technica
engineering
Member of KPIT Group

THE MODERN IN-VEHICLE NETWORK VALIDATION PLAYBOOK(2026)

HOW VALIDATION TEAMS DEBUG SECURE, SYNCHRONIZED & OBSERVABLE AUTOMOTIVE ETHERNET TEST BENCHES, WITHOUT BREAKING THEM

What to expect and who this playbook is for

Automotive MACsec is increasingly mandatory in the development of vehicles, not just at Start of Production (SOP). While this improves cybersecurity, it makes traditional logging, tapping, rest-bus simulation, and diagnostics, unusable without breaking the secure link.

MACsec-protected ethernet links are designed for endpoints that actively participate in trust establishment. Passive taps and external loggers, which do not engage in the MACsec key agreement (MKA), lose the possibility to inject traffic and could even lose visibility if encryption is activated. Although limited observation may still be possible with certain tap devices, any scenario that requires establishing a MACsec/MKA session with a new active participant or decrypting traffic for analysis, requires the use of an in-path device capable of secure communication participation.

Shared time is the second foundation of modern validation, and often, what benches end up accidentally disturbing.

Distributed ECUs and sensors rely on a common time base (e.g., gPTP/802.1AS or PTPv2)[1] so events can be correlated, ordered, and measured across the system. Even when time seems to work, subtle issues—asymmetry, added delay, role changes, or mixed profiles—can introduce silent drift that produces repeatable but incorrect results. Well-designed Capture Modules (CMs) and time-aware taps can preserve timing while providing visibility. However, any instrumentation that introduces asymmetry, additional delay, or unintended role changes can silently degrade synchronization accuracy. This is why timing and observability must be designed together in time-aware networks.

Modern IVN validation succeeds when three realities are handled together:
Security, shared time, and non-disruptive observability.

Note on networks: Many labs operate two interconnected networks: (1) the In-vehicle Network (IVN), where ECUs communicate and where MACsec and gPTP apply; (2) the measurement / tool network between Capture Modules and tools, where loggers and lab clocks are connected and PTPv2 time sources are common. In development, benches often bridge timing and visibility between these contexts (PTPv2 ⟷ gPTP at the boundary).

This playbook provides an overview-level path to achieve all three without disturbing your bench topology, so you can debug faster, prove results, and move confidently to production.

[1] gPTP stands for generalized Precision Time Protocol; and PTPv2 stands for Precision Time Protocol version 2

This playbook is written for:

- IVN architects
- Validation teams/leads /engineers/architects
- HIL leads/HIL bench owners
- ADAS integration managers and cybersecurity teams
- Engineering leaders

This playbook reflects recurring validation patterns observed by subject matter experts at Technica Engineering, across multiple Automotive Ethernet programs. It focuses on validation architectural decisions, not tool configuration or step-by-step instructions.

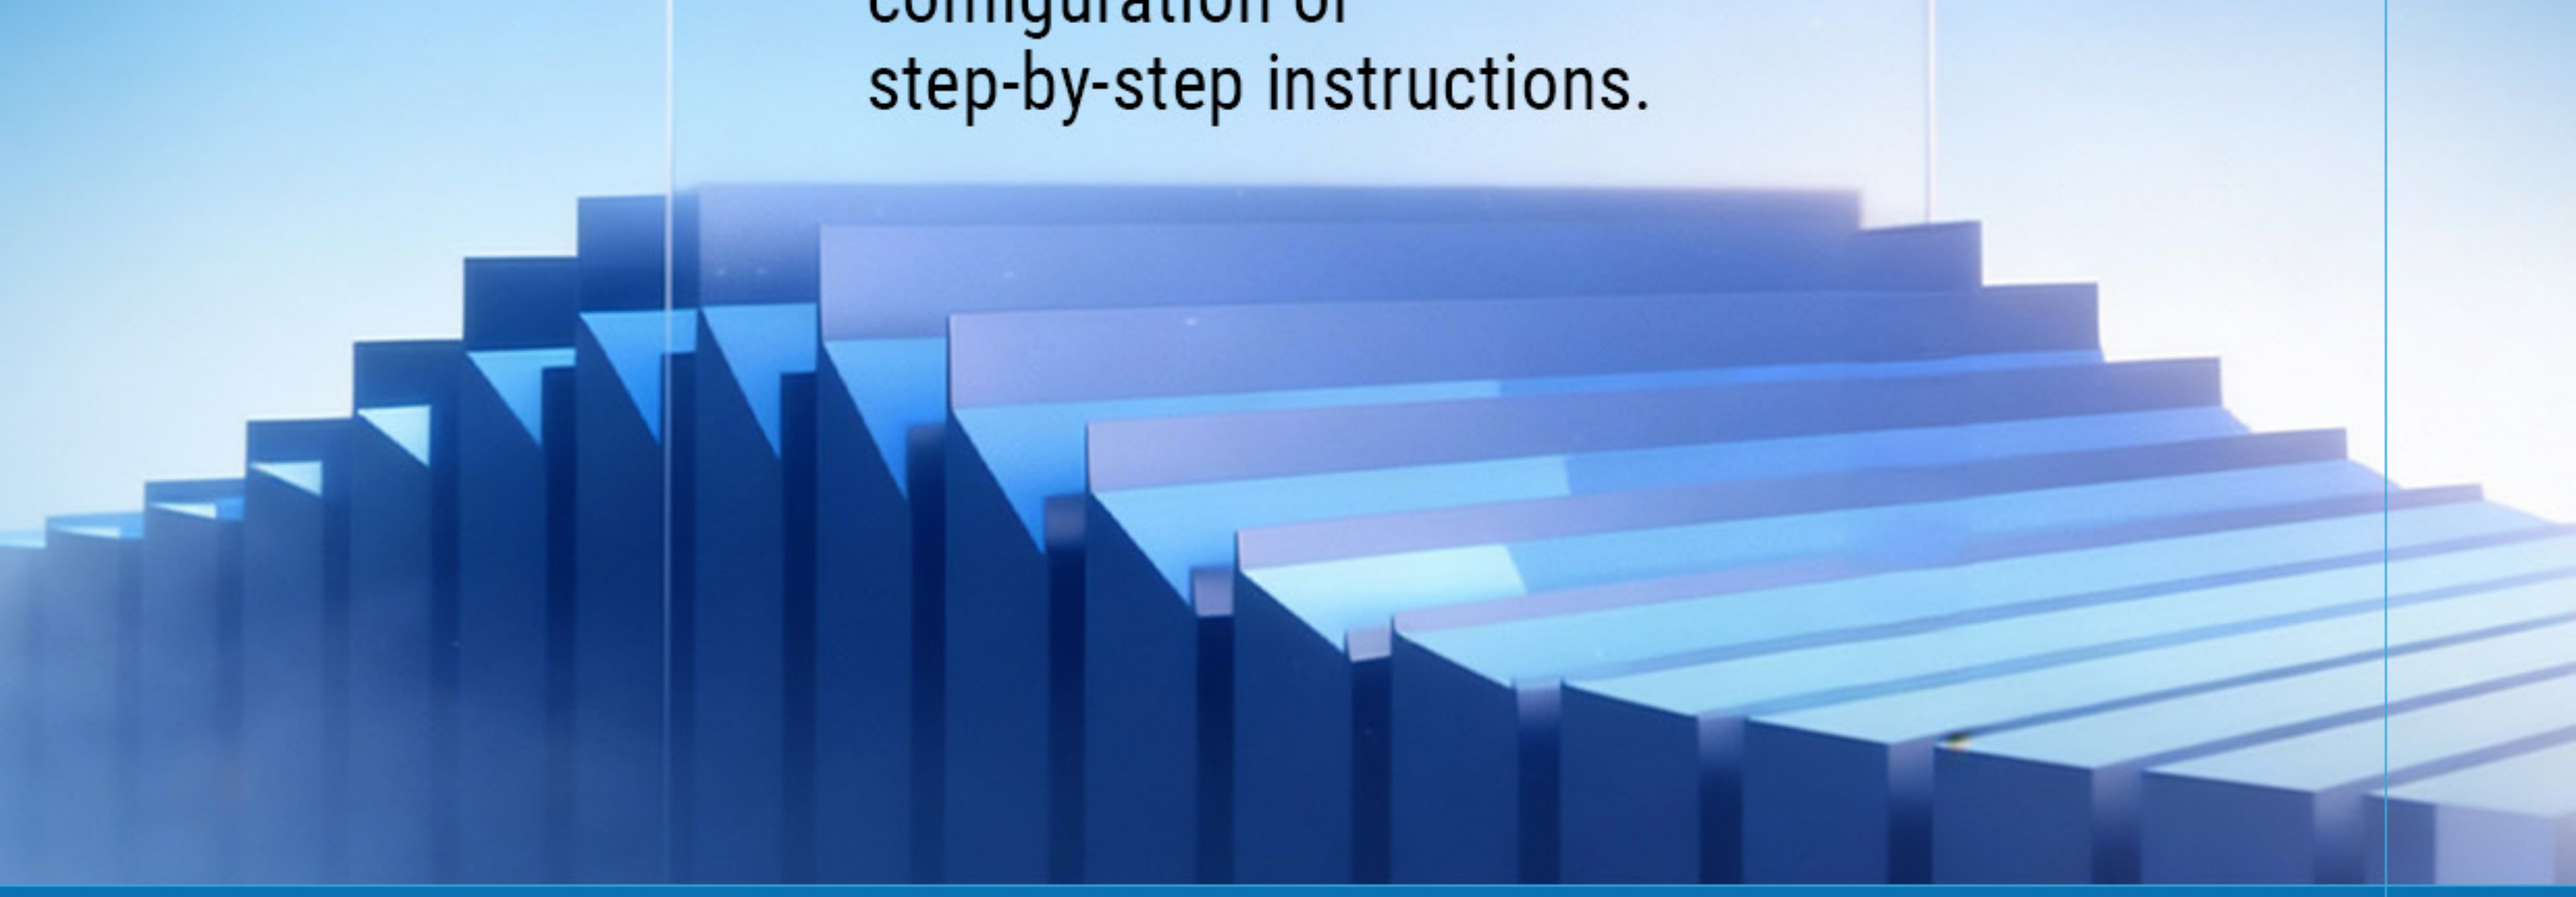

Index

Why timing & security matter (before you wire)

technica engineering
Member of KPIT Group

In modern Automotive Ethernet benches, timing, security, and observability cannot be treated independently.

Security

Automotive MACsec provides data integrity, origin authentication and optionally data confidentiality in production-grade links, but can hide the very payload you must analyze.

Timing

Shared time ensures sensors and ECUs operate on a common clock; without it, time-sync and latency checks fail.

Observability

Observability must preserve security and timing behavior; gain insight via a mirrored analysis interface without weakening security or disturbing timing roles.

