

Decentralized AI Agent and Model Federation Without a Central Registry: Adaptive Indexing for Cross-Organization Discovery and Addressing

Organizations want to discover, address, and federate AI models and agents across trust boundaries, but every available approach forces a central registry, a shared ledger, or a single naming authority that no participant fully controls. This article describes how that discovery-and-addressing problem is solved by Adaptive Indexing, disclosed in United States Patent Application 19/326,036. The framework supplies a decentralized, anchor-governed namespace in which each organization governs its own scope, model and agent identifiers stay stable across moves and version changes, and federation happens through scoped delegation rather than network-wide coordination.

What This Application Specifies

Adaptive Indexing, disclosed in United States Patent Application 19/326,036, is a distributed indexing and resolution architecture that operates across heterogeneous, decentralized environments without requiring global consensus, static directories, or centralized trust anchors. Its core data structure is an adaptive index: a set of entries organized in a parent-child hierarchy, where each entry corresponds to a unique semantic scope identified by a structured alias, and each scope is governed by one or more anchors.

Anchors are the governance units of the framework. They perform two primary roles within a defined scope: they cache content and they execute scoped voting procedures for structural mutation. An anchor encodes mutation policy, alias mapping, and access control metadata. Crucially, anchors are not data hosts. They maintain index metadata, permissions, and lineage references, while actual storage and delivery are performed by participating nodes. A single node may host multiple anchors, and any given anchor may reference a set of nearby or trusted nodes capable of fulfilling fetches.

The framework names AI federation directly. The specification describes peer-to-peer AI networks that share models or support federated training, where a versioned checkpoint resolves via an alias such as a model path so that anchors manage routing and replication dynamically based on topic, location, or demand. It also routes requests originating from semantic agents and instantiates caches according to real-time telemetry and anchor policy. Applied to cross-organization AI federation, an organization registers its models and agents as assets within anchor-governed containers, and other organizations resolve and address them through scoped delegation rather than through a directory any single party owns.

Why It Matters

Federating AI across organizational boundaries runs into a naming and governance problem before it ever reaches an inference call. Existing decentralized infrastructures suffer from structural rigidity due to globally replicated state, monolithic consensus enforcement, and externalized mutation control, and current indexing mechanisms such as DNS, distributed content addressing, and contract-based registries rely on static alias mappings, centralized delegation hierarchies, or cryptographic immutability. None of these fit a setting where independent organizations each want to publish, version, retire, and govern their own models and agents on their own terms while still being discoverable by partners.

A central model registry forces every participant to trust an operator who can rename, deprioritize, or revoke entries. A shared ledger forces global agreement on every catalog change and makes high-frequency model versioning expensive. A single naming authority becomes a chokepoint and a single point of failure. Adaptive Indexing matters here because it permits organic, trust-divergent growth and reorganization over time: independently governed subtrees can evolve, delegate authority, or rekey entries without requiring network-wide coordination, and anchors coordinate mutation decisions only within their jurisdictional boundaries. Each organization keeps sovereignty over its own scope while remaining resolvable across the federation.

How It Composes With the Domain

Each participating organization owns one or more semantic scopes in the adaptive index, identified by structured aliases of the form that the specification defines, namely a top-level domain, domain and subdomain, nested subindices, and an asset. A model or agent becomes an addressable asset under that scope. The specification describes model aliases that resolve down a hierarchy such as a vendor, a model family, and a versioned checkpoint, and it describes resolving requests that originate from semantic agents. Federation across organizations follows the same nested resolution path: an alias is interpreted stepwise, each segment relative to its parent scope, enabling traceable, recursive delegation of authority. If an alias is not found locally, it can be escalated to trusted peers or delegated upward through the nesting structure, which the specification identifies as the mechanism that allows seamless federation between systems or organizations.

Stable identity is what makes addressing durable. Each alias resolves to a unique identifier that remains stable even as the alias is renamed, delegated, or restructured. A model can move between index paths, change alias bindings, or migrate across physical infrastructure while its underlying identifier stays fixed, and anchors store metadata, permissions, and version pointers keyed to that identifier. Versioning is first-class: a mutation to an asset creates a new version entry under its identifier, with prior versions

retained for audit, rollback, or comparison. For an agent or model that ships new checkpoints frequently, this means a partner organization can pin a stable reference and still receive continuity across releases.

Governance is local and policy-driven. When an organization restructures its catalog, for example splitting an overloaded model family into child subindices or merging dormant ones, the change is validated by the anchor group governing that scope through quorum voting under a pre-registered policy. Quorum thresholds are policy-defined and can vary by sensitivity, so routine catalog updates may use a low quorum while a policy rekey may require full anchor participation. Propagation beyond a zone boundary requires an elevated quorum validation, ensuring that inter-organization changes occur only under explicit policy authorization. Each approved mutation appends a cryptographically committed lineage record capturing the prior anchor map, the justification, and the quorum configuration, so federation participants get a verifiable audit trail of how a shared model namespace evolved.

Discovery and delivery are decentralized end to end. Anchors maintain an index of nodes serving a given asset, annotated with proximity, latency, load, and a trust score derived from performance history and policy compliance. A requesting agent resolves an alias to a stable identifier and a set of candidate host nodes, and the routing layer selects the most performant and reliable path, rerouting automatically when a node degrades. Caching is proximity-aware and instantiated on demand: when a model or agent endpoint sees repeated access from a region, eligible nodes instantiate local replicas and register with the responsible anchor group. Anchors may also integrate with federated identity providers using token adapters that translate externally issued credentials into scoped, temporary access rights governed by anchor-local policy, which is the hook for letting one organization's identity system authorize access to another's models.

What This Enables

A cross-organization federation built on this framework can offer a single resolvable namespace for models and agents without a registry operator. Each organization governs its own scope, publishes and retires entries autonomously, and remains discoverable through scoped delegation. Partners can address a counterpart's agent by alias and receive resolution to a stable identifier even after the publisher reorganizes, renames, or migrates the underlying model.

Because mutation governance is scoped and quorum-based, organizations can co-own shared scopes, for example a consortium catalog of jointly governed agents, with quorum thresholds set to the consortium's trust model and elevated quorum required for changes that cross into another member's zone. Access control evaluates dynamically at resolution time based on identity, device context, request provenance, and trust score, so a model can be exposed to one partner under edit rights and another under read-only, with time-bound policies that grant access to a checkpoint only during an approved window and then expire without administrative intervention.

The framework also tolerates the messy reality of federation. Anchors may accept mutation proposals asynchronously and validate them later upon quorum availability, and anchor groups can temporarily form isolated quorums and maintain index responsiveness in disconnected or high-latency conditions, reconciling on reconnection. For privacy-sensitive federation, hybrid consensus modes can integrate cryptographic attestations such as zero-knowledge proofs, letting anchors validate claims without accessing private content. And because the framework is designed as a structural overlay, the specification states that peer-to-peer AI networks sharing models or supporting federated training can be retrofitted with scalable, trust-scoped resolution without altering their core protocols, so an existing federation effort can adopt the namespace incrementally.

Boundary Conditions

The framework supplies discovery, addressing, governance, and routing for models and agents. It does not specify the inference runtime, the model architecture, the training procedure, or the agent reasoning logic. A model is treated as an addressable, versioned asset and an agent as an originator and target of resolution. The value of the framework is in how those assets are named, governed, found, and delivered, not in what they compute.

Resolution and governance are only as trustworthy as the policies and anchor groups that organizations define. Quorum thresholds, trust-weighting functions, and admission criteria are policy choices; a poorly configured policy or a compromised quorum within a scope can still approve unwanted structural mutations within that scope, which is why the framework scopes authority tightly and requires elevated quorum to cross zone boundaries. The specification provides legacy fallback, where an alias that fails to resolve within the network may fall back to a corresponding legacy domain, but full benefit assumes participants actually adopt anchors and aliases. The specification describes mechanisms and architecture; it does not assert performance benchmarks, and no throughput, latency, or scale figures should be inferred from this article. Real deployments will be shaped by node availability, network conditions, and the trust relationships among the federating organizations.

Disclosure Scope

The technical capabilities described here, namely the adaptive index, anchor-governed scopes, structured aliases resolving to stable unique identifiers, quorum-based mutation governance with lineage continuity, trust-weighted and proximity-aware routing, on-demand caching, and federated-identity token adapters, are disclosed in United States Patent Application 19/326,036, including its express description of peer-to-peer AI networks that share models or support federated training and of requests originating from semantic agents. The decentralized AI agent and model federation

domain framing, including consortium catalogs, cross-organization governance models, partner access patterns, and references to general categories of registries, ledgers, and naming systems, is external application context provided to illustrate an enabling implementation; it is not part of the disclosed invention and does not expand the scope of the application. Any reference to standards, regulatory regimes, or domain practices is included only as faithful external context.

Adaptive Indexing ([/adaptive-indexing](#))

[All 40 steps](#) → ([/inventive-steps](#)).

Resolution without global consensus. Anchor-governed self-organization.

[U.S. 19/326,036](#) ([/patents/19-326036](#))

PRIMARY TECHNICAL DISCLOSURE

- [The Adaptive Index: A Scalