

# Agent interoperability protocols as the next standard essential patent frontier

By Roozbeh Gorgin, Esq., and Ananth J. Josyula, Sterne, Kessler, Goldstein & Fox PLLC

APRIL 24, 2026

Artificial intelligence (AI) is moving beyond a phase of stand-alone chatbots to “agentic AI.” Agentic AI systems do not merely generate outputs. They can take actions, invoke tools, retrieve information, update records, and coordinate multi-step workflows across software environments.

As that shift accelerates, a less discussed but legally significant development is also taking place: the AI ecosystem is beginning to standardize — not around any particular model, but around the protocols that allow AI agents to interact with tools, data sources, enterprise software systems, and one another.

That architectural shift has important patent consequences; when industries converge on shared interoperability mechanisms, patents covering those mechanisms can become commercially unavoidable. In other words, agent interoperability protocols may become the next terrain on which standard-essential patents disputes emerge.

This article addresses emerging issues as agentic AI systems increasingly rely on shared interoperability protocols to act across software boundaries, and explains how those protocols may become the next frontier for standard essential patents (SEPs).

## Key takeaways

- **Agentic AI is driving standardization at the interoperability layer.** As agentic AI systems increasingly interact with tools, data sources, enterprise software systems, and other agents, shared interoperability protocols are becoming more important than any single model.
- **AI patent strategy should adapt now.** Patent value may increasingly depend on claiming protocol mechanics, such as discovery, permission controls, invocation, state and output handling, and fallback behaviors, rather than focusing only on model internals.
- **Interoperability protocols may create the next SEP battleground.** If adoption hardens around common interoperability protocols, patents covering protocol mechanics may become difficult to avoid in practice.

- **Licensing and enforcement may be more complex than in traditional SEP settings.** In agent ecosystems, implementation is often distributed across multiple parties, raising difficult questions about infringement, licensing targets, and royalty allocation.
- **Open ecosystems may intensify FRAND-like tensions (fair, reasonable, and non-discriminatory).** Top-down interoperability initiatives may increase adoption while also creating pressure around licensing transparency, access, and patent monetization.

## Agentic AI favors protocols

Earlier AI products were largely self-contained; a user submitted prompts, the model generated responses, and any iteration occurred entirely within that closed loop. Agentic AI changes that architecture. Its value lies in acting outside the model — invoking external capabilities through standardized interfaces to query a database, call enterprise software, generate and send reports, update CRM records, or trigger internal workflows.

Once AI systems operate across software boundaries, interoperability becomes essential. Agents must know what tools and capabilities are available, what permissions apply, how to invoke them, how state and outputs are handled, and what happens when an action fails. Those are not optional features; they are foundational requirements of a functioning agent ecosystem.

Custom integrations may solve these issues in isolated instances, but they do not scale. Instead, the predictable market response is protocolization: shared technical methods governing capability discovery, permission controls, invocation, auditing, and coordination across multiple participants.

That process is already underway; Model Context Protocol (MCP) is emerging as an open standard for connecting models and agentic applications to external tools and data sources. Agent2Agent (A2A) is aimed at a related but distinct problem of how agents discover one another, exchange messages and artifacts, and collaborate across systems.

Together, they illustrate that standardization pressure is landing not on the model layer, but on the interoperability layer.

### What may actually become standardized

Standardization is unlikely at the model layer itself. Models may remain proprietary, differentiated, and fast-moving.

Instead, standardization pressure is emerging around recurring interoperability functions, including:

- capability discovery;
- permission controls;
- invocation;
- state and output handling; and
- coordination across sessions, systems, and environments.

From an IP perspective, those protocol mechanics matter because they create top-down compatibility expectations across vendors and systems. Once those expectations become commercially necessary, patents covering such protocol mechanics may begin to look functionally essential — even before formal adoption by a standards body.

