

# How to Meet Regulatory Software-Update Rules for Autonomous Vehicles

If you ship software to autonomous vehicles or field robots, regulators increasingly demand that every update be authorized, traceable, and reversible before it changes behavior on a public road. This guide describes an architecture for delivering updates as governed, signed artifacts that a receiving unit validates against policy and continuity before the update takes effect. It teaches an approach disclosed in U.S. Provisional Application No. 64/049,409, not a shipping library. The home inventive step is the Spatial Adaptation inventive step.

---

## What You Are Building

You are building the control layer that sits between "an update is available" and "the update is running on the vehicle." Regulatory regimes for automotive software updates increasingly require that a manufacturer prove, per update and per unit, that a change was authorized by the right party, that its effect on vehicle behavior was assessed before it went live, that the change can be undone, and that the whole sequence left an audit trail. Whether you are subject to a formal software-update management framework, a type-approval regime, or an internal safety case, the engineering problem is the same: a raw firmware or model update landing on a moving vehicle and silently changing how it behaves is exactly what regulators do not want.

This guide is for engineers who build the update path for autonomous units: self-driving vehicles, warehouse robots, port vessels, airfield ground vehicles, or wearable-equipped personnel. It describes an architecture in which an update is not a file you push, but a governed, signed artifact the receiving unit must evaluate and admit before it can take effect. The approach is disclosed in U.S. Provisional Application No. 64/049,409; you implement it yourself.

## **Why the Obvious Approaches Fall Short**

The conventional over-the-air update path is a centralized distribution server, a signed package, and a client that checks the signature and installs. This is a real and necessary baseline, and code-signing with public-key infrastructure genuinely establishes that a package came from a holder of a given key. The gap is not in the cryptography; it is in what the signature means and what happens after verification.

A publisher signature answers "did this come from us." It does not answer "is this party the correct authority to change behavior on this class of unit in this region," because a flat signing key carries no hierarchy: an operational key and a safety-regulator key verify identically. It does not answer "what will this change do before it runs," because in the common pattern the package is verified and then applied, with any behavioral assessment having happened offline in the lab against a fleet the update may no longer match. And a plain install is often a one-way door: rolling back means shipping another update, and the record of who changed what, when, and under which policy lives in server logs rather than on the unit that has to answer for its own behavior.

The centralized server itself is also a structural dependency. Each unit must reach it and hold a credential for it, which is fragile precisely in the sparse or intermittent connectivity where field vehicles operate. None of this is a flaw in signing; it is a mismatch between "verify and install" and what a regulator is actually asking you to demonstrate.

## The Architecture

The disclosed approach reframes an update as a **governed observation**: a signed, self-describing artifact that travels through the same governed mesh the units already use to exchange information, and that each receiving unit evaluates through a single admissibility gate before it takes effect. The specification covers three artifact kinds under one propagation pattern: governance-policy updates, firmware updates, and skill-adapter (model-adaptation) updates.

**Authority-credentialed distribution.** An update is published by a deploying authority as a governed observation carrying the update's version, its content, the deploying authority's signature, and an applicability-scope specification identifying which devices are eligible to receive it. Critically, the credential is not flat. The disclosure describes an authority taxonomy: a hierarchy of credentialed authority levels defined for an operational domain. In the roadway example the levels include a regulatory-infrastructure authority (national, state, or municipal transportation authorities), an emergency-preemptive authority, an operational authority, and an un-credentialed level. A receiving unit therefore evaluates not just whether a signature is valid but what authority level stands behind it, and higher-authority policy supersedes a conflicting policy set by a lower authority.

**A single admissibility gate before effect.** When a unit receives an update, it does not install it. It routes the update through a composite admissibility evaluator, the same evaluator the architecture uses for every other admission decision. For a policy update the evaluator verifies the deploying authority's credential, the evidential weight that the receiving unit's authority taxonomy assigns to that authority, and the consistency of the update with the unit's prior policy state. The evaluator does not return a binary yes or no; it returns one of several outcomes: admit, gate (admit under additional constraints), defer (hold pending corroboration, with an expiration), solicit (query for

corroborating evidence), reject (with a reason classification), or escalate. Every determination is itself emitted as a governed record noting the inputs, the weights applied, the outcome, and the governance-policy version in force.