The missing piece in modern IVN validation

technica engineering
Member of KPIT Group

Most validation benches were built incrementally. Tools were added to solve individual problems:

- A tap for visibility
- A converter for media mismatch
- A timing device for sync
- Scripts to glue everything together

What was missing was a single control point inside the network. Validation tools were added to solve individual problems, but the overall system behaviour of the bench was never designed as a whole.

Why IVN validation environments break

Automotive Ethernet was meant to simplify in-vehicle networks. In validation labs, the opposite often happens. Modern IVN benches must now deal with:

- **MACsec-protected links** that block payload inspection and rest bus simulation for non-participating tools
- **Time-aware networks (gPTP)** that can lose accuracy when instrumented without time-aware design (e.g., added delay/asymmetry)
- **Mixed timing domains** (PTPv2 ⟷ gPTP)
- **High-bandwidth sensors** (camera, radar, LiDAR) that stress logging and analysis pipelines
- **Multi-sensor timestamp alignment (<1 µs)** - ensuring that camera, radar, and LiDAR data streams can be correlated accurately
- **Simulation routing,** fault injection, selective filtering
- **Hybrid T1 + RJ45 architectures** and media mismatches

When security, timing, and observability are handled independently, their interactions create failure modes that are difficult to detect and harder to diagnose.

Capture Modules (CMs) and taps remain a proven approach for high-fidelity visibility — CMs were created to provide a constant minimal delay and minimize the impact in synchronizated links. The missing piece addressed here is an in-path control point for cases where secure participation (MACsec), time-aware timing, and controlled traffic experimentation must be preserved together.

Introducing Enhanced Ethernet Switches (EES)

technica engineering
Member of KPIT Group

Enhanced Ethernet Switches (EES) are Layer 2 managed Automotive Ethernet switches designed for validation environments. Unlike traditional switches or external tools that sit beside the network, an EES is placed inside the communication path, in this playbook, EES is used as the reference platform to illustrate how modern validation benches can:

- Participate in MACsec-protected communication
- Preserve time synchronization (gPTP / PTPv2)
- Control traffic behaviour

without disturbing ECUs or network topology.

Rather than observing the network from the outside, control is introduced inside the communication path, allowing validation without altering ECU behaviour.

This playbook uses EES as the reference platform in modern IVN benches to illustrate:

- Secure observability
- Deterministic timing
- Controlled experimentation

How benches typically evolve

Tool driven benches	EES driven benches
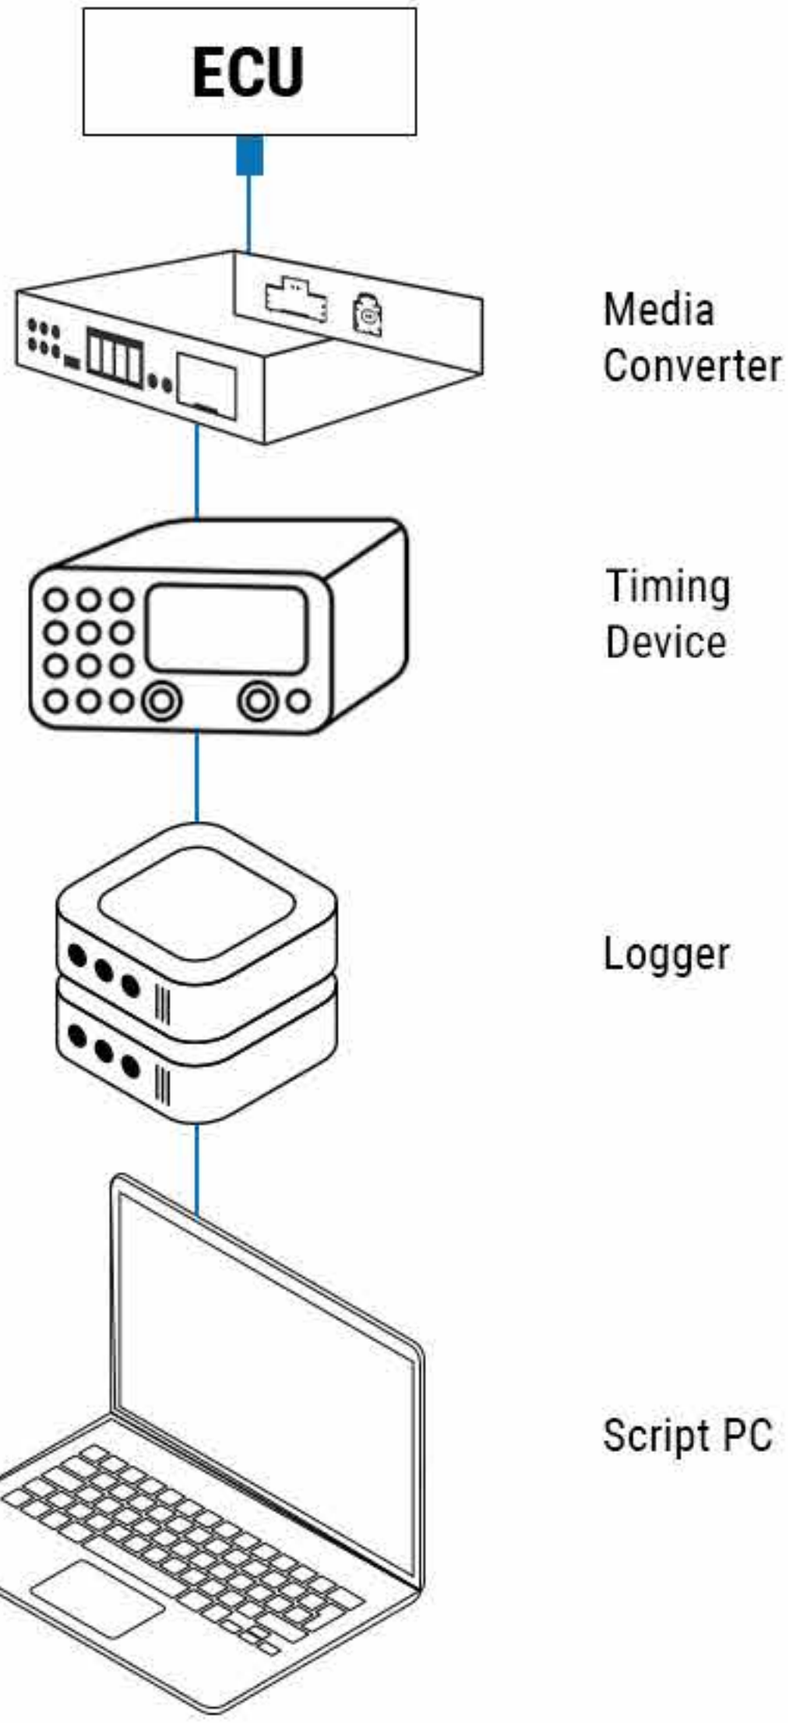	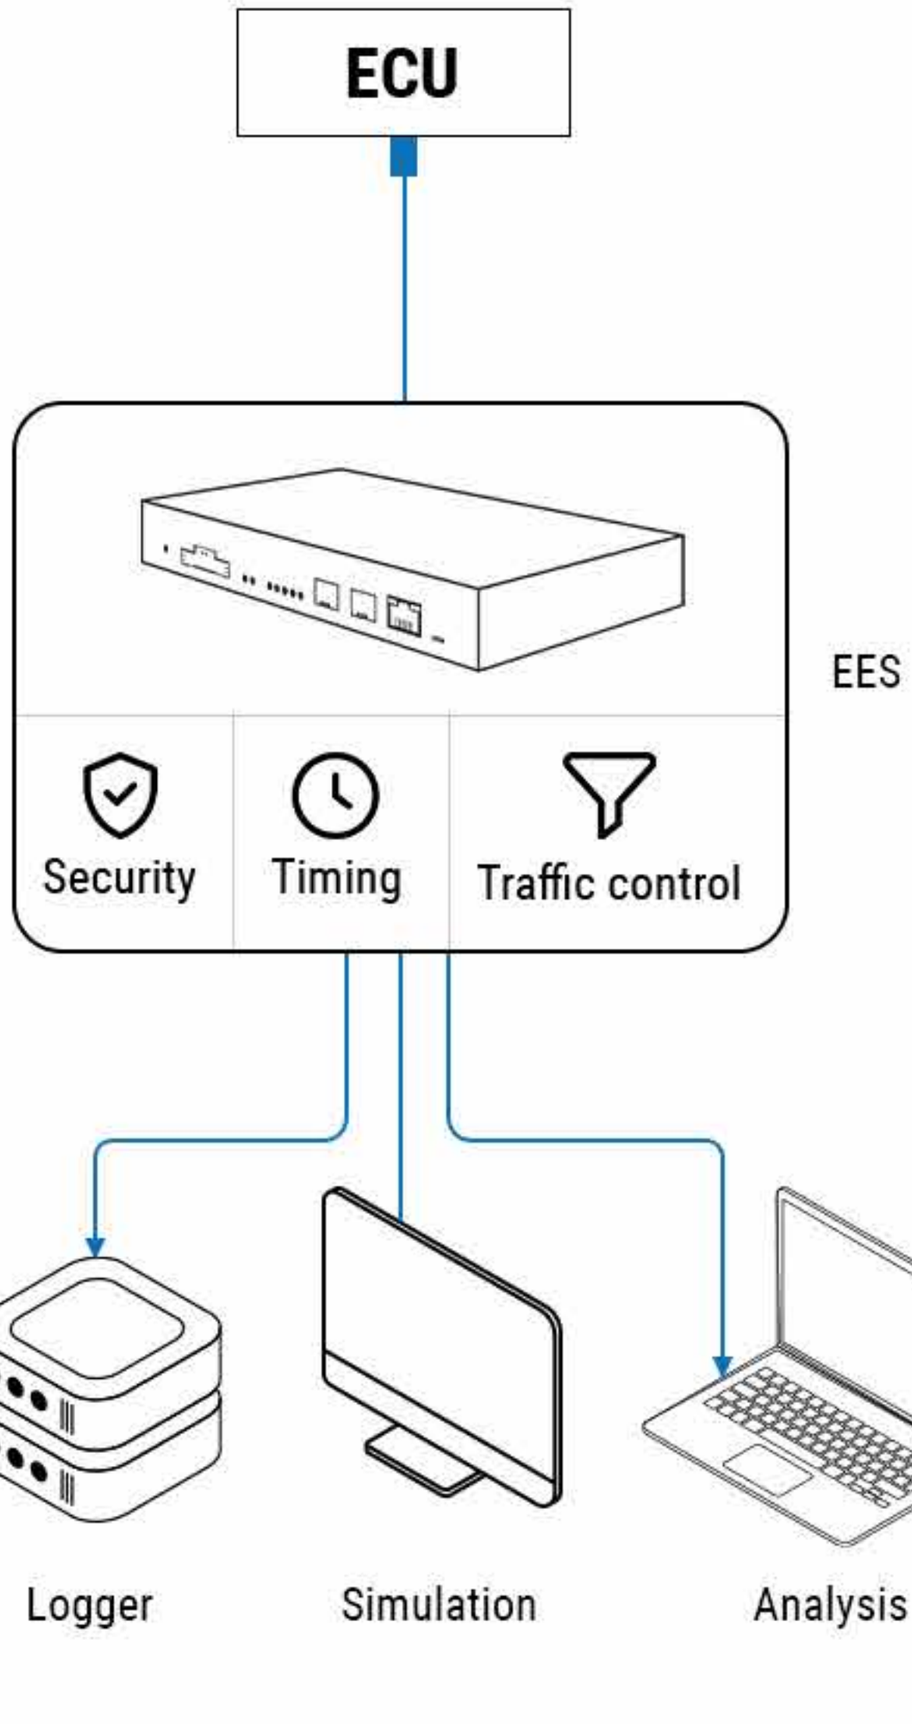
⚠ Each new requirement adds another failure point	✓ Same validation intent ✓ Fewer tools ✓ Predictable behaviour

The difference is Layer-2 control inside the bench.