---

### *When industries converge on shared interoperability mechanisms, patents covering those mechanisms can become commercially unavoidable.*

---

That is a familiar pattern; many important standards begin as *de facto* conventions, gaining traction through developer adoption, reference implementations, or platform momentum before any institution-led governance structure emerges. In other words, commercial dependence often arrives before formal standardization.

### Why this matters for patent drafting

For patent prosecutors and portfolio managers, the rise of agent interoperability standards has a practical implication: AI patent value may increasingly turn not on the model itself, but on the technical rules that compliant systems must follow to participate in an ecosystem. That changes how patent applications should be drafted from the outset.

In the standards context, a patent becomes valuable not merely because it covers a useful feature, but because it reads on functionality that implementers cannot realistically omit and still remain compliant.

Therefore, the drafting exercise is not just about claiming an invention at a high level; it is about identifying which parts of a protocol are likely to become mandatory, or at least commercially unavoidable, and then describing those parts with enough technical specificity to later support meaningful claims.

That point is often missed in AI filings, with many applications drafted around broad concepts such as “an AI agent invoking a tool” or “an autonomous system performing a multi-step workflow.” While those formulations may sound expansive, they do little to position a portfolio around compliance-critical mechanics that are most likely to harden into standardized requirements.

---

### *Once AI systems operate across software boundaries, interoperability becomes essential.*

---

If a protocol later emerges, essentiality usually will not attach to the general idea of tool use. Instead, it will attach to the required mechanics of interoperability — how capabilities are discovered, how permissions are checked, how invocation requests are structured, how state and outputs are managed, how failures are handled, how security and auditability are preserved, and how coordination is achieved across implementations.

Current protocol efforts help make that concrete; if an ecosystem converges on standardized protocol mechanisms for capability discovery, permissions, invocation, state and output handling, and coordination across systems, then patents covering those protocol mechanics may matter far more than patents on a general “AI assistant using software tools” concept. The patent value lies in the mandatory handshake, not the abstract aspiration.

Accordingly, applications should be drafted with an eye toward the control points where interoperability standards are most likely to harden. The points below highlight drafting considerations that preserve long-term standard-essentiality value.

First, the specification should describe the protocol-facing architecture in concrete terms. Rather than stopping at a functional narrative, the application should explain how the agent interacts with the discovery layer, permission controls, invocation mechanisms, state and output handling components, and any coordination, validation, or logging subsystems.

Standards tend to form around interfaces between components. Applications that define those interfaces clearly are better positioned to support claims that later map onto standard-mandated behavior.

Second, applications should emphasize required sequences and state transitions. Many standards are defined not only by what data is exchanged, but by when and in what order steps must occur.

For example, a protocol may require capability discovery before invocation, permission checks or confirmations before execution of sensitive actions, credential validation before data access,

and structured error reporting before retry. Those sequencing requirements can become the most important claim hooks.

Third, applications should distinguish between mandatory behaviors and optional enhancements. Standards typically define a baseline compliance layer that implementers must satisfy, with discretionary implementation choices layered on top.

Claims directed only to high-level optional or vendor-specific improvements may still have value, but they are less likely to become essential. By contrast, claims that track baseline compliance functions have greater SEP potential.

---

*The patent value lies in the mandatory handshake, not the abstract aspiration.*

---

Fourth, applications should include embodiments focused on error handling, compatibility negotiation, and security enforcement. These issues are often underdeveloped in drafting, but they become central during adoption.

A standard may permit flexibility in basic invocation or request format, while tightly constraining how a compliant implementation handles malformed inputs, conflicting permissions, unavailable tools, version mismatches, or unsafe actions.

Fifth, drafters should think in terms of claim sets that target different actors in a distributed ecosystem. Agent interoperability rarely resides in a single box.

For example, one actor may operate an agent host, another actor may expose a callable tool, another actor may manage identity or permissions, and an enterprise customer may configure policy enforcement. Applications that describe only one side of those interactions may introduce attribution — and correspondingly enforcement — challenges later.

