

Beyond Model Nationality: What Makes an AI Supplier Acceptable?

An executive guide to origin, operational control and the evidence behind enterprise approval

Enterprise AI Adoption Gap · Research brief 10
October 5, 2026

© 2026 Eldris Inc. All rights reserved.

Executive brief

A U.S. enterprise considering a foreign AI supplier faces several decisions at once. Is the proposed use permitted? Can the system do the work? Who receives the information? Who controls changes? Who takes responsibility when the service fails?

Nationality can matter to those decisions. It can identify a relevant jurisdiction, trigger a buyer's sourcing policy or influence perceptions. But it cannot, by itself, describe the data flows and obligations of a particular deployment. A foreign-developed model accessed through its developer's service and a separately operated copy of its published weights present different operating arrangements.

The unit of approval should be a defined supplier arrangement for a defined use. Record the model and version, contracting party, hosting and access conditions, permitted data, operational responsibilities and evidence supporting the decision.

This report addresses U.S. commercial enterprises considering an internal assistant that drafts answers from company policy documents, with employees checking answers before use. It does not establish eligibility for government procurement, regulated workloads or any named supplier. Those decisions require their own applicable rules and evidence.

For buyers, the immediate value is a clearer comparison: distinguish a binding exclusion from a concern that better controls or documentation could address. For foreign suppliers, the value is a more precise market-entry question: which eligible buyers could approve which arrangement, and what is preventing that approval?

Three principles guide the assessment:

- Separate model origin, service operation, data access and contractual responsibility.
- Treat mandatory exclusions as eligibility gates; assess eligible options against workflow requirements.
- Test whether proposed changes resolve the buyer's actual objection before investing in new infrastructure or positioning.

Public sources explain why these distinctions matter. They do not tell us how many U.S. buyers would approve a foreign supplier or what change would persuade them. That remains a bounded empirical question.

Specify what the buyer is approving

Start with the workflow, not a country label. In this report's example, employees ask an assistant questions about internal company policies. The assistant retrieves approved documents, drafts an answer and identifies supporting passages. A person checks the result before acting on it.

Define the document categories, user roles, permitted questions and consequences of a wrong answer. Exclude sensitive employee records from the initial scope if they are unnecessary. An assistant using public policy material is a different proposition from one receiving confidential investigations or individual personnel files.

Then identify the actual offering. A buyer might procure a developer-operated API, an application that sends requests to several model providers, a managed deployment operated by another vendor, or published weights run in its own environment. Each arrangement creates a different set of counterparties and responsibilities.

Record model developer and ownership separately from the organization providing the service. Identify the application vendor, cloud operator, relevant subcontractors and support access. A U.S. reseller address does not establish where inference occurs or who can inspect the content.


NIST's Generative AI Profile recommends use-case-based supplier assessment, defined contracts and service requirements, inventories of third parties with access to organizational content, and contingency planning. These are governance recommendations, not certification that a supplier is acceptable. [NIST profile](#).

Decision implication: Require a deployment description precise enough for business, security and procurement teams to evaluate the same proposition.

Four questions hidden inside model nationality

Four questions behind supplier acceptance

Assess a defined deployment for a defined workflow



Origin remains relevant. Hosting and assurance answer different questions.
Original Eldris framework; no country rankings or measured buyer preferences.

© Eldris Inc

Figure 1. An original Eldris assessment map. The dimensions are distinct but can interact. This is not a ranking of countries or a measured model of buyer preferences.

Origin and ownership

Who developed the model, who controls the supplier and which jurisdictions are relevant to those entities? Identify dependencies on upstream developers and the provenance of the selected version. Origin can remain material even when inference occurs elsewhere, including where a buyer's policy expressly covers models or ownership.

A country label may also stand in for an unspoken concern about recourse, continuity or content behavior. Ask which concern the decision maker means. That clarification does not remove an applicable restriction; it makes the assessment actionable.

Operation and access

Where are inference, storage, backups and support performed? Who can access prompts, retrieved documents, outputs and diagnostic records? How do routing, application integrations and optional features change those paths?