Timing without breaking the chain

technica engineering
Member of KPIT Group

Modern in-vehicle systems depend on precise time alignment of raw sensor streams:

- Camera frames
- Radar measurements
- Other time-stamped sensor data, combined through sensor-fusion algorithms

The challenge: Time-aware links cannot be tapped casually.

In gPTP networks:

- Master and slave roles are defined within a time synchronization domain.
- A single Grandmaster exists per domain, with multiple slave devices synchronizing to it.
- Topology changes affect synchronization accuracy.

A hidden nightmare

In many cases, synchronization does not fail outright. Instead, accuracy degrades silently, producing repeatable but incorrect results.

Multi-partition timing: Observing without disturbing

technica engineering
Member of KPIT Group

Creative Visualization

What is an EES time partition?

A partition is an isolated timing domain inside the switch.

Each partition:

- Has **exactly one slave port**
- Can forward timing to multiple master ports
- Is fully isolated from other partitions

Why do partitions matter in validation?

Partitions allow engineers to:

- Tap **both sides** of a timing-critical link
- Preserve original timing roles
- Mirror traffic safely
- Keep ECUs unaware of test equipment

Multi-partition timing: observing without disturbing

technica engineering
Member of KPIT Group

Example: Tapping both sides of a LiDAR-ECU link

Partition 1	Slave: LiDAR Masters: ECU
Partition 2	Slave: CM Master: Logger

Outcome:

- Original timing roles are preserved on both sides of the link.
- Traffic is mirrored safely for analysis.
- Neither ECU detects the presence of test equipment or altered network behavior.

Mixed timing domains in development

technica engineering
Member of KPIT Group

Development benches rarely start clean. Many lab-grade Grandmasters speak **PTPv2.** Automotive ECUs expect **gPTP.** Direct connection is not possible.

The bridging pattern:

PTPv2 Grandmaster ⟶ EES (Timing Bridge) ⟶ gPTP ECUs

Context note: This pattern is commonly used at the boundary between the Measurement/Tool Network (PTPv2 time source) and the In-Vehicle Network (gPTP domain). MACsec remains an IVN feature; timing bridging connects lab reference clocks to the IVN during development.

In this pattern, the EES acts as a timing bridge, consuming PTPv2 on one side and distributing gPTP on the other. This allows lab-grade clocks and existing infrastructure to be reused during development.

This pattern enables:

- Reuse of high-precision lab clocks
- Stable mixed-domain development setups
- Gradual architecture convergence

Mixed timing domains are common in development and validation due to legacy infrastructure and labgrade reference clocks. This setup is recommended for development and validation environments. Production (SOP) architectures should avoid mixed timing domains wherever possible. When mixed timing cannot be avoided, an EES can be used to bridge profiles in a controlled and deterministic way.

Development vs Production

Mixed timing domains are common in development and validation due to legacy infrastructure and lab-grade reference clocks.

Production architectures typically converge to a single timing domain, but validation must operate correctly during the transition.

Mixed timing domains should be treated as a **transitional state** during development. Validation environments must remain flexible during architecture evolution, while production networks should converge to a single timing profile.

Testing MACsec-protected traffic

technica engineering
Member of KPIT Group

Automotive MACsec allows protection of all ethernet traffic providing frame data integrity, data origin authentication and optionally user data confidentiality through encryption.

MACsec-protected traffic:

- Cannot be injected/modified unless done through an active participant
- Can appear as uncomprehensible data to standard analysis tools.

This creates a major validation challenge.

Validation must preserve production security assumptions. Disabling MACsec and/or MKA or weakening trust boundaries with configuration changes, could hide performance or system integration issues and invalidate results.

The active participant in the hop-by-hop MACsec pattern

EES enables controlled observability:

1. **MACsec sessions terminate on EES ports.**
2. **MACsec traffic is unprotected inside the switch**
3. **A diagnostic copy is mirrored to a logger.**
4. **Frames are protected again on egress.**

Key points:

- **MKA negotiation is fully handled by the EES.**
- **External tools do not need to manage or access CAK / CKN material.**
- **No DUT modification required.**
- **No security compromise.**

Validation teams can now participate in secured communication without breaking it. They don't even have to know about MACsec/MKA to implement them in the network, EES switches can handle it within the device itself.

Building a secure IVN with MACsec/MKA

technica engineering
Member of KPIT Group

Automotive MACsec protects ethernet frames at layer 2.
EES implements MACsec in hardware, with full line-rate support.

MKA (MACsec Key Agreement Protocol)

Handled inside EES:

- CAK & CKN provided in UI.
- Taking both roles as server or client
- Web GUI provides feedback about MKA negotiation status

Challenge: Rest bus simulation with MACsec-protected ECUs – MACsec-protected traffic cannot be simulated by normal tooling.

Solution: The active participant in the hop-by-hop MACsec pattern

EES supports:

- Forward protected traffic normally.
- Mirror an unprotected copy to: SFP+ → Logger → Wireshark.
- MKA negotiation fully handled by the EES, external tools don't need to worry about CAK/CKN.
- EES protects your traffic so you can use your existing tools as in a scenario without MACsec

This enables troubleshooting without breaking security.

Traffic engineering for validation

Validation is not about forwarding all traffic indiscriminately.
It is about deliberate, selective control of specific flows during test scenarios.
EES provides TCAM-based traffic engineering to support targeted validation experiments.

Pattern 1: Rest-bus simulation

Validation benches are often built before all ECUs are available, or long before the full vehicle network exists.

To continue development and validation, parts of the network must be **functionally simulated.**

- Traffic for specific ECUs or services is selectively routed to a simulation device
- Responses are modified or generated by the simulator and re-inserted into the live network
- All other Ethernet traffic continues on its original path without interruption
- No physical rewiring or topology changes are required

Used when:

- Parts of the vehicle network are missing, or unavailable
- ECU behaviour must be exercised against realistic network responses, Bench rewiring would slow iteration or introduce instability

Pattern 2: Selective logging

Validation benches carry a mix of high-bandwidth sensor traffic and low-bandwidth control, audio, or diagnostic flows. Logging everything is often impractical and unnecessary.

With selective logging:

- Only explicitly defined streams (e.g. audio, control, diagnostics) are mirrored to the logging path
- High-bandwidth video or raw sensor traffic continues directly between ECU
- Storage usage and post-processing effort are significantly reduce
- Original timestamps and timing relationships are preserved

Used when:

- Full-bandwidth logging exceeds storage or I/O limits
- Analysis focuses on specific subsystems rather than raw sensor payloads
- Timing correlation must be maintained without logging the entire bench

Traffic engineering for validation

technica engineering
Member of KPIT Group

Pattern 3: Failure injection

With failure injection:

- Intentionally drop or modify selected frames to provoke defined fault conditions.
- Counters quantify how often events occur
- ECU and system behavior is observed under repeatable, controlled failures./scenarios
- Tests remain deterministic and repeatable across runs

Used when:

- Verifying timeout handling and error recovery mechanisms
- Testing robustness against packet loss or malformed traffic
- Validating system behavior under degraded network conditions

The current EES implementation supports up to **48 traffic rules per device** with roadmap expansion driven by customer use cases.

Pattern 4: TSN Configuration Exploration

During development, multiple TSN configurations are often evaluated before a final architecture is committed. Iteration speed matters more than static correctness at this stage.

With TSN configuration exploration:

- Traffic scheduling, prioritization, and queue behavior can be adjusted in real time.
- System behavior is observed immediately without rewiring or ECU reprogramming.
- Multiple configurations can be evaluated quickly and compared objectively.
- Engineers gain confidence in TSN choices before production freeze.

Used when:

- Evaluating alternative TSN strategies during development
- Comparing latency, jitter, and prioritization effects
- Tuning schedules and queues prior to SOP decisions

Common bench problems & non-disruptive fixes

technica engineering
Member of KPIT Group

Problem A:

MACsec-protected links block your existing validation tools.

Overview fix:

Keep MACsec enabled and mirror an unprotected diagnostic copy to a logging path.

Problem B:

Mixed timing domains (PTPv2 vs gPTP) in development.

Overview fix:

Use a timing bridge in the lab to stabilize synchronization, then plan convergence to a single profile for SOP.

Problem C:

Dual-side tapping breaks timing roles.

Overview fix:

Create isolated timing domains (partitions) so both sides can be observed without changing roles.

Problem D:

Simulation detours require physical rewiring.

Overview fix:

Apply selective routing rules to divert only target flows to a simulation host and reinject them dynamically

3 simple frameworks you can adopt immediately

technica engineering
Member of KPIT Group

Modern IVN validation does not require rewriting your bench architecture.

It requires adopting a few repeatable design frameworks that scale across timing, security, and traffic control challenges.

Framework 1:
Secure + Visible

Goal: Maintain MACsec protection while gaining payload visibility and injection on a diagnostic path.

Why this matters:
MACsec protects in-vehicle networks but may invalidate traditional validation tools.

Disabling MACsec and/or MKA invalidates security-relevant test results

How (overview):

- EES deals with MACsec and MKA.
- MACsec-protected traffic is forwarded between ECUs adding an additional active participant "hop-by-hop".
- All egress traffic is protected again before leaving the switch.
- An unprotected diagnostic copy is mirrored to a highspeed logging interface.
- DUTs continue to send and receive only protected frames.
- Validation tools operate as they did before MACsec, without managing keys or security roles.

Framework 2:
Synchronize

Goal: Share one time across cameras, radar, lidar and ECUs so that events align.

Why this matters:
Precise validation depends on understanding when events occur relative to each other.

Timing issues often emerge only when logging, simulation, or diagnostics are introduced.

How (overview):

- A single Grandmaster provides time to the bench.
- Timing is distributed through the network without altering ECU roles.
- Timing health is monitored continuously to ensure consistency across participants.

