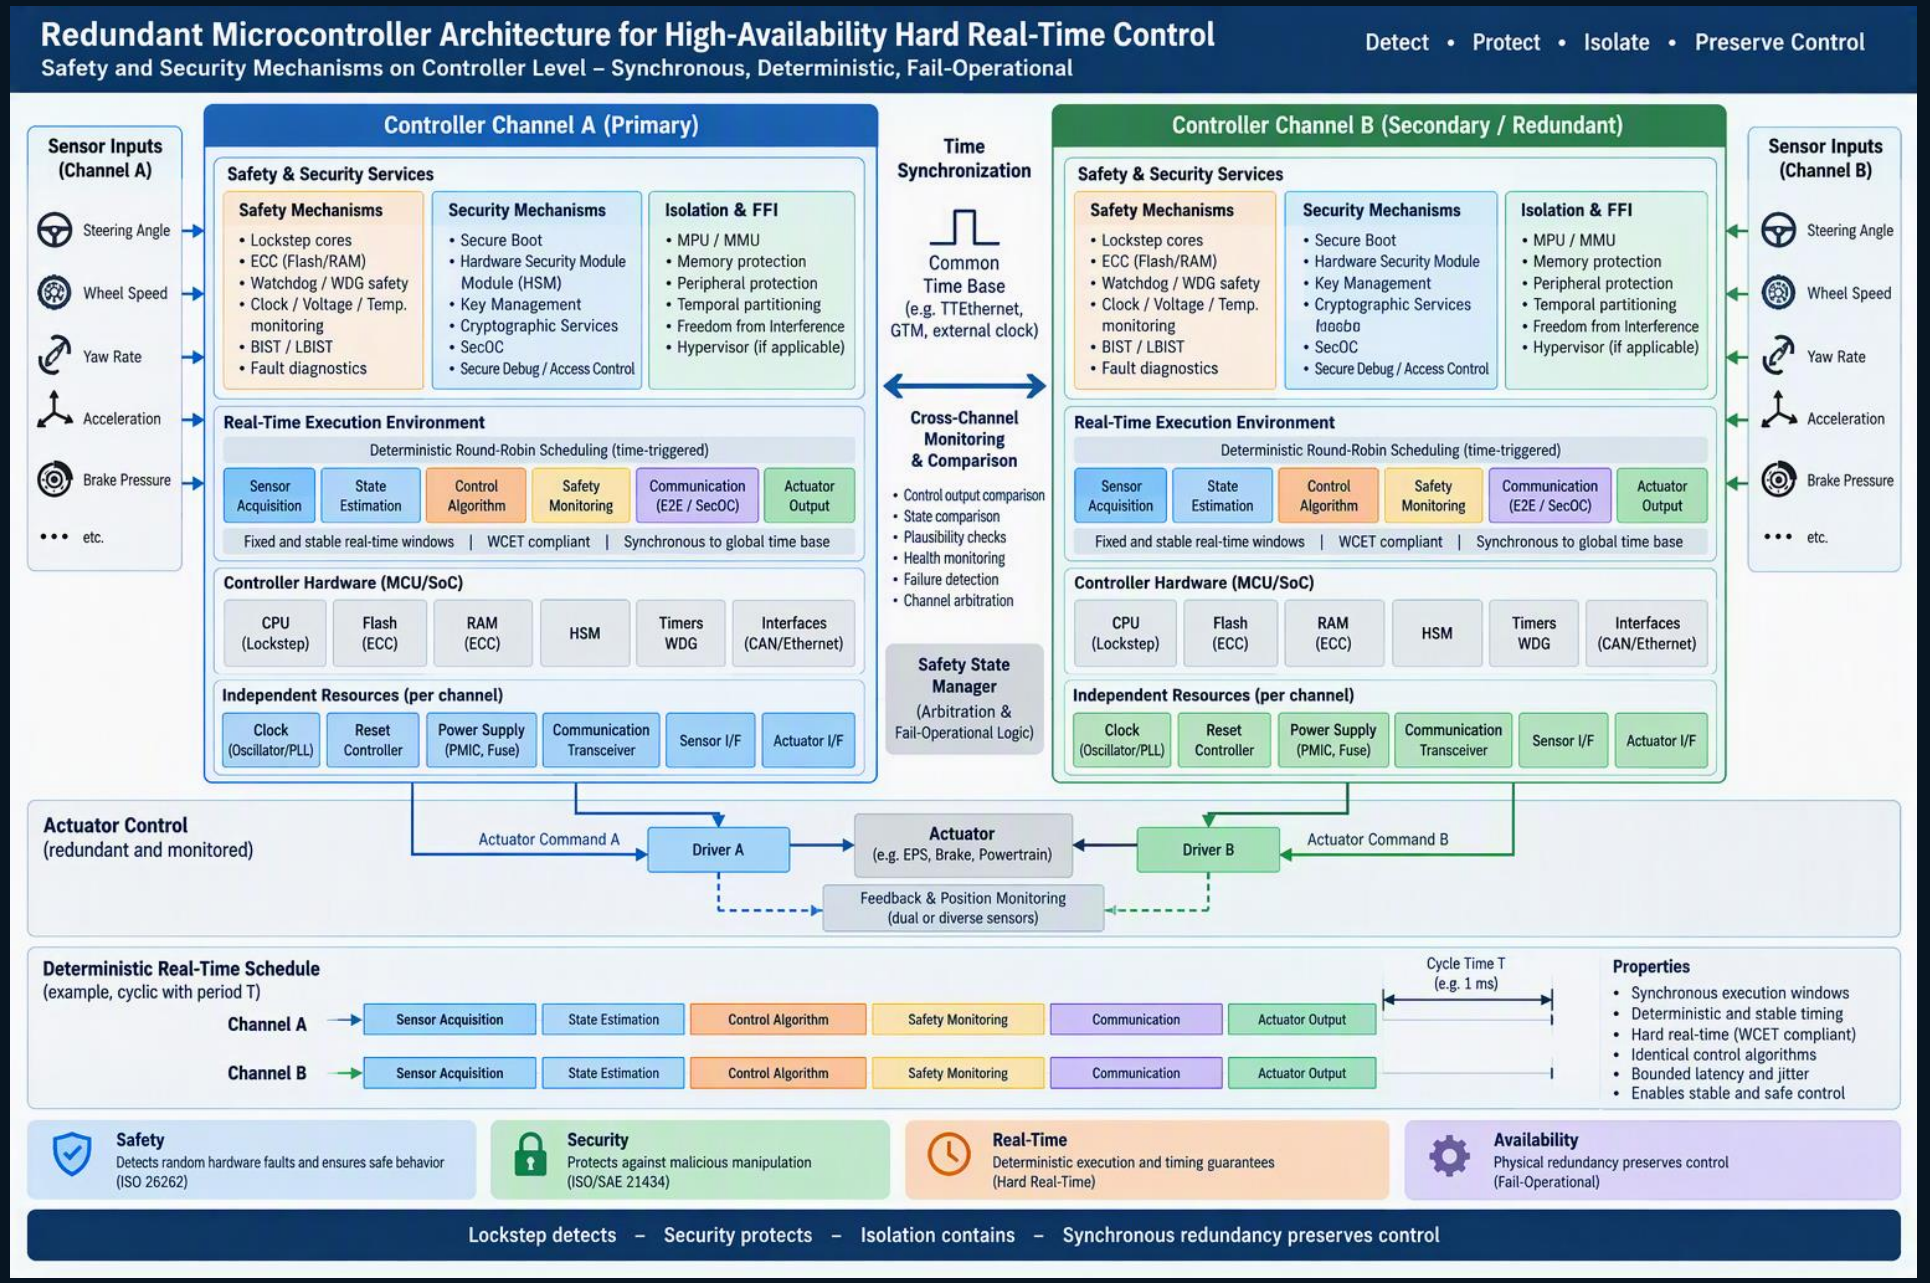
Redundant Microcontroller Architecture for High-Availability Hard Real-Time Control
Safety and Security Mechanisms on Controller Level – Synchronous, Deterministic, Fail-Operational
Detect • Protect • Isolate • Preserve Control
Sensor Inputs (Channel A)
Steering Angle
Wheel Speed
Yaw Rate
Acceleration
Brake Pressure
etc.
Controller Channel A (Primary)
Safety & Security Services
Safety Mechanisms
• Lockstep cores
• ECC (Flash/RAM)
• Watchdog / WDG safety
• Clock / Voltage / Temp. monitoring
• BIST / LBIST
• Fault diagnostics
Security Mechanisms
• Secure Boot
• Hardware Security Module Module (HSM)
• Key Management
• Cryptographic Services
• SecOC
• Secure Debug / Access Control
Isolation & FFI
• MPU / MMU
• Memory protection
• Peripheral protection
• Temporal partitioning
• Freedom from Interference
• Hypervisor (if applicable)
Real-Time Execution Environment
Deterministic Round-Robin Scheduling (time-triggered)
Sensor Acquisition
State Estimation
Control Algorithm
Safety Monitoring
Communication (E2E / SecOC)
Actuator Output
Fixed and stable real-time windows | WCET compliant | Synchronous to global time base
Controller Hardware (MCU/SoC)
CPU (Lockstep)
Flash (ECC)
RAM (ECC)
HSM
Timers WDG
Interfaces (CAN/Ethernet)
Independent Resources (per channel)
Clock (Oscillator/PLL)
Reset Controller
Power Supply (PMIC, Fuse)
Communication Transceiver
Sensor I/F
Actuator I/F
Time Synchronization
Common Time Base (e.g. TTEthernet, GTM, external clock)
Cross-Channel Monitoring & Comparison
• Control output comparison
• State comparison
• Plausibility checks
• Health monitoring
• Failure detection
• Channel arbitration
Safety State Manager