A location claim should identify its scope. A U.S. inference region does not, on its own, establish the location of every log or support activity. Nor does it establish that no other entity can access data.

Assurance and accountability

What evidence verifies the promised behavior? Which controls can the buyer inspect or enforce? Who is responsible for incidents, updates, deletion and recovery? Which obligations are contractual, and which are statements in marketing or documentation?

An audit report must cover the relevant entity, service, period and controls. A badge or broad company certification cannot answer every deployment-specific question.

Workflow performance and continuity

Does this version answer the buyer's policy questions accurately, reveal its sources and handle unsupported requests appropriately? Can the buyer monitor changes, maintain service and move to an alternative when needed?

A model benchmark does not establish performance inside the buyer's retrieval system. A capable model with an unsuitable operating arrangement may be ineligible; an eligible arrangement with inadequate task performance may be unusable.

Decision implication: Convert each concern into a requirement, evidence request or exclusion. Avoid a single trust score that allows strength in one dimension to obscure a failure in another.

A model and its hosted service are different objects

DeepSeek provides a concrete illustration of the distinction. Its privacy policy, last updated February 10, 2026, identifies a China-registered service controller and states that personal data is collected, processed and stored in the People's Republic of China for the covered services. The policy has a defined scope; it is not a description of every application using DeepSeek-derived weights. [DeepSeek policy](#).

Separately, the DeepSeek-R1 model card publishes instructions for local serving and identifies the repository and weights as MIT-licensed. It also identifies different upstream origins for distilled variants. A buyer must inspect the exact artifact and applicable terms rather than extrapolate from the family name. [DeepSeek-R1 model card](#).

These sources establish that a developer-operated service and separately served weights are distinguishable. They do not establish that a self-hosted installation is secure, that every variant has identical obligations or that a named deployment qualifies for a buyer's use.

A buyer-controlled installation can change who operates inference and receives requests. Its actual boundaries still depend on the surrounding application, serving software, network access, telemetry and maintenance. The buyer takes on responsibilities for configuration, updates, security and capacity. Published weights do not supply those functions automatically.

One model lineage, different operating arrangements

Generic comparison · Verify each actual offering

DEVELOPER-OPERATED	SEPARATELY MANAGED	BUYER-OPERATED
Developer supplies the model service.	Another operator serves the selected model.	Buyer runs selected published weights.
Check service terms, data access, retention and incident obligations.	Check operator controls, upstream dependencies and contract coverage.	Verify serving software, network paths, access and telemetry.
Buyer still owns the application, permitted use and integration.	Responsibilities depend on the service and the configured deployment.	Buyer assumes operation, security, updates and capacity responsibilities.

Model provenance remains a separate assessment in every arrangement.
These are not security rankings or endorsements of named products.

Original Eldris analysis. Published weights do not provide managed operations automatically.

© Eldris Inc

Figure 2. Generic deployment arrangements, not descriptions or endorsements of named products. There is no implied ordering from least to most secure. Verify the actual data paths and responsibilities in each offering.

Managed platforms further illustrate why endpoint-specific evidence matters. AWS's Bedrock retention documentation distinguishes retention settings and states that content is not currently shared with model providers. It also describes circumstances involving AWS retention or review and cross-region processing. Those statements concern the documented service behavior; they are not a promise of zero retention for every model, setting or connected application. [AWS documentation](#).

The business lesson is to verify the proposed service, model, region, features and settings. A model developer's country does not fully specify that configuration. A platform's domestic identity does not establish that every configuration meets the buyer's requirements.

Decision implication: Compare complete deployment arrangements. Moving inference changes some risks and responsibilities while leaving model provenance and other dependencies to assess.

Separate exclusions from remediable objections

A supplier should know whether a prospect cannot approve an arrangement or has not yet received enough evidence to approve it. Those conditions call for different investments.

A binding law, procurement rule, contractual obligation or internal sourcing policy may exclude the proposed entity, model, origin or use. Verify applicability with the buyer's responsible teams. Do not assume that moving hosting resolves a restriction whose scope concerns something else.


This report does not catalog current U.S. restrictions or give legal clearance. It recommends making the eligibility determination explicit before testing preference or presenting a technical workaround.

Among eligible options, an objection may concern an evidence gap: unclear retention, unidentified support access, missing incident commitments or uncertain update practices. Another may concern a capability gap: inadequate performance, unavailable controls or service capacity. A third may concern perception after requirements are satisfied.

These distinctions are investigative categories, not a claim about the prevalence of each barrier. The same buyer can have several objections, and an organization can have a stricter standard than any individual respondent.

From eligibility to bounded approval

Resolve the condition that prevents the next decision



A buyer preference is research evidence, not institutional authorization.
Original Eldris decision sequence; not a legal clearance or validated predictor.

© Eldris Inc

Figure 3. An original Eldris decision sequence. Resolve mandatory eligibility first, then requirements and evidence, then bounded deployment. Passing a preference study does not authorize enterprise use.

For the policy assistant, missing documentation about data access may justify requesting evidence. A service that cannot provide a required boundary may require a different architecture or supplier. A demonstrated exclusion requires a different eligible proposition, if one exists.

Decision implication: Identify the objection and its owner before paying to fix it. Better messaging cannot substitute for a missing control or remove a binding exclusion.

What foreign suppliers can change

Market-entry planning should connect each investment to an eligible segment and an approval requirement. Establish which parties make the technical, contractual and business decisions, and who can veto the arrangement.

Make the operating arrangement legible

Provide a current deployment description, data-flow inventory and responsibility matrix. Explain model provenance, service entities, locations, retention, remote access and relevant integrations. Label optional features and distinguish public defaults from negotiated terms.

The useful test is whether the buyer can identify what will happen to its documents and who is accountable. A general statement about trust or sovereignty leaves that work to the prospect.

Offer a verifiable arrangement for the intended workload

A managed U.S. deployment, buyer-hosted option or reduced-data workflow may address particular requirements. Each option must actually exist and be supportable. Describe the remaining dependencies and operating costs.

Do not market a domestic endpoint as a universal resolution. If the objection is model provenance, ownership or a mandated sourcing rule, relocating infrastructure may leave it intact. If the objection is operational access, a verified boundary may matter more than a change in branding.

Build evidence for the buyer's decision

Supply task-relevant evaluations, incident procedures, change commitments and accessible documentation. Let the buyer test relevant configuration and inspect the supporting evidence within appropriate access arrangements.

An evaluation should include realistic questions, incomplete sources, conflicting policies and requests that should be escalated. Test the complete assistant and its review process, not only the underlying model.

Make continuity and exit credible

Show how the customer can preserve required records, stop access, recover service and transfer the workflow. Identify which components depend on the supplier and what an alternative would require. Evaluate transition effort rather than treating portability as a statement of intent.

These changes can make an offering easier to assess. This report does not establish their sales lift. Investment should follow evidence that the change resolves a material barrier for buyers the supplier can serve.

Decision implication: Fund improvements that connect to a specific approval requirement and validate the result. Separate technical eligibility, institutional approval and commercial attractiveness.

Research the decision that could change

A commissioned study should begin with one sponsor decision. For example: should a foreign supplier invest in a U.S.-operated deployment for internal policy assistants sold to U.S. commercial enterprises?

Define the target segment, eligible use and actual architecture options. Exclude buyers whose restrictions cannot be addressed by those options from the primary approval comparison, while reporting how that screening limits the addressable segment. Otherwise, an apparent preference result may mix remediable barriers with buyers who cannot purchase.

Recruit people with relevant, verified responsibilities. Workflow owners, security reviewers and procurement participants answer different parts of the decision. General U.S. adults cannot represent an enterprise approval committee merely because they use AI at work.

Use discovery interviews to identify objections and approval processes. Then compare a small number of realistic offers, varying origin, operation and assurance only where the combinations are feasible and interpretable. Present the same demonstrated task capability when estimating the effect of those conditions. A comparison that changes performance and hosting together cannot isolate a hosting effect.