This framework applies both to native gPTP environments and mixed PTPv2 ⟷ gPTP development setups.

Framework 3:
Orchestrate

Goal: Control how selected traffic flows during validation without disturbing the rest of the network.

Why this matters:
Validation rarely requires observing or manipulating all traffic.

Uncontrolled logging, simulation rewiring, or scripting introduces side effects that hide root causes.

How (overview):

- Match conditions are defined for selected traffic flows.
- Matching frames are mirrored, forced to a specific port, or dropped.
- Counters track how often rules are applied.
- All non-matching traffic flows remain unchanged.

This enables selective logging, failure injection, rest-of-bus simulation, and TSN exploration without topology changes.

Best-practice bench architectures

(Combining timing, security, filtering, hybrid ports, high-speed logging.)

technica engineering
Member of KPIT Group

Modern validation benches must support secure communication, precise timing, and controlled observability at the same time. The following reference architectures illustrate how these requirements can be combined without introducing bench instability.

1. Hybrid bench

Architecture includes:
T1, RJ45, SFP+, EES

Use EES hybrid when:

- Automotive T1 ECUs (cameras, sensors, domain controllers) must be connected directly.
- RJ45-based hosts, analysis PCs, or simulation systems are part of the bench.
- Highspeed logging paths are required without disturbing production links.

Key characteristics:

- T1 ports connect directly to ECUs and sensors.
- RJ45 ports connect tools, hosts, and simulation devices.
- SFP+ ports provide dedicated, high-bandwidth diagnostic or logging paths.
- Rest-of-bus simulation traffic can be injected or reintegrated without rewiring.