**Sandbox evaluation before activation.** Firmware and skill-adaptor updates add a step that closes the "what will it do" gap. Before a firmware update is applied, it is evaluated in a sandboxed execution environment on the receiving unit to verify it does not violate governance-policy-defined safety, integrity, or capability-envelope properties; an update that fails sandbox evaluation is not applied and the failure is recorded. For model-adaptation artifacts the disclosure is more detailed: a sandboxed copy of the unit's cognitive substrate is instantiated separately from the operational one, representative evaluation contexts are generated (from the unit's current task context, from standard evaluation suites, and from its own observation store), the candidate artifact is applied to the sandbox, and a behavior evaluator, performance-preview generator, governance-summary generator, and risk-projection evaluator feed an activation-gate controller that produces an admit-or-reject outcome. The result is written as a certification record and can escalate to a deeper, stricter second-tier evaluation when the first tier is uncertain.

**Self-describing artifacts with provenance.** A model-adaptation artifact is specified as a structured object carrying, among other fields, an artifact identifier, the adaptation-technique identifier, the content, a capability-scope specification, a compatibility specification, a licensing specification, a dependency specification, a certification record, a provenance-lineage record (training inputs, methodology, contributing authorities, antecedent artifacts), an authority credential, a version identifier and version-lineage chain, and a cryptographic integrity attestation over those fields. The audit trail regulators ask for is thus carried by the artifact itself rather than reconstructed from server logs.

**Atomic application, lineage, and rollback.** On admission, a policy update is applied as an atomic transition from the prior state to the new state, and the transition is written to the unit's lineage record. A rollback mechanism reverts to the prior state on detected admission failure or on receipt of a later revocation. For adaptation artifacts, rollback returns the unit to a prior version upon detection of post-activation policy violations attributable to the current version, recording the rollback, the triggering violation, and the restored artifact.

**Lifecycle and regulator-driven deprecation.** Artifacts move through publication, certification, active consumption, version-revision, deprecation, and retirement, each transition lineage-recorded. Deprecation is governance-policy-configurable and explicitly includes security-driven deprecation (consumption rejected immediately on a security-relevant defect) and regulatory-driven deprecation (consumption constrained in response to regulatory changes), alongside soft and hard deprecation with transition windows.

**Working without a central server.** Because updates propagate through the governed mesh with multi-hop relay and mobile store-and-forward carriage, an update emitted at an ingress point can be carried by units traversing a sparse-connectivity region and rebroadcast to units there, without each unit holding an out-of-band credential for a distribution server. Scope specifications (geographic, temporal, device-class) bound where an update applies.

## **How to Approach the Build**

1. **Define your authority taxonomy first.** Before any transport work, decide the ordered authority levels for your domain and which real-world party holds each (regulator, emergency services, operator, unknown). This taxonomy is what turns a signature into a decision. Bind each level to a credential and encode supersession: higher authority overrides lower.

2. **Make the update a self-describing artifact, not a file.** Give every update a structured envelope: version and version-lineage link, content, applicability scope, issuing authority credential, and an integrity attestation over the whole. For model updates, add capability scope, compatibility, dependencies, and provenance. Treat these fields as the audit record.
3. **Build one admissibility gate and route everything through it.** Implement a single evaluator that takes an update and returns admit / gate / defer / solicit / reject / escalate, and have it emit a signed determination recording inputs, weights, outcome, and the policy version in force. Resist per-update-type bespoke logic; uniformity is what makes the audit story provable.
4. **Insert sandbox evaluation before activation.** Stand up an isolated copy of the runtime (or cognitive substrate) separate from the operational one. Generate representative contexts, apply the candidate, and evaluate compatibility, behavior, and risk into an activation gate. Nothing reaches the operational substrate until the gate admits it. An illustrative interface, faithful to the disclosure:

```
# Illustrative sketch, not a shipping API
outcome = admissibility.evaluate(update, policy_state, authority_taxonomy)
if outcome == ADMIT:
    cert = sandbox.evaluate(update, contexts) # isolated substrate
    if cert.activation_gate == ADMIT:
        prior = state.snapshot()
        state.apply_atomic(update); lineage.record(update, cert)
# rollback path: state.restore(prior) on later violation/revocation
```

5. **Apply atomically and keep the prior state.** Snapshot before applying so rollback is a restore, not another update round-trip. Wire rollback to two triggers: admission failure and a later revocation or violation.

