

# Cross-Model Skill Portability Across Base-Model Substitutions

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

## What Compatibility Metadata Specifies

Each adaptation artifact declares its compatibility scope: which base-model classes it has been validated against, which adaptation techniques it uses, and what compatibility evidence supports the claim. The compatibility evidence is itself a credentialed observation — an authority (the artifact author, an independent compatibility-testing authority, the consumer's own validation) signs the claim that this artifact functions correctly with this base model.

Compatibility is not a binary property. An artifact may be compatible with one base model at full functionality, with another at reduced functionality, and incompatible with a third. The metadata captures this graduated compatibility, and the admissibility gate evaluates it during activation.

## Why Vendor Lock-In Is the Default Today

Current agent platforms produce vendor-locked artifacts by structural default. An OpenAI Custom Action targets GPT-class models. An Anthropic Skill targets Claude-class models. A Google Gemini Extension targets Gemini. Migrating between vendors requires re-authoring the artifacts, which destroys the value of accumulated artifact investment.

The lock-in is architecturally unnecessary. The technical content of many adaptation artifacts (RAG indices, prompt configurations, even most LoRA fine-tunes) is largely vendor-independent. The lock-in lives in metadata and platform conventions rather than in the artifacts themselves.

## **How the Architecture Achieves Portability**

Cross-model portability is achieved through standardized artifact metadata, declared compatibility evidence, and consumer-side certification that operates uniformly across vendors. When a consumer migrates between base models, the admissibility gate re-evaluates each artifact's compatibility metadata against the new base model. Artifacts compatible with both continue to operate; artifacts incompatible with the new model deactivate cascade-style.

Replacement routing is automatic where alternative artifacts are available. The consumer's policy specifies preferred replacement criteria (same authoring authority, same functional scope, similar compatibility evidence); the architecture identifies candidates and proposes activation. The migration is configuration rather than re-engineering.

## **What This Enables for AI Sovereignty**

Enterprises increasingly want vendor-independence in their AI infrastructure. Multi-vendor agent strategies (using Claude for some tasks, GPT for others, Gemini for others) are blocked by vendor-locked artifacts. Cross-model portability removes the structural barrier.

The pattern also supports model migration over time. As base models are deprecated, replaced, or upgraded, accumulated artifact investment migrates rather than being abandoned. The patent positions the primitive that the multi-vendor enterprise agent ecosystem requires for the durability that current platforms cannot provide.