Diagrammatic representation on the next page

Diagrammatic representation of
Hybrid bench architecture

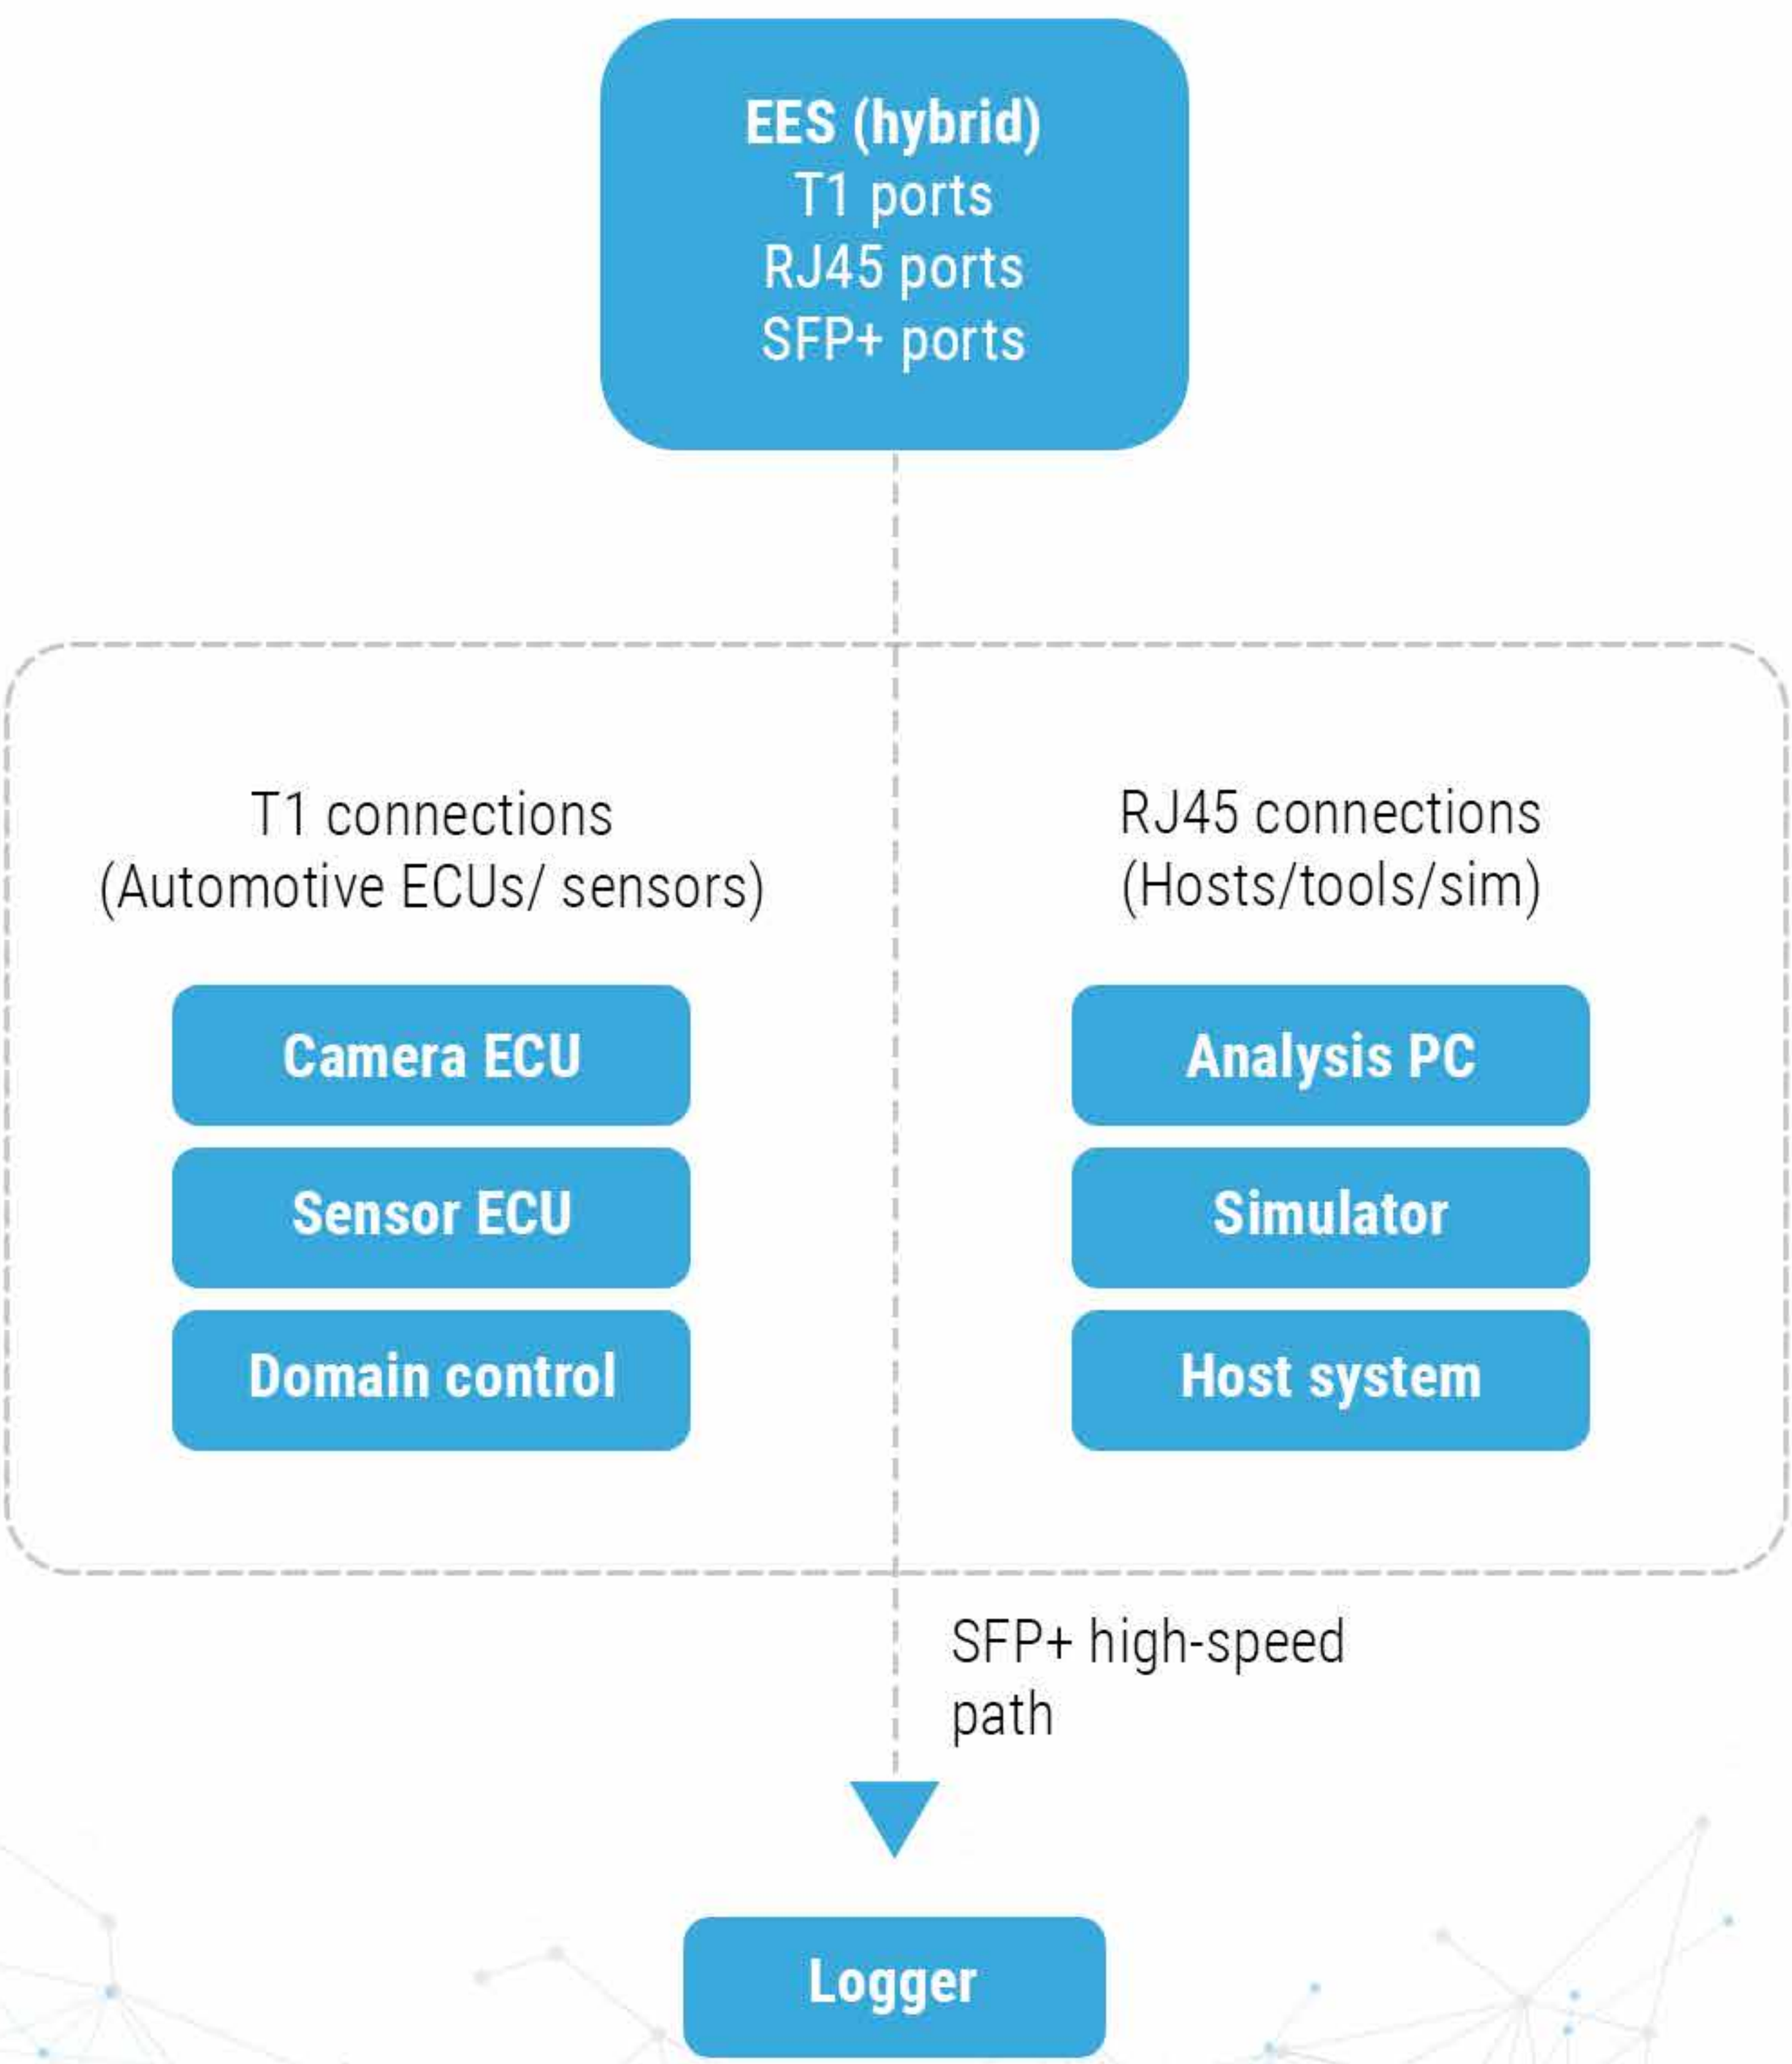

2. Secure high-speed logging path

Architecture pattern:

Production Link (MACsec ON) → **EES** → Diagnostic Mirror → SFP+ → Logger

Why this matters:

Highbandwidth logging is often required for deep analysis, but production links must remain MACsec-protected and unchanged.

Key Properties:

- MACsec remains enabled on all production facing ports.
- MACsec-protected traffic is forwarded normally between ECUs adding an additional active participant “hop-by-hop”.
- All egress traffic is protected again.
- An unprotected diagnostic copy is mirrored internally to a highspeed logging interface.
- Logging and analysis do not influence timing or security behavior.