For stated-choice experiments, clearly label hypothetical offers. Test comprehension of what hosting and access claims mean. Record conditional shortlist decisions, required evidence and willingness to proceed to evaluation. Measure whether respondents accept a proposition only after an improvement that the supplier can actually deliver.

An origin disclosure comparison can help distinguish an origin-related response from reactions to operating conditions. Brand familiarity and country disclosure need separate consideration. A disclosed country effect does not establish its cause: respondents may infer unstated quality, legal or support differences. Ask about those interpretations and report the limit.

Use behavioral evidence where possible: voluntary technical evaluation, completion of a procurement step or a documented conditional approval. An individual survey response is a research outcome, not an authorized purchasing commitment. A panel can support screening and decision experiments; actual committee processes and customer evaluations are needed to establish institutional approval.

The deliverable should distinguish three questions: what is eligible, what the buyers studied would advance, and what the sponsor can economically deliver. Report uncertainty, exclusions and remaining objections. Do not translate a panel's approval percentage directly into a U.S. market-size forecast.

Decision implication: Commission research to choose among real investments. Favor evidence that changes an architecture, assurance or segment decision over a broad sentiment ranking of foreign AI.

An approval record the business can use

Prepare a short record for each proposed arrangement:

1. **Scope:** workflow, users, permitted data and consequences of failure.
2. **Identity:** model version, provenance, supplier ownership, service entities and dependencies.
3. **Eligibility:** applicable restrictions, buyer policies, decision owner and documented determination.
4. **Operation:** locations, access, routing, retention, updates and tested configuration.
5. **Evidence:** task performance, assurance scope, unresolved gaps and contractual commitments.
6. **Economics and continuity:** complete operating cost, support responsibilities, fallback and exit effort.
7. **Decision:** authorized use, conditions, accountable owner and reassessment triggers.

For the internal policy assistant, approval might be limited to specified documents, employee review and a verified deployment configuration. Expansion to new data, autonomous action or another endpoint is a new proposition to evaluate. A supplier's acceptance is conditional on the arrangement that earned it.

The strategic question is therefore more precise than whether U.S. enterprises trust foreign AI: **which eligible buyers will approve which supplier arrangements, for which work, on what evidence?** Answering that question can guide market entry and procurement. A nationality ranking cannot supply those decisions alone.

Method and limitations

This brief selectively reviews primary sources checked on October 5, 2026. It contains original Eldris analysis of supplier assessment and research design. It includes no proprietary buyer survey, measured approval rates, market forecast or independently tested supplier configuration.

Provider documentation establishes stated policies and published deployment options, not independently verified operation or contractual suitability. The DeepSeek example illustrates the difference between a covered service and separately served weights. The AWS example illustrates configuration-specific retention and access questions. Neither is an endorsement or an eligibility determination. Policies, product behavior and available offerings can change.

NIST provides governance guidance. The assessment map, deployment comparison, decision sequence and approval record are analytical aids, not validated predictors of buyer behavior. They do not assign country risk scores or imply that all origins are interchangeable.

The scope is a bounded U.S. commercial enterprise workflow. Applicable legal, regulatory, procurement and contractual requirements need current review for the actual buyer and offering. The report does not establish that domestic hosting removes jurisdictional exposure, that published weights are inherently safe or that one supplier's documentation applies to another.

The next evidence step is a study of an actual approval or market-entry decision, followed by verification of the arrangement the buyer would use.

Sources

1. **NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1.** July 2024. GOVERN 6.1 and 6.2 address third-party assessment, contracts, access inventories and contingencies. [Profile](#).
2. **DeepSeek. Privacy Policy.** Last updated February 10, 2026; checked October 5, 2026. Scope, service controller and personal-data processing and storage location. [Policy](#).
3. **DeepSeek-AI. DeepSeek-R1 Model Card.** Checked October 5, 2026. Published serving instructions, license statement and distilled-model provenance. [Model card](#).
4. **AWS. Data retention - Amazon Bedrock.** Checked October 5, 2026. Retention modes, provider access, AWS review and cross-region handling. [Documentation](#).