6. **Record every transition in on-unit lineage.** The unit, not a server, should be able to reconstruct what changed, under whose authority, and under which policy version. This is the artifact regulators can inspect.
7. **Model the lifecycle, including regulatory deprecation.** Support version-revision with lineage links and deprecation modes, so a regulator- or security-driven change can constrain or reject a now-noncompliant version fleet-wide within a defined transition window.
8. **Consider pre-certification for scale.** Let a credentialed certification authority run the sandbox evaluation once and publish a certification observation, so units can admit a widely-vetted update under explicit trust parameters instead of every unit re-evaluating from scratch.

## **What This Does Not Give You**

This is an architecture, not a drop-in library. There is no package to install and nothing here "just works" out of the box; you build the evaluator, the sandbox, the artifact envelope, the mesh transport, and the credentialing yourself, and you own the hard parts the disclosure treats as design surface, notably what makes a sandbox evaluation representative of on-road conditions and how you set evidential weights and policy thresholds.

The approach is disclosed in a patent filing. It is not benchmarked, not a certified compliance product, and not proof of conformance to any specific regulation; nothing here is legal advice, and mapping the architecture onto the exact clauses of a given jurisdiction's software-update rules is work you must do with your own regulatory and safety teams. The disclosure describes mechanisms (admissibility outcomes, sandbox gating, rollback, lineage) but does not supply performance numbers or safety guarantees, and you should not claim any. Finally, the design assumes units that

participate in a governed mesh and carry credentials and a cognitive-substrate runtime; a fleet without those primitives needs that foundation built first, and domains outside governed autonomous units may not fit the model at all.

## Disclosure Scope

The architecture described in this guide is disclosed in U.S. Provisional Application No. 64/049,409, which sets out the governed-observation update model, the authority-credentialed distribution and admissibility evaluation, the sandbox pre-activation evaluation, and the versioning, deprecation, and rollback mechanisms summarized above, under the Spatial Adaptation inventive step. This guide is educational: it explains an approach a developer can build, and it is neither a warranty, a certification of regulatory compliance, nor an offer of software. Any third-party technologies or standards referenced (for example, public-key infrastructure) are described for context only.

---

## **Spatial Adaptation Artifacts** (</spatial-adaptation>) [All 40 steps → \(/inventive-steps\)](#)

Runtime signed-skill loading. Sandbox-certified. Admissibility gate as skill router.

Provisional application

### **PRIMARY TECHNICAL DISCLOSURE**

- [Spatial Adaptation Artifacts: Runtime Skill Loading With Admissibility Gating \(/articles/spatial-adaptation-artifacts-runtime-skill-loading-with-admissibility-gating\)](/articles/spatial-adaptation-artifacts-runtime-skill-loading-with-admissibility-gating)

### **SECONDARY TECHNICAL**

- [Runtime-Signed Adaptation Artifacts \(/articles/spatial-adaptation/runtime-signed-artifacts\)](/articles/spatial-adaptation/runtime-signed-artifacts)
- [Sandbox Pre-Activation Certification \(/articles/spatial-adaptation/sandbox-pre-activation-certification\)](/articles/spatial-adaptation/sandbox-pre-activation-certification)

- [Admissibility as Skill Router \(/articles/spatial-adaptation/admissibility-as-skill-router\)](/articles/spatial-adaptation/admissibility-as-skill-router).
- [Always-Active Personal Layer \(/articles/spatial-adaptation/always-active-personal-layer\)](/articles/spatial-adaptation/always-active-personal-layer).
- [Cascade Deactivation Dependencies \(/articles/spatial-adaptation/cascade-deactivation-dependencies\)](/articles/spatial-adaptation/cascade-deactivation-dependencies).
- [Cross-Model Adaptation Portability \(/articles/spatial-adaptation/cross-model-portability\)](/articles/spatial-adaptation/cross-model-portability).
- [Federated Skill Training \(/articles/spatial-adaptation/federated-skill-training\)](/articles/spatial-adaptation/federated-skill-training).
- [Composite Licensing Intersection \(/articles/spatial-adaptation/composite-licensing-intersection\)](/articles/spatial-adaptation/composite-licensing-intersection).
- [Decentralized Mesh Adaptation Distribution \(/articles/spatial-adaptation/decentralized-mesh-distribution\)](/articles/spatial-adaptation/decentralized-mesh-distribution).

