The Agentic AI Foundation: When OpenAI and Anthropic 'Donate' Open Source Standards, Who Really Owns the Protocol?

The biggest AI companies just gave away their most valuable infrastructure protocols for free. Here's why that should terrify you more than if they'd kept them locked up.

The Gift That Keeps on Taking

On December 9, 2024, something unprecedented happened in the AI industry. OpenAI, Anthropic, and Block stood together under the Linux Foundation banner and announced they were *donating* their core agentic AI protocols to a new neutral governance body: the Agentic AI Foundation.

The press releases read like a masterclass in tech philanthropy. Words like "interoperability," "open standards," and "community-driven development" peppered every paragraph. Anthropic contributed their Model Context Protocol (MCP). OpenAI handed over AGENTS.md. Block donated their goose framework. Three foundational technologies that will define how autonomous AI agents communicate, coordinate, and execute tasks for the next decade—all wrapped up with a bow and placed under "neutral" stewardship.

If you felt a warm glow reading that, I need you to pause and consider a single number: $350,000.

That's the annual price tag for a platinum membership in this new foundation. And that platinum membership comes with something money can't buy anywhere else: a guaranteed seat on the Governing Board.

Eight companies ponied up immediately. AWS. Anthropic. Block. Bloomberg. Cloudflare. Google. Microsoft. OpenAI. The complete founding roster reads like a who's-who of the entities that already control the AI supply chain from compute to deployment.

So here's the question nobody in the breathless press coverage seemed to ask: When the companies donating the protocols are the same companies paying for board seats that control those protocols, what exactly got donated?

Understanding What's Actually at Stake

Before we dissect the governance structure, let's establish why these particular protocols matter so much that their ownership structure should concern every developer, enterprise, and regulator paying attention.

The Model Context Protocol (MCP)

Anthropic's MCP isn't just another API specification. It's a protocol for standardizing how AI models handle and share contextual information across multi-agent systems. Think of it as the TCP/IP of agentic AI—the foundational layer that determines how autonomous systems pass state, maintain memory, and coordinate actions.

The stated goal is preventing vendor lock-in. The reality is that whoever controls the MCP specification controls the compatibility layer for every multi-agent system built on top of it.

AGENTS.md Standards

OpenAI's AGENTS.md has already achieved remarkable traction. Released in August 2024, over 60,000 projects have adopted it, including GitHub Copilot and Cursor. It's becoming the de facto standard for how AI agents describe their capabilities, limitations, and interaction patterns.

This isn't speculative infrastructure. This is already-dominant infrastructure being moved into a governance structure that its creator will continue to control.

The Goose Framework

Block's goose provides the execution layer—the scaffolding for building, deploying, and managing AI agents in production environments. Combined with MCP for communication and AGENTS.md for self-description, you have a complete stack for agentic AI development.

Together, these three projects represent the full protocol layer for autonomous AI systems. Not the models themselves, but something potentially more valuable: the rails those models run on.

The Linux Foundation Governance Model: A Primer

To understand what's really happening here, you need to understand how Linux Foundation projects actually work. This isn't the Apache Software Foundation's meritocratic model or the IETF's rough consensus approach. The Linux Foundation operates on a fundamentally different principle: **pay-to-play influence stratification.**

Under this model, corporate members purchase board seats and committee privileges rather than earning them through code contributions or community standing. The projects themselves are technically "open"—the code can be forked, the specifications can be read, anyone can submit patches. But the strategic direction, the roadmap priorities, the technical committee compositions—these are shaped by whoever is writing the largest checks.

This model has worked reasonably well for projects like Kubernetes, where multiple hyperscalers have aligned incentives in maintaining a healthy ecosystem. It's worked less well for projects where founding companies have divergent commercial interests from the broader community.

The Agentic AI Foundation is the latter scenario with steroids.

How Influence Actually Flows

Let me walk you through the mechanics:

- **Governing Board:** The eight platinum members each get a seat. This board sets foundation strategy, budget allocation, working group charters, and approval for new technical initiatives. It's not a rubber stamp—it's the actual steering wheel.
- **Technical Advisory Committees:** While technically open to community participation, these committees are invariably staffed and led by engineers from founding companies. They wrote the original code. They understand the architectures. They have the institutional knowledge. Even with the best intentions, this creates inherent structural advantage.
- **Working Groups:** Day-to-day specification work happens in working groups. Participation is open, but agenda-setting power flows from above. If the Governing Board decides MCP should prioritize enterprise authentication over peer-to-peer agent coordination, guess which working groups get funded and staffed?
- **Specification Approval:** Final sign-off on specification changes typically routes through technical committees that report to the board. Community input is solicited, but ultimate authority rests with the governance structure—which is controlled by the platinum members.

The genius of this model is that it's technically open at every level while being structurally closed where it matters most.

The Regulatory Capture Playbook

Let's call this what it is: a potential regulatory capture of open source AI infrastructure.

Regulatory capture traditionally refers to the phenomenon where regulatory agencies become dominated by the industries they're supposed to oversee. The regulators start acting in the interest of the regulated rather than the public interest.

What OpenAI, Anthropic, and their cohort have constructed is the open source equivalent. They've created a governance body that *appears* to represent community interests while structurally ensuring their own interests remain paramount.

Consider the incentive alignment:

For the Founding Companies

1. **Legitimacy without loss of control:** Moving protocols to a foundation provides the legitimacy of "open governance" while board seats ensure they maintain effective control over technical direction.
2. **Competitive moat via standardization:** Once their protocols become the official standard, every competitor must build to their specifications. The playing field tilts toward whoever understands those specifications best—the original authors.
3. **Regulatory defense:** When regulators come calling about AI concentration, these companies can point to their "open source" contributions and "neutral" foundation participation. See? We're good actors!
4. **Ecosystem lock-in disguised as interoperability:** Standards that technically enable interoperability but practically require their platforms for optimal implementation. MCP might be open, but whose infrastructure runs it best?

For Everyone Else

1. **Theoretical participation:** You can submit proposals, comment on specifications, contribute patches. Your voice is heard.
2. **Structural powerlessness:** But without a board seat ($350,000), without engineers embedded in technical committees, without the institutional knowledge of the original codebase, your influence is marginal at best.
3. **Dependency without recourse:** Build your agentic AI platform on these standards, and you're building on foundations you don't control. If the specification evolves in ways that disadvantage your use case, your options

are: adapt, fork (and lose compatibility), or abandon ship.

The "Open Source AI" Deception

This moment with the Agentic AI Foundation crystallizes a broader problem in the AI industry: the systematic appropriation of "open source" terminology to describe arrangements that share almost nothing with traditional open source values.

The reality is that many models labeled 'open source AI' include restrictive licenses that limit commercial use, scale, or industries. Meta's LLaMA 2, frequently cited as an example of open source AI, bars companies with over 700 million monthly active users from using it commercially without separate licensing. That's not open source. That's marketing.

As industry analysts have warned, the gap between "open source" as a marketing term and open source as a licensing philosophy has become a chasm. And the Agentic AI Foundation sits right in the middle of that chasm, using the language of openness while implementing a governance structure that concentrates power among its wealthiest members.

The Compute Reality

There's another dimension to this that the foundation structure obscures: you can have the most open protocols in the world, but if you can't run the models, you can't participate meaningfully in the ecosystem.

Critics have noted that concentration of compute, models, and data access allows large vendors to shape informal standards and de facto governance through APIs, safety SDKs, and proprietary toolchains. The protocols might be open. The practical ability to implement them at scale? That remains firmly concentrated.

Independent open source AI projects face competitive pressure from proprietary models' scale and hosted agent platforms, with higher barriers to match enterprise convenience. A two-person startup can read the MCP specification all day long. Without the compute to run sophisticated multi-agent systems, without the training data to build competitive models, without the enterprise relationships to deploy at scale, that specification knowledge is largely theoretical.