(Arbitration & Fail-Operational Logic)
Controller Channel B (Secondary / Redundant)
Safety & Security Services
Safety Mechanisms
• Lockstep cores
• ECC (Flash/RAM)
• Watchdog / WDG safety
• Clock / Voltage / Temp. monitoring
• BIST / LBIST
• Fault diagnostics
Security Mechanisms
• Secure Boot
• Hardware Security Module
• Key Management
• Cryptographic Services
• SecOC
• Secure Debug / Access Control
Isolation & FFI
• MPU / MMU
• Memory protection
• Peripheral protection
• Temporal partitioning
• Freedom from Interference
• Hypervisor (if applicable)
Real-Time Execution Environment
Deterministic Round-Robin Scheduling (time-triggered)
Sensor Acquisition
State Estimation
Control Algorithm
Safety Monitoring
Communication (E2E / SecOC)
Actuator Output
Fixed and stable real-time windows | WCET compliant | Synchronous to global time base
Controller Hardware (MCU/SoC)
CPU (Lockstep)
Flash (ECC)
RAM (ECC)
HSM
Timers WDG
Interfaces (CAN/Ethernet)
Independent Resources (per channel)
Clock (Oscillator/PLL)
Reset Controller
Power Supply (PMIC, Fuse)
Communication Transceiver
Sensor I/F
Actuator I/F
Sensor Inputs (Channel B)
Steering Angle
Wheel Speed
Yaw Rate
Acceleration
Brake Pressure
etc.
Actuator Control
(redundant and monitored)
Actuator Command A
Driver A
Actuator
(e.g. EPS, Brake, Powertrain)
Driver B
Actuator Command B
Feedback & Position Monitoring
(dual or diverse sensors)
Deterministic Real-Time Schedule
(example, cyclic with period T)
Cycle Time T
(e.g. 1 ms)
Channel A
Sensor Acquisition
State Estimation
Control Algorithm
Safety Monitoring
Communication
Actuator Output
Channel B
Sensor Acquisition
State Estimation
Control Algorithm
Safety Monitoring