Diagrammatic representation on the next page

Diagrammatic representation of
Secure high-speed logging path

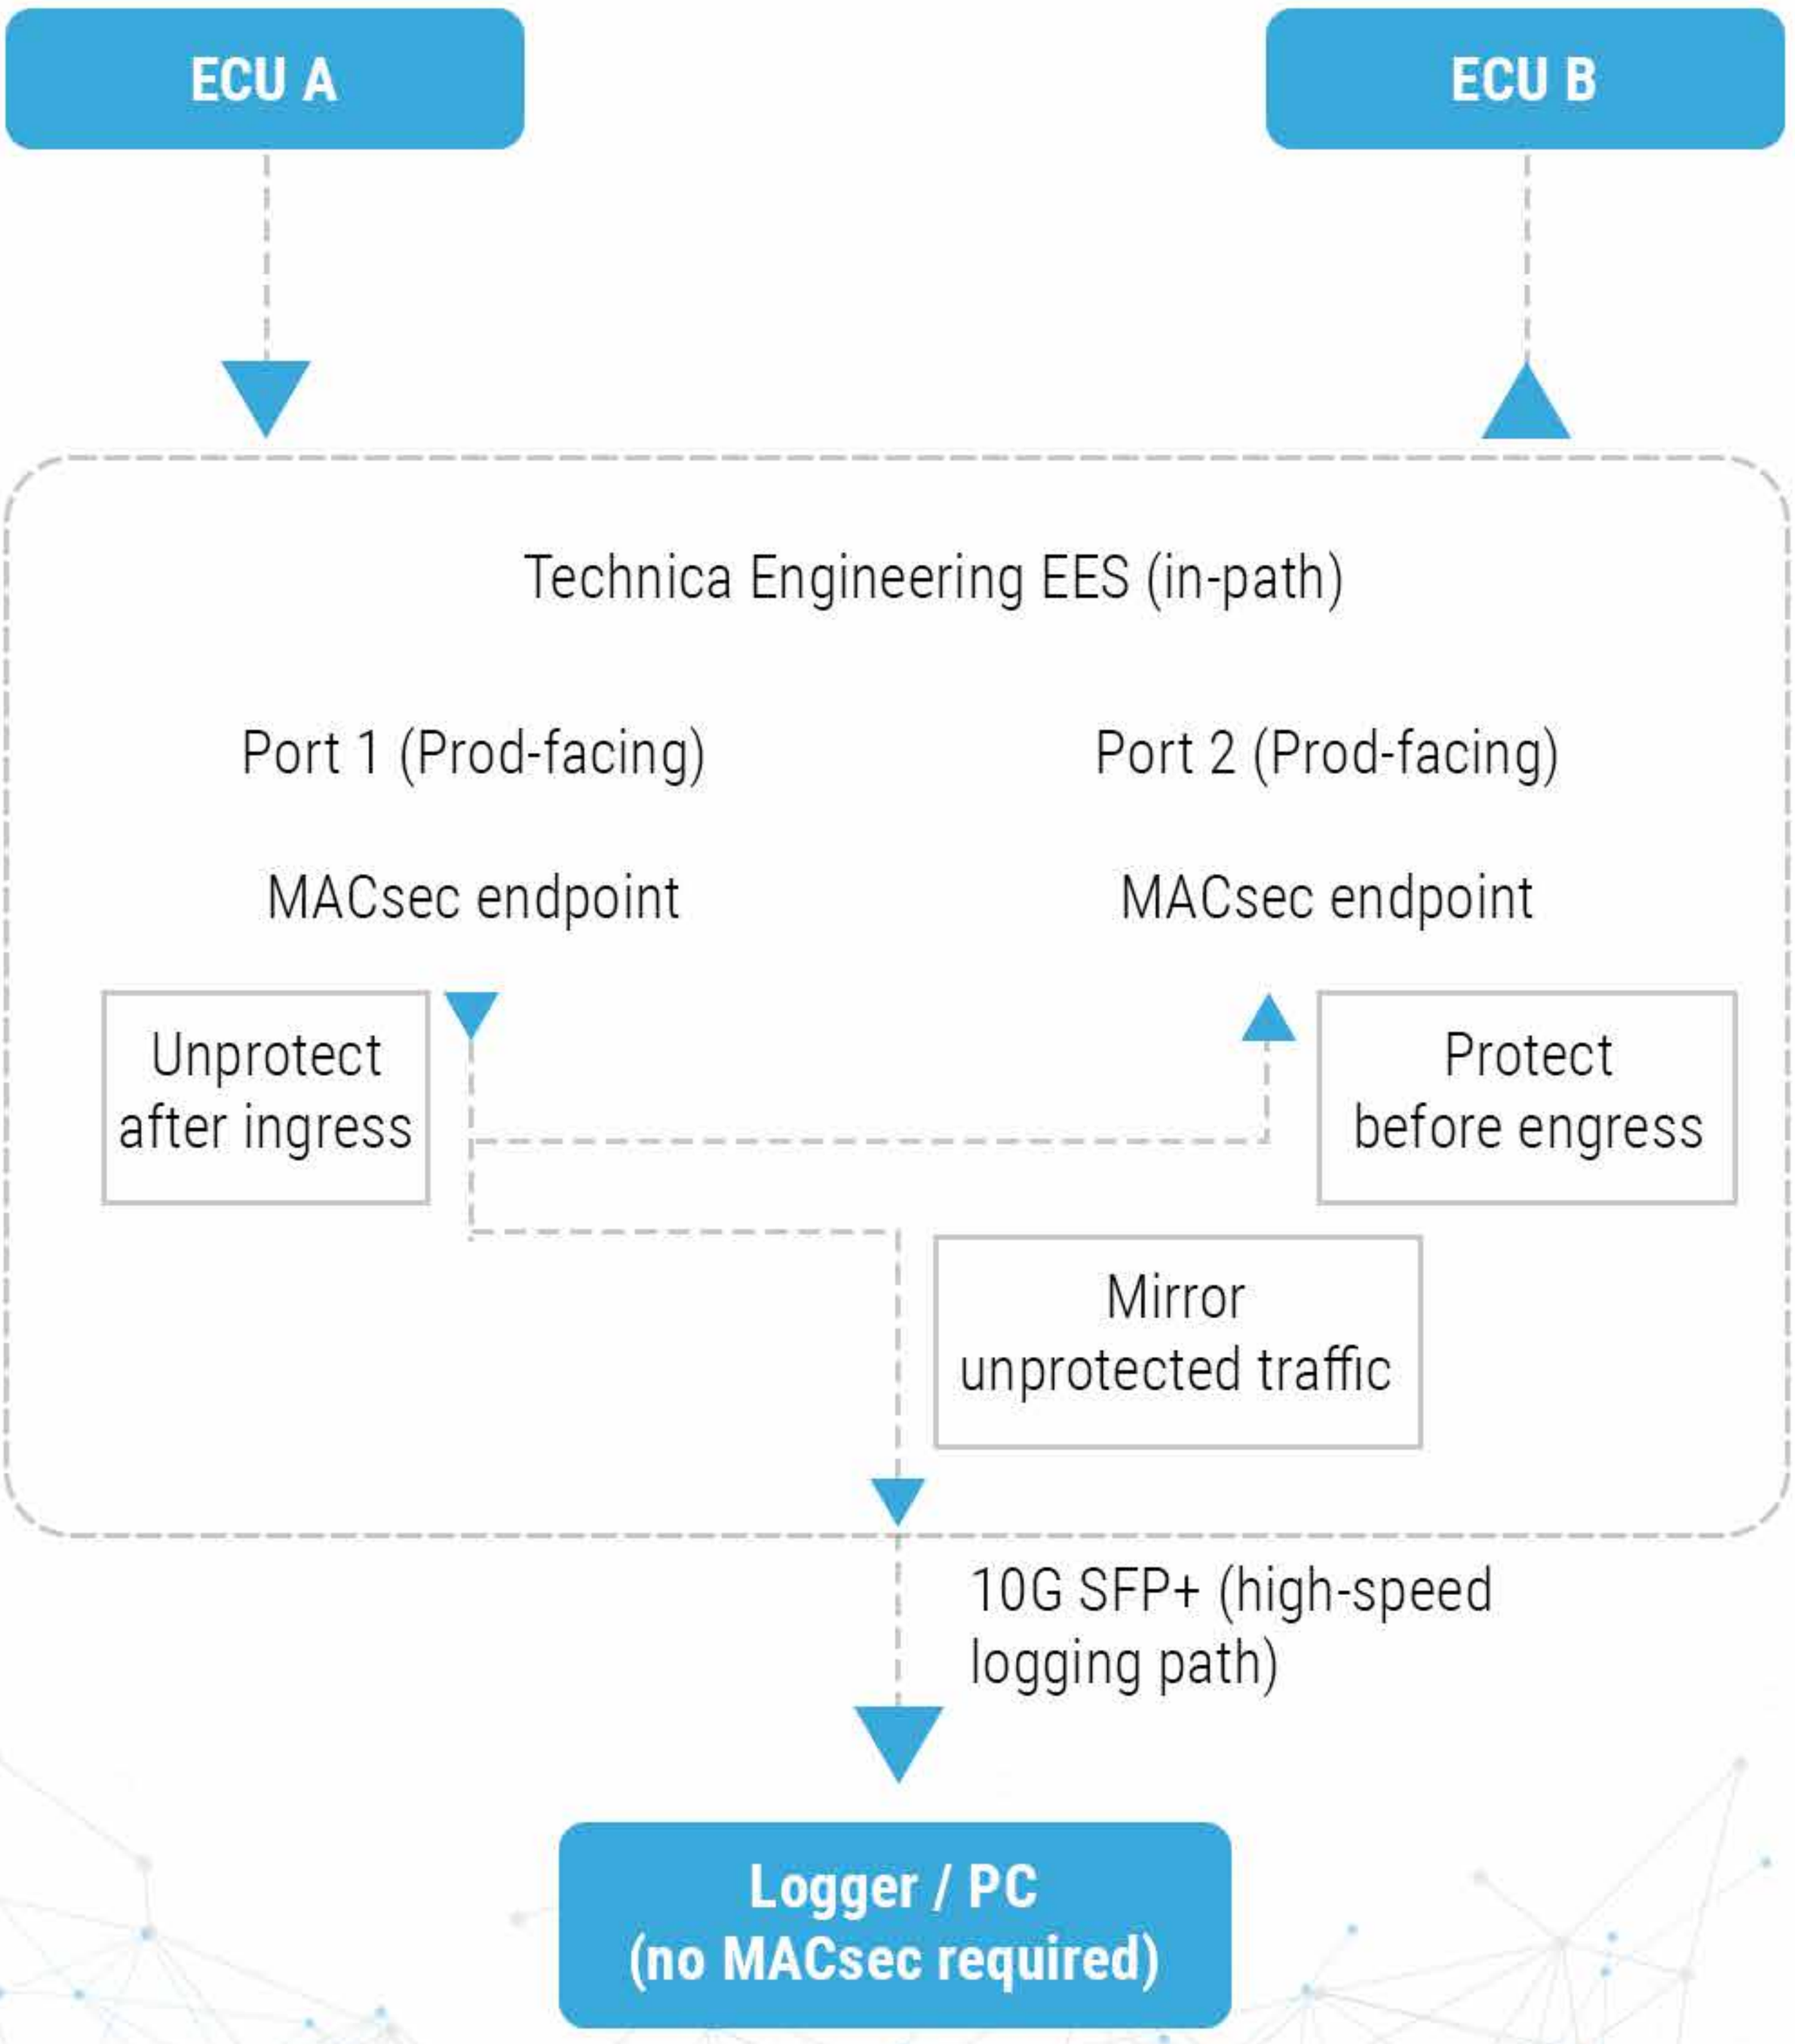

3. Multi-sensor timing distribution tree

Architecture pattern:

External Grandmaster → **EES** → Multiple Sensors / ECUs

Why this matters:

Validation benches often aggregate multiple sensors whose data must be correlated across ECUs.

Key characteristics:

- Time is distributed deterministically from a single source.
- Sensors and ECUs participate in the same timing domain.
- Timing behavior remains stable when logging or diagnostics are enabled.