The Agentic AI Foundation addresses none of these structural advantages. It may, in fact, entrench them by creating a legitimacy wrapper around protocols optimized for hyperscaler deployment patterns.

The Security Dimension Nobody's Discussing

Agentic AI systems present security challenges fundamentally different from traditional software. We're talking about autonomous systems that can take actions, maintain state across sessions, and coordinate with other agents. The attack surface is enormous:

- **Memory poisoning:** Corrupting the contextual state that agents use to make decisions
- **Tool misuse:** Manipulating agents into using their capabilities in unintended ways
- **Privilege compromise:** Escalating agent permissions beyond intended boundaries
- **Non-deterministic behavior:** The inherent unpredictability of AI systems makes security guarantees nearly impossible

These security challenges require robust governance. But robust governance by whom?

Under the Agentic AI Foundation structure, security standards will be set by technical committees staffed primarily by engineers from the platinum members. Those members also happen to operate the largest deployed AI systems in the world. Their security priorities may not align with smaller players. Their threat models may assume resources and capabilities that others lack.

A security framework designed by AWS and Microsoft engineers, approved by an OpenAI and Anthropic-controlled board, will inevitably reflect the security assumptions and capabilities of AWS, Microsoft, OpenAI, and Anthropic. That might be fine for enterprises deploying on those platforms. It might be catastrophic for everyone else.

The Gartner Warning

Here's a statistic that should give everyone pause: Gartner predicts that 40%+ of agentic AI projects will be cancelled due to technical immaturity and governance challenges.

That’s not a prediction about technology failing. That’s a prediction about governance failing. And the Agentic AI Foundation is being positioned as the solution to that governance challenge.

But consider: if the governance solution is controlled by the same companies selling the models, the compute, and the deployment platforms, what happens when “governance” decisions happen to benefit those companies’ commercial interests?

When a startup’s agentic AI project fails because the security requirements proved too expensive to implement, is that a technical failure or a governance failure? When a enterprise cancels their multi-agent initiative because the protocol evolution went in a direction that required re-architecting their entire system, is that immaturity or capture?

What Genuine Open Governance Would Look Like

I want to be clear: I’m not arguing that the Agentic AI Foundation is illegitimate or that its founding members have malicious intent. The people involved are, by all accounts, genuinely committed to advancing agentic AI capabilities and believe open standards will accelerate the field.

But good intentions don’t fix structural problems. And the structure here is problematic.

Genuine open governance for these critical protocols would look different:

Board Composition

- **Independent seats:** At minimum, half of governing board seats should be independent of commercial AI interests. Academia. Civil society. Public interest technologists. Voices whose incentives don’t align with maximizing revenue from AI platforms.
- **Use-based representation:** Seats reserved for representatives of communities actually building on these protocols, not just companies selling into those communities.
- **Rotating leadership:** Board chair and key committee positions should rotate among member classes, not remain perpetually with platinum members.

Technical Committee Reform

- **Merit-based participation:** Technical committee membership based on demonstrated expertise and contribution, not employment by a founding company.
- **Conflict disclosure:** Mandatory disclosure when committee members’ employers have commercial interests in specification decisions.
- **Delayed implementation:** Founding companies prohibited from shipping specification features for a fixed period (90 days? 180 days?) after approval, giving the broader community time to implement.

Economic Model

- **Sliding scale membership:** Membership tiers based on company revenue or AI market position, not flat fees that price out smaller players.
- **Community sponsorship:** Pool of funded seats for independent developers and small organizations to participate meaningfully.
- **Alternative influence mechanisms:** Non-financial pathways to board-level influence based on code contribution, implementation success, or community leadership.

None of this is radical. These are governance patterns that exist in other open source foundations. They were deliberately not chosen here.

The Path Forward: What You Can Do

> If you’re building on these protocols, you need to understand that you’re building on political infrastructure, not just technical infrastructure.

Here’s my practical guidance:

For Enterprises Evaluating Agentic AI

1. **Assess governance risk as technical risk:** Your vendor evaluation should include analysis of foundation governance. Who controls the standards your system depends on? What are their commercial incentives?
2. **Demand abstraction layers:** Build with abstraction layers that could,

theoretically, swap underlying protocol implementations. Yes, this adds complexity. It also adds optionality if governance evolves unfavorably.