Communication
Actuator Output
Properties
• Synchronous execution windows
• Deterministic and stable timing
• Hard real-time (WCET compliant)
• Identical control algorithms
• Bounded latency and jitter
• Enables stable and safe control
Safety
Detects random hardware faults and ensures safe behavior (ISO 26262)
Security
Protects against malicious manipulation (ISO/SAE 21434)
Real-Time
Deterministic execution and timing guarantees (Hard Real-Time)
Availability
Physical redundancy preserves control (Fail-Operational)
Lockstep detects – Security protects – Isolation contains – Synchronous redundancy preserves control

INTERNAL

Hans-Leo Ross (TX-SB) | October 2026

Mission: Adopt ten integrated design rules for trustworthy vehicle control

The final architecture is a coordinated control-and-assurance system, not a stack of independent mechanisms.

01 Start from vehicle capability and controllability.

02 Allocate mechanisms according to loop speed and consequence.

03 Use spatial and temporal separation as the basis of containment.

04 Preserve physical independence where common-cause effects matter.

05 Keep safety, security, timing and health evidence distinguishable.

06 Protect monitors and confidence evidence from the evaluated function.

07 Authenticate commands, then validate state, limits and physical plausibility.

08 Recover at the smallest effective fault-containment region.

09 Use watchdog reset or shutdown only as the final escalation.

10 Validate the combined mechanism chain in a time-accurate physical closed loop.

SUMMARY Separate → observe → build evidence → contain locally → reconfigure → preserve control → degrade only when required.

INTERNAL