## **APPLICATIONS · GENERAL**

- [Governed In-Field AI Model Adaptation for Defense Operations \(/articles/spatial-adaptation/defense-field-adaptation\)](/articles/spatial-adaptation/defense-field-adaptation).
- [Safe Runtime Updates for AI-Driven Industrial Robots: Governed Adaptation Under ISO 10218 and the EU Machinery Regulation \(/articles/spatial-adaptation/industrial-robotics-adaptive-update\)](/articles/spatial-adaptation/industrial-robotics-adaptive-update).
- [Governing Adaptive Medical Device and SaMD Updates Under FDA PCCP and EU MDR \(/articles/spatial-adaptation/medical-device-adaptive-update\)](/articles/spatial-adaptation/medical-device-adaptive-update).
- [Safe Rapid Security Updates for Safety-Critical Systems \(/articles/spatial-adaptation/cybersecurity-rapid-update\)](/articles/spatial-adaptation/cybersecurity-rapid-update).
- [Regulatory-Aware LLM Adaptation: Verifiable Governance for EU AI Act and FDA Compliance \(/articles/spatial-adaptation/regulatory-aware-llm-adaptation\)](/articles/spatial-adaptation/regulatory-aware-llm-adaptation).
- [IEC 62304 Compliance for Continuously Adapting Medical Device Software \(/articles/spatial-adaptation/iec-62304-medical-software\)](/articles/spatial-adaptation/iec-62304-medical-software).
- [Enforcing NIST AI RMF Compliance at Runtime for AI Model Fleets \(/articles/spatial-adaptation/nist-ai-rmf\)](/articles/spatial-adaptation/nist-ai-rmf).
- [UNECE R156 SUMS Compliance for In-Vehicle Software Updates \(/articles/spatial-adaptation/unec-r156-software-update\)](/articles/spatial-adaptation/unec-r156-software-update).

## **APPLICATIONS · SPECIFIC**

- [Governed Adaptation Beyond Anthropic Skills: Admissibility-Gated, Reversible Capability Loading \(/articles/spatial-adaptation/anthropic-skills\)](/articles/spatial-adaptation/anthropic-skills).
- [OpenAI Fine-Tuning vs Governed Model Adaptation \(/articles/spatial-adaptation/openai-fine-tuning\)](/articles/spatial-adaptation/openai-fine-tuning).
- [Tesla FSD Updates vs Governed Adaptation Artifacts \(/articles/spatial-adaptation/tesla-fsd-updates\)](/articles/spatial-adaptation/tesla-fsd-updates).

- [AWS Bedrock vs Governed Adaptation: The Certifiable-Artifact Layer \(/articles/spatial-adaptation/aws-bedrock\)](/articles/spatial-adaptation/aws-bedrock)
- [Databricks Mosaic AI vs Governed Adaptation Artifacts \(/articles/spatial-adaptation/databricks-ai\)](/articles/spatial-adaptation/databricks-ai)
- [Google Vertex AI vs. Governed Adaptation: Who Owns Model-Adaptation Governance? \(/articles/spatial-adaptation/google-vertex-ai\)](/articles/spatial-adaptation/google-vertex-ai)
- [Hugging Face Hub Alternative for Governed Model Adaptation \(/articles/spatial-adaptation/huggingface-hub\)](/articles/spatial-adaptation/huggingface-hub)
- [Governed Agent Adaptation vs Anthropic Claude MCP \(/articles/spatial-adaptation/anthropic-claude-mcp\)](/articles/spatial-adaptation/anthropic-claude-mcp)

---

[Spatial Adaptation Artifacts overview → \(/spatial-adaptation\)](/spatial-adaptation)