Diagrammatic representation on the next page

Diagrammatic representation of
Multi-sensor timing distribution tree

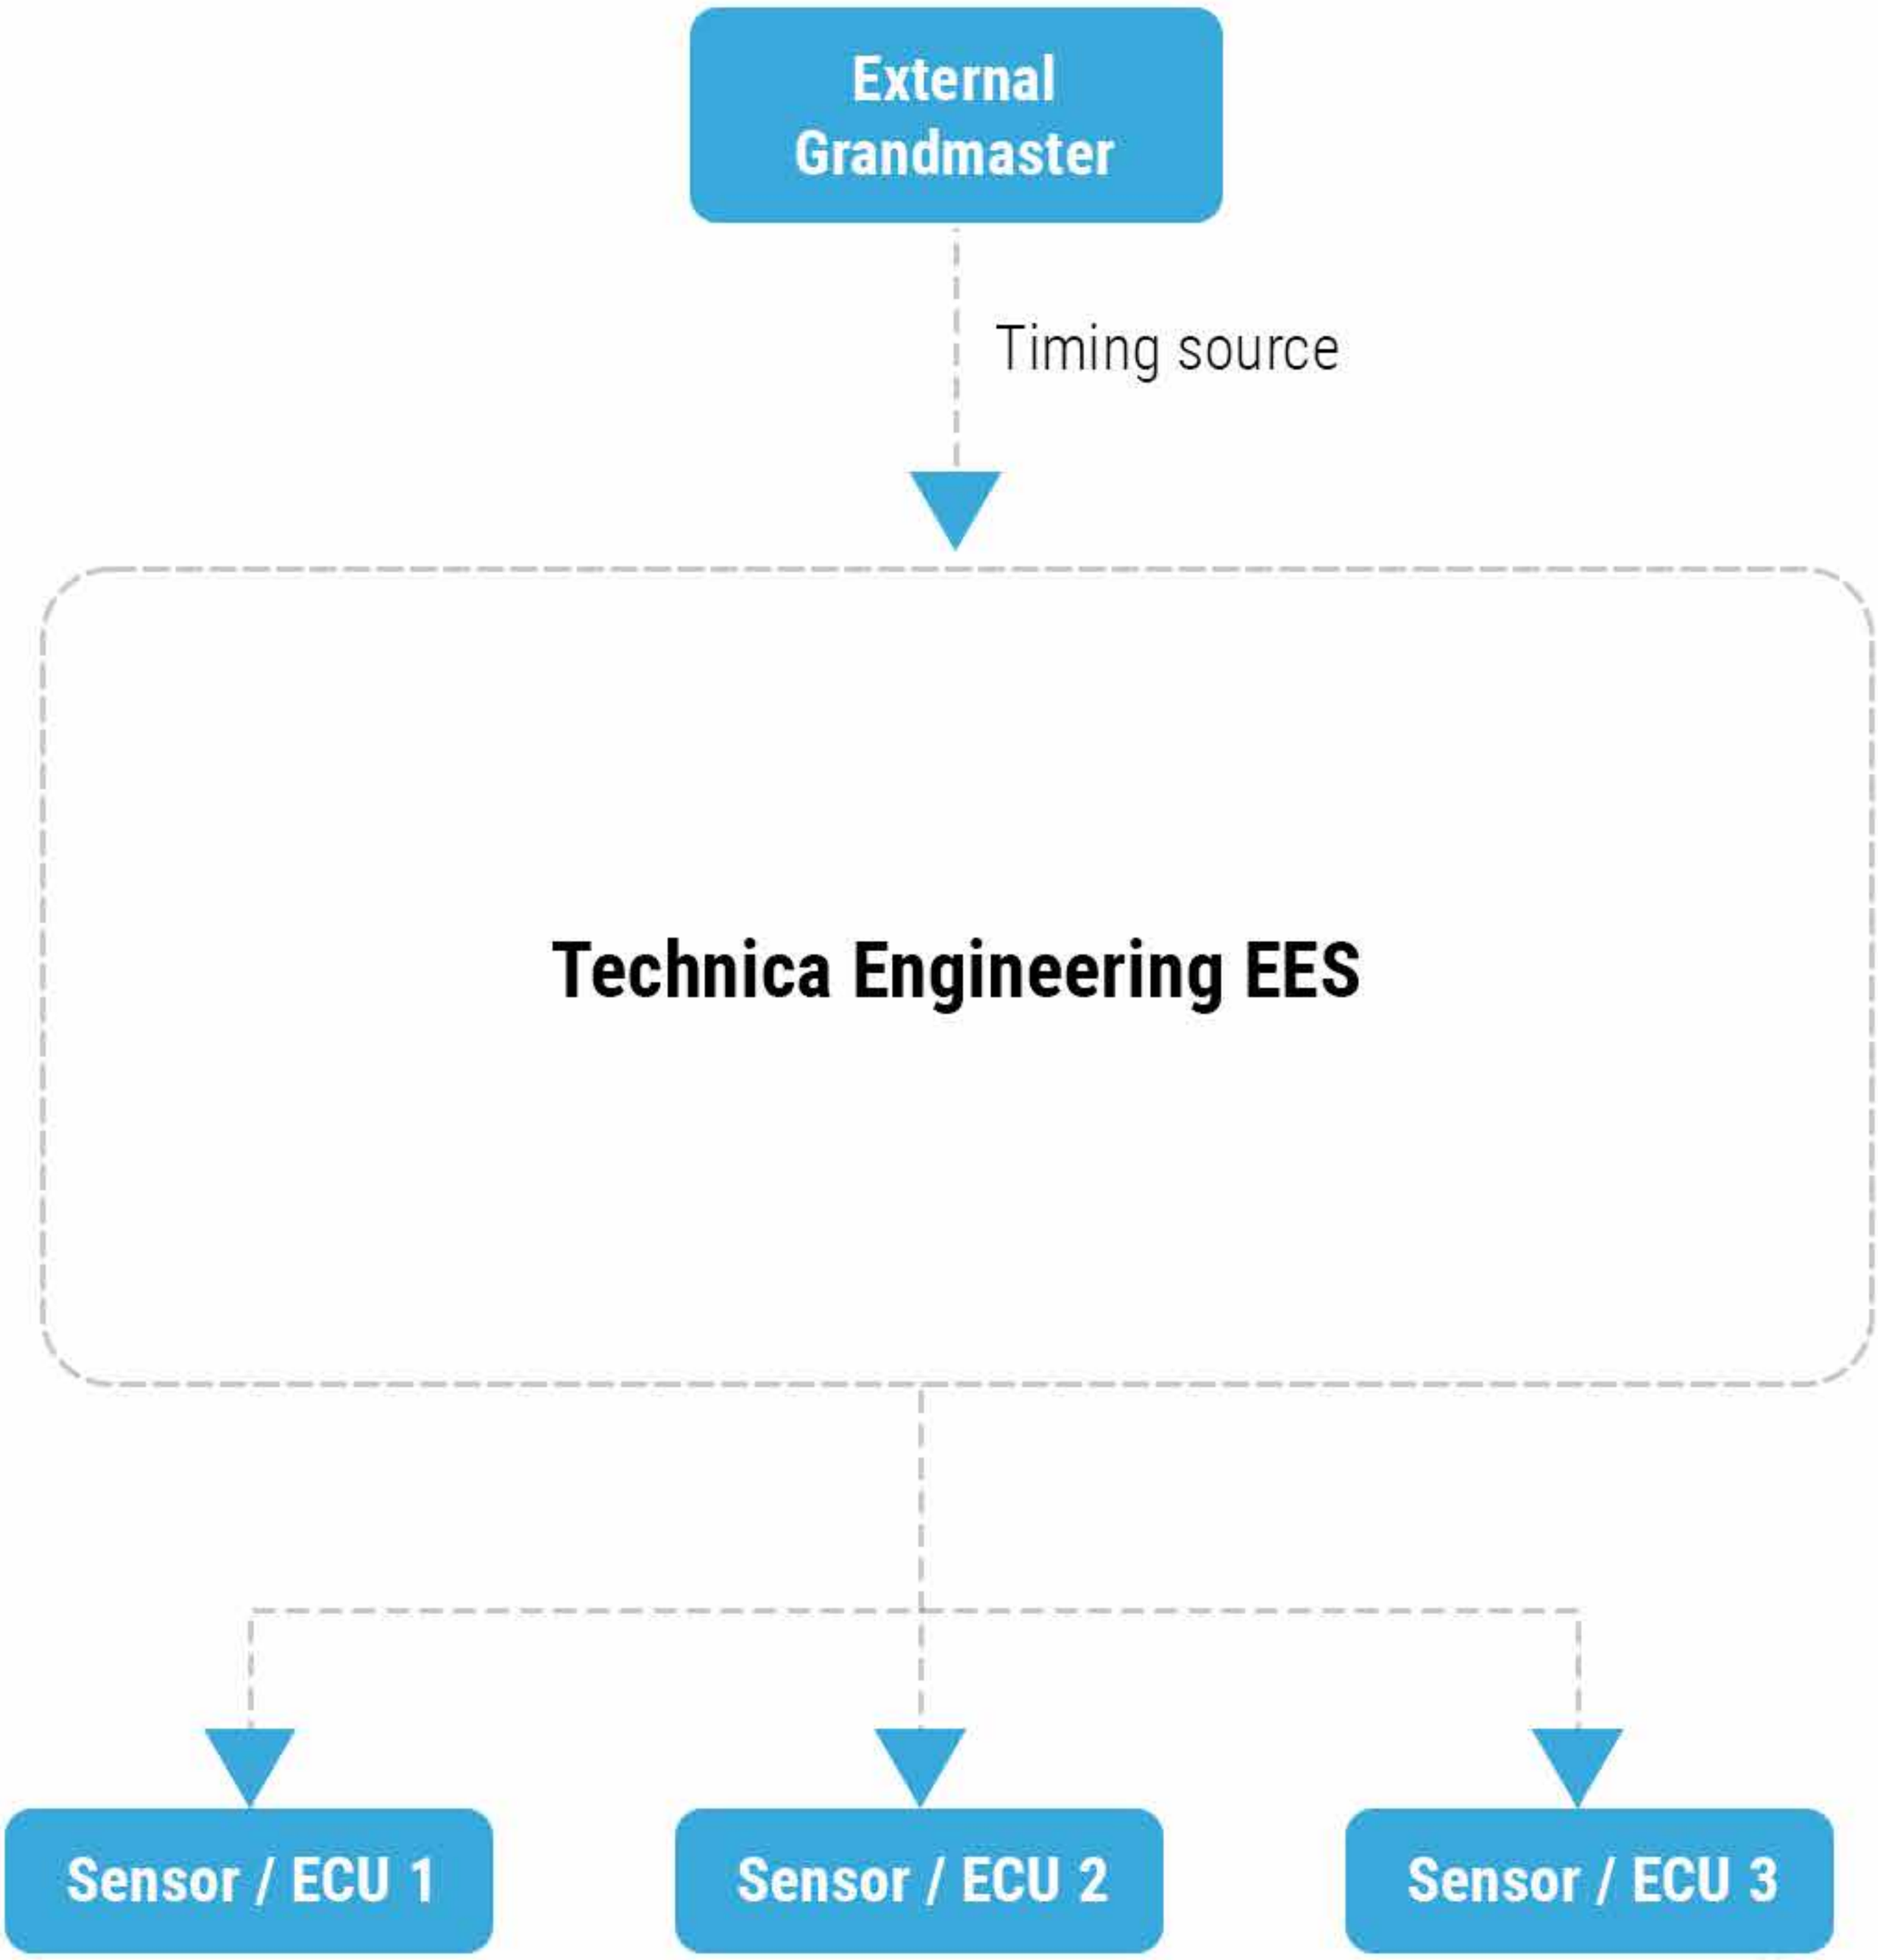

4. Mixed timing domain development bench

Architecture pattern:

PTPv2 Grandmaster → **EES (Timing Bridge)** → gPTP Sensor ECUs

Why this matters:

Development environments often combine lab-grade reference clocks with automotive ECUs that expect gPTP.

Key characteristics:

- Existing PTPv2 infrastructure can be reused during development.
- gPTP-based ECUs operate normally without modification.
- Timing domains remain isolated and controlled.

Guidance:

- Recommended for development and validation environments.
- Production (SOP) architectures should converge to a single timing profile.

Diagrammatic representation on the next page

Diagrammatic representation of
Mixed timing domain development bench

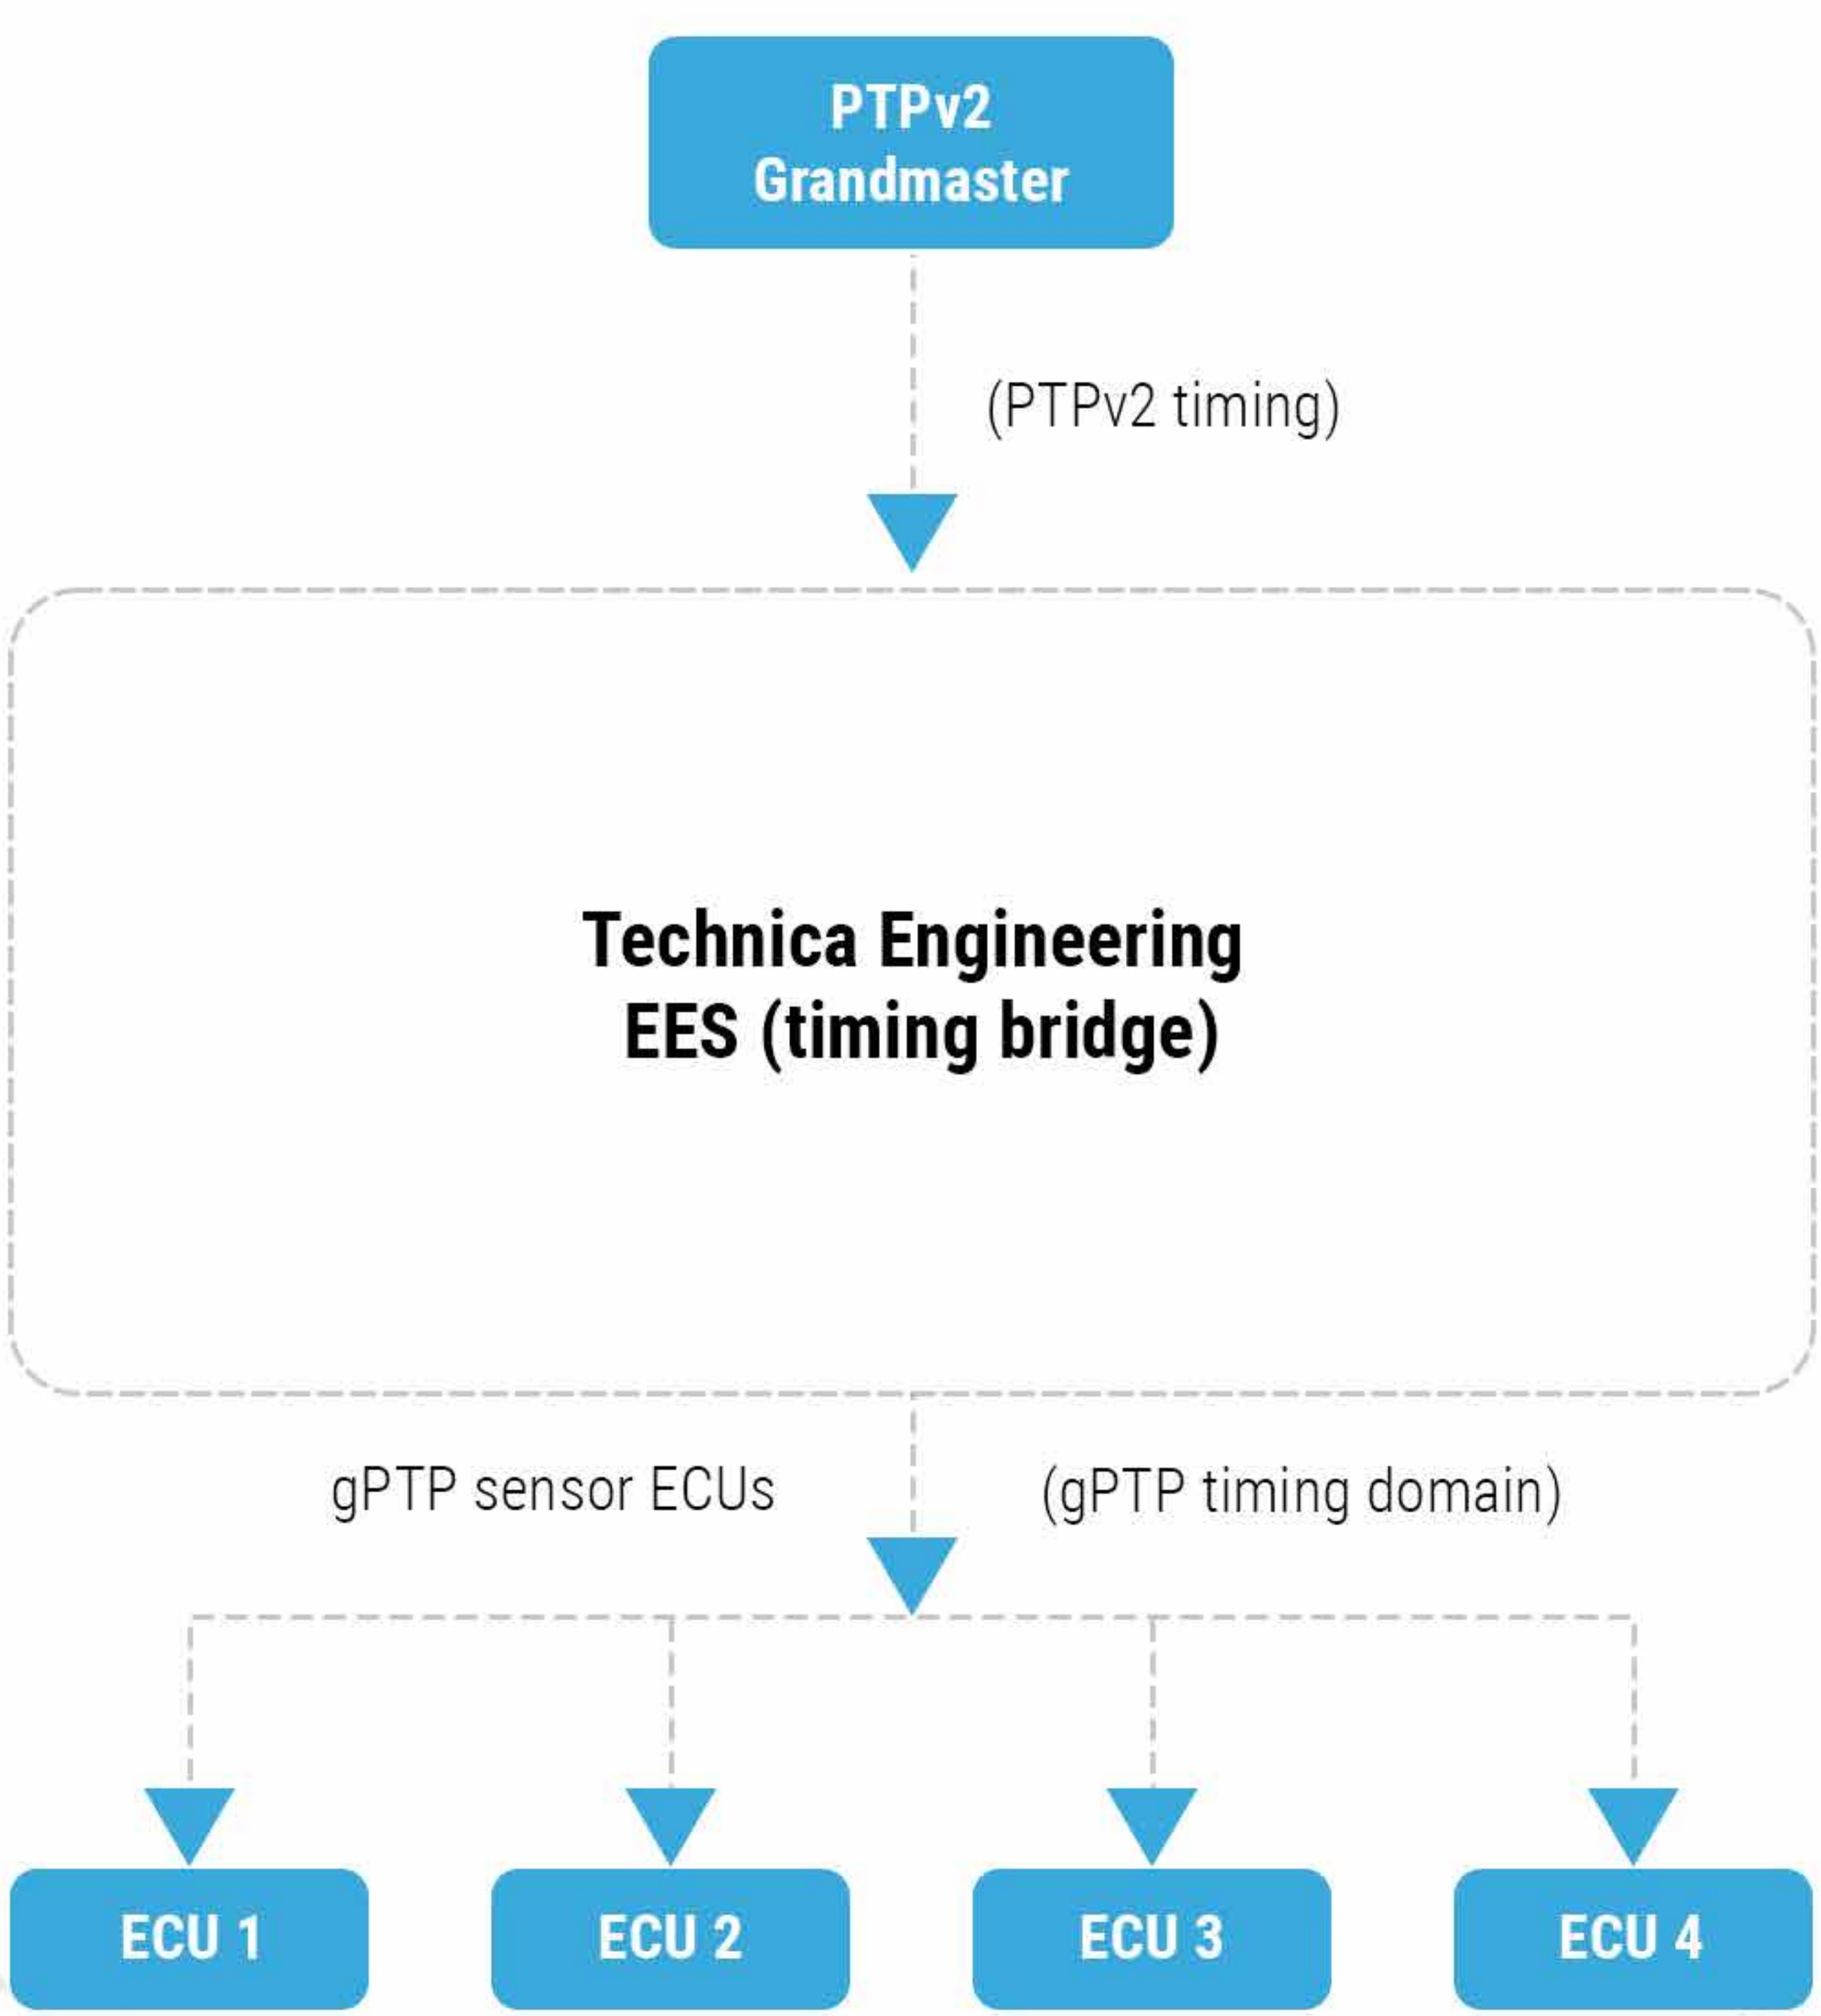

Validation challenges in modern IVN test benches

(What has become normal — and why it's difficult)

technica engineering
Member of KPIT Group

Automotive Ethernet validation environments have evolved rapidly.
What were once edge cases are now routine challenges in day-to-day validation work.

The sections below outline where modern benches struggle, not because of poor design, but because security, timing, bandwidth, and toolchains are now tightly coupled.

Timing & synchronization challenges

Modern validation benches rely on precise time alignment across distributed ECUs and sensors. Issues rarely appear in isolation; they emerge when tools are added or topologies evolve.

Common challenges include:

- ☐ Adding a tap or logger may impact in gPTP behavior
- ☐ Timing roles shift unexpectedly when new devices are introduced
- ☐ Mixed PTPv2 and gPTP domains are present during development
- ☐ Synchronization appears stable until logging or simulation is enabled
- ☐ Correlating events across sensors and ECUs becomes unreliable

These issues are often subtle: time synchronization may not fail outright but degrade silently, producing repeatable yet invalid results.

Security & MACsec-protected network challenges

Security is no longer optional – MACsec is increasingly enabled early in development. While necessary, it impacts the traditional validation workflows.

Common challenges include:

- ☐ MACsec-protected links block payload inspection by standard tools
- ☐ Requirements for Injection of traffic, for example Rest-Bus Simulation, becomes a challenge
- ☐ Security teams restrict access to keys and trust boundaries
- ☐ Validation equipment operates outside the MACsec-protected domain
- ☐ Debugging workflows require disabling or bypassing MACsec protection

Note: Limited diagnostic traffic may be permitted via MACsec bypass lists, but this does not address use cases requiring full participation in a MACsec-protected link or payload inspection.

Validation challenges in modern IVN test benches

(What has become normal — and why it's difficult)

technica engineering
Member of KPIT Group

Traffic control & bench complexity challenges

As benches grow, workloads diversify and bandwidth demands increase. Without deliberate traffic control, validation setups become fragile and difficult to reason about.

Common challenges include:

- ☐ Logging all traffic exceeds storage or I/O capacity
- ☐ Simulation requires physical rewiring of the bench
- ☐ Failure injection relies on scripts or ad-hoc workarounds
- ☐ Small changes introduce unintended side effects
- ☐ The bench has grown organically rather than architecturally

In many cases, the bench itself becomes harder to validate than the system under test.

To see how these patterns are implemented on a real test bench, including live configuration, metrics, and troubleshooting, join our upcoming webinar.

See Enhanced Ethernet Switches in action - Live

technica engineering
Member of KPIT Group

This playbook introduces the design patterns behind secure, synchronized, and observable IVN validation. To see how these patterns work on a real test bench, we invite you to join our live webinar. This is where the concepts in this playbook will be executed, observed, and explained in real time.

FREE WEBINAR

Upcoming Webinar:

Secure. Synchronized. Easy.
Modern IVN Validation with Technica Enhanced Ethernet Switches

Date: March 31, 2026
Time: 01:00 PM Stuttgart | 04:30 PM Mumbai | 07:00 AM Detroit
Format: Live technical walkthrough + demos

In the webinar, you'll see:

- A practical "one-use-case" IVN validation workflow: connect ECUs, time sync, log safely, and introduce simulation paths. All within one coherent bench storyline.
- MACsec testing without going blind: how to add an additional active participant "hop-by-hop", and configure authentication/encryption to enable observation/ diagnostics/rest-bus simulation while maintaining security strategy.
- Traffic control that supports validation experiments: advanced filters and actions (mirror/drop/count/force-route) for targeted inspection or redirection.
- Simulation interaction without bench rework: redirect/select frames to a simulation PC and inject secure traffic back through the EES path.
- What "simplification" looks like in reality: examples of reduced devices and improved scalability/control from real deployments.
- You'll also experience this with a live demo, along with some practical tips.

Bring your bench topology, timing profile, or MACsec questions.
The session is built around real validation problems.

REGISTER NOW

AN INITIATIVE BY