Taken together, the drafting goal is to focus not only on what is novel, but on what other implementers will eventually need to do in the same manner to remain interoperable.

In practice, that typically means supporting claims directed to compliance-critical exchanges and baseline behaviors, such as capability registration, version negotiation, confirmation requirements for privileged actions, propagation of authorization attributes, and defined fallback behavior on failure.

It also means populating the specification with protocol-oriented examples such as discovery sequences, permission exchanges, handshakes, and failure modes so later claim amendments can map onto an emerging standard without presenting new-matter problems.

Thus, in an interoperability-driven AI market, patent drafting cannot remain model-centric; portfolios built only around model internals may miss the layer where ecosystem lock-in and licensing leverage actually emerge.

## Open ecosystems and FRAND tension

Another complicating factor is that many agent interoperability efforts are emerging in open-source or institution-led ecosystems. Such efforts include Linux Foundation initiatives like the Agentic AI Foundation (governing MCP, goose, and AGENTS.md) alongside the A2A Project, AGNTCY, and Agentgateway.

In practice, developers and vendors treat such open interoperability projects like shared infrastructure. As a result, they expect the core compatibility rules to be easy to implement, broadly available, and not blocked by surprise licensing demands.

That expectation can collide with patent owners' incentives to monetize once patents read on protocol features that become widely adopted — and therefore hard to avoid in practice. At that point, FRAND-like tension can emerge for patent owners as a technology becomes functionally unavoidable and stakeholders begin to expect licensing behavior that looks like fair, reasonable, and non-discriminatory access.

That tension creates strategic risk; an aggressive assertion posture of any such patents early on may reduce the uptake of the very protocol that would otherwise make a patent position more valuable. Patent owners in these situations may want to develop licensing strategies that balance adoption versus monetization.

Nevertheless, if agent interoperability protocols ultimately become governed by formal standards bodies or industry consortia, the bylaws and IP policies adopted at that stage may significantly dictate what patenting, disclosure, and licensing practices are permitted.

## Patent policy and remedies implications

If agent protocols become essential infrastructure before clear disclosure and licensing norms develop, the agent ecosystem could replay familiar SEP problems — hold-up concerns, royalty stacking, opaque licensing demands, and antitrust scrutiny — but in a faster-moving, software-centric market.

Remedies may also look different than in classic product disputes. In software ecosystems, exclusion may take the form of disabling compatibility, withholding certification, denying access to shared registries, or otherwise breaking interoperability. Those measures can be commercially powerful even without traditional product injunctions.

All of that suggests a need for earlier governance — including patent disclosure expectations, baseline licensing commitments, and a clearer separation between required interoperability layers and discretionary extensions.

## Conclusion

Agentic AI is not simply a more capable version of the chatbot era. It represents a different system architecture that depends on interoperability at scale. As the industry converges on

shared protocols to make that interoperability possible, those shared protocols are likely to become commercially unavoidable. That makes them fertile ground for the next generation of SEP disputes.

For lawyers advising on AI patent strategy, licensing, competition, or policy, the critical insight is this: the next SEP frontier in AI may be defined less by model intelligence than by the protocols that allow autonomous systems to work together.

### About the authors



**Roozbeh Gorgin (L)** is counsel in **Sterne, Kessler, Goldstein & Fox PLLC**'s electronics practice group. His practice focuses on patent prosecution, post-grant proceedings and pre-litigation counseling. Based in Washington, D.C., he can be reached at [rgorgin@sternekessler.com](mailto:rgorgin@sternekessler.com). **Ananth J. Josyula (R)** is a technical specialist in the firm's electronics practice group, where he assists in patent preparation and prosecution. His areas of expertise include artificial intelligence, data science, software systems and blockchain technologies. He is also based in Washington, D.C., and can be reached at [ajosyula@sternekessler.com](mailto:ajosyula@sternekessler.com).

This article was first published on Reuters Legal News and Westlaw Today on April 24, 2026.