3. **Participate actively:** If you're building significant systems on these protocols, participate in the foundation. Show up to working group meetings. Submit comments on specifications. Make your voice heard, even if your vote doesn't count.

For Developers and Startups

1. **Document dependencies explicitly:** Know exactly where your stack touches foundation-controlled specifications. Map those dependencies. Understand your exposure.
2. **Watch the working groups:** Specification changes often provide months of warning through working group discussions. Monitor those discussions for changes that could affect your systems.
3. **Build coalitions:** Individual developers have no voice in this structure. Developer coalitions might. Find others with similar concerns. Organize.

For Regulators and Policymakers

1. **See through the "open source" narrative:** Open code and open governance are not the same thing. Protocols can be technically open while being structurally captured. Evaluate AI infrastructure concentration based on governance reality, not marketing claims.
2. **Consider governance mandates:** If these protocols become critical infrastructure (they will), governance structure becomes a matter of public interest. Foundation governance could be subject to the same scrutiny as critical infrastructure governance in other sectors.
3. **Fund alternatives:** Public investment in genuinely independent AI protocol development could provide competitive pressure on captured foundations. The EU has done this in other technology areas. AI infrastructure may warrant similar attention.

The Bigger Picture

What's happening with the Agentic AI Foundation is not unique to AI. It's the latest iteration of a pattern we've seen repeatedly in technology: the infrastructure layer consolidates first, then the value chain above it consolidates around whoever controls that infrastructure.

IBM dominated computing when they controlled the hardware architecture. Microsoft dominated the PC era when they controlled the operating system. Google dominated the web era when they controlled the search index. AWS dominates cloud computing because they control the deployment infrastructure.

Era	Infrastructure Layer	Dominant Controller	Consolidation Mechanism
Mainframe	Hardware Architecture	IBM	Proprietary Standards
PC	Operating System	Microsoft	OEM Licensing
Web	Search Index	Google	Data Accumulation
Cloud	Deployment Infrastructure	AWS	Service Integration
Agentic AI	Communication Protocols	???	Open Source Capture?

The agentic AI era is next. And the infrastructure layer—how agents communicate, coordinate, and interoperate—is being set right now. The companies positioning themselves to control that layer are doing so not through proprietary lock-in (which draws regulatory scrutiny) but through a more sophisticated mechanism: governance capture of nominally open standards.

It's elegant. It's legal. It's probably even well-intentioned, at least partially.

And it should worry anyone who cares about the long-term structure of the AI ecosystem.

Conclusion: The Question That Matters

I started this piece with a question: When the companies donating the protocols are the same companies paying for board seats that control those protocols, what exactly got donated?

After analyzing the governance structure, the incentive alignment, and the historical patterns, I have an answer: What got donated was the obligation to maintain the protocols alone.

The founding companies transferred maintenance burden to a shared foundation while retaining effective control through governance mechanisms. They gained legitimacy, regulatory cover, and ecosystem lock-in potential. They gave up… the costs of going it alone.

That's not necessarily malicious. Shared maintenance of critical infrastructure is genuinely valuable. Open specifications that enable interoperability are genuinely good.

But let's not pretend this is liberation. Let's not pretend this is democratization. Let's not pretend that $350,000 board seats and founding-company-staffed technical committees represent open governance in any meaningful sense.

The Agentic AI Foundation may deliver genuine technical value. The protocols it governs may enable real innovation. The interoperability it promises may materialize.

But the governance structure ensures that the value, the innovation, and the interoperability will flow through channels controlled by a handful of companies who already dominate the AI landscape. That's not a bug in the foundation design. It's the feature.

The protocols are open. The governance is not. And in the long run, governance always wins.

The most dangerous lock-in isn't technical—it's governance capture disguised as open source contribution, and the Agentic AI Foundation represents the most sophisticated implementation of this strategy the tech industry has ever seen.