

Spatial Adaptation Artifacts: Runtime Skill Loading With Admissibility Gating

by [Nick Clark](#) | Published April 25, 2026

Skill Marketplaces Without an Architecture

Anthropic Skills, OpenAI Custom Actions, Google Gemini Extensions, and Microsoft Copilot Studio are converging toward a common shape: runtime-loadable adaptation artifacts that extend a base model's behavior. The shape is the right one — production AI deployments need this, and the alternative (retraining the base model for every task) is structurally infeasible.

But the platforms ship without an architecture for the questions that production deployment requires: which skill applies in this context, who certified the skill is safe, what other skills must be active for this skill to function correctly, what happens when the certifying authority revokes the skill, and how do we audit skill-routing decisions after the fact?

The current answer is that these questions are platform-internal: the platform operator handles skill admission, the operator's policy decides which skills are active, and the operator's logs record what happened. This is the same platform-operator model that the governed marketplace primitive (Article 9) eliminated for commodity exchange. It has the same problems here: operator failure compromises every consumer, operator policy preferences distort the market, no cross-platform skill portability.

1. The Primitive: Signed Adaptation Artifacts With Consumer-Side Certification

Spatial adaptation artifacts are runtime-loadable behavioral modifications signed by their authoring authority. The artifact format is technique-agnostic: LoRA fine-tuning weights, RAG retrieval indices, prompt-injection configurations, mixture-of-experts adapter routing, hybrid combinations, and emerging techniques.

Each artifact carries: the authoring authority's credential, the artifact content (the actual weights, prompts, or configurations), declared dependencies on other artifacts, declared model compatibility (which base models can host it), declared scope (which tasks the artifact applies to), and declared training provenance (which data was used and under what governance).

Critically, certification is consumer-side rather than authoring-side. The consumer (the system loading the artifact for inference) runs the artifact through a sandbox evaluation against the consumer's own admissibility policy before activation. The artifact's authoring authority signs what it is; the consuming authority certifies whether to activate it.

2. Admissibility Gate as Skill Router

At inference time, a request enters the system, and the admissibility gate routes the request to the appropriate set of active skills. The gate is the same composite admissibility evaluator used elsewhere in the architecture: it consumes the request, the available active skills, the consumer's policy, and operational context to produce a routing decision.

Routing is graduated rather than binary: a request may be routed to multiple active skills with weighted contribution, may activate a contextually-appropriate skill that wasn't already loaded (subject to admissibility), may defer to a higher-authority skill when policy requires, or may decline routing when no available skill is admissible.

Admissibility-as-router unifies what current platforms split between two layers: skill selection (which skills can fire) and inference routing (which skills do fire on this input). The unified gate evaluates both questions in a single deterministic step.

3. Always-Active Personal Layer

Every consumer maintains an always-active personal layer that is exempt from the de-weighting that admissibility may apply to other skills. The personal layer carries the consumer's own preferences, identity, history, and authority — the irreducible 'self' of the consuming system — and contributes to every inference at full weight.

Personal-layer carve-out solves a recurring problem in marketplace-style skill ecosystems: third-party skills can dominate, override, or even adversarially manipulate the consumer's intent. The personal layer prevents this structurally: even if a third-party skill is fully admitted, the personal layer remains sovereign over the consumer's intent and will modulate the third-party skill's contribution accordingly.

The personal layer is itself a signed adaptation artifact, but signed by the consumer's own authority and held in a privileged position by the admissibility evaluator. Its content is the consumer's own; its authority is the consumer's own; its admissibility is the consumer's own.

4. Dependency Chains and Cascade Deactivation

Real skill ecosystems have dependencies. A medical-coding skill may depend on a clinical-vocabulary skill, which depends on a base medical-language adaptation. A legal-research skill may depend on jurisdictional-corpus skills. A code-review skill may depend on language-specific syntax skills.

Each artifact declares its dependencies in its credentialed metadata. The admissibility gate evaluates dependency satisfaction before activation: a skill cannot fire if its declared dependencies are unmet.

Cascade deactivation handles dependency revocation: when an authority revokes a skill, all skills that declared dependency on it deactivate as well, transitively. The cascade is recorded in lineage; consumers see the cascade as a credentialed observation and can re-activate alternatives, request replacements, or operate in degraded mode.

5. Cross-Model Portability

Adaptation artifacts authored against one base model often need to apply across model versions, vendors, or substitutions. A regulatory compliance skill authored for one LLM should remain usable when the deployment migrates to another. Current platforms handle this poorly: artifacts are typically locked to a specific base model and require re-authoring at migration.

The governed primitive supports cross-model portability through declared compatibility metadata: the artifact specifies which base-model classes it is compatible with, which adaptation techniques it uses, and what compatibility evidence supports the claim. Compatibility is itself a credentialed observation; an authority can sign 'this artifact tested compatible with these models' for downstream consumers.

When a consumer migrates between base models, compatible artifacts continue to operate without re-authoring. Incompatible artifacts deactivate (cascade deactivation) and are replaced by compatible alternatives or degraded-mode operation.

6. Federated Skill Training

Skills improve through use. The governed primitive supports federated skill training: a deployed skill records its performance against admissibility-evaluated outcomes; the records propagate (with privacy and rights governance) back to the authoring

authority; the authority improves the skill and publishes a new credentialed version; consumers re-evaluate the new version under their admissibility policies.

Federated training is governance-credentialed, not blockchain-mediated. The training authority and the contributing consumers operate under the same governance-chain framework that admits any other observation. Privacy preservation is structural: the contributing consumers control what observations they release, the training authority signs what it received, and the audit trail covers every contribution.

This produces an evolving skill ecosystem where authority-credentialed skills continuously improve under operating use, rather than being frozen at training time and replaced wholesale.

7. Decentralized Mesh-Distributed Distribution

Skill artifacts distribute through the governed mesh (Article 1). There is no central skill marketplace operator. Authoring authorities publish credentialed artifacts; consumers subscribe to authorities they admit; artifacts propagate through fixed infrastructure relay, peer-to-peer transmission, and mobile store-and-forward.

Mesh distribution composes with intentional-disconnect (Article 13): a consumer in an isolated environment can pre-stage credentialed artifacts before disconnect, operate during disconnect with the staged artifacts, and reconcile any updates after reconnect. This serves expeditionary, defense, maritime, and other operational contexts where centralized skill distribution is infeasible.

Distribution authority is decentralized: any authority with relevant standing can publish artifacts, and consumers choose which authorities to admit. No platform operator gates the skill economy.

8. What This Is Not

This is not the App Store / Google Play / HuggingFace Hub. Those have a single platform operator and centralized policy gating. The governed primitive operates without an operator.

This is not Anthropic Skills, OpenAI Custom Actions, Google Gemini Extensions, or Microsoft Copilot Studio as currently shipped. The governed primitive could underpin those products with the architecture they currently lack: consumer-side certification, dependency-chained cascade deactivation, cross-model portability, and admissibility-gate-as-skill-router.

This is not LoRA / PEFT / sigstore alone. Those are component techniques that the governed primitive composes; the architecture is broader than any single technique.

Conclusion

Spatial adaptation artifacts provide runtime skill loading with consumer-side sandbox certification, admissibility-gate-as-skill-router, always-active personal layer, dependency chains with cascade deactivation, cross-model portability, federated training, and decentralized mesh distribution.

Disclosed under USPTO provisional 64/049,409, the primitive provides the missing architecture for the AI-agent skill marketplace that current platforms are building ad hoc. It composes with the governed mesh, mesh-distributed firmware updates, and the five-property governance chain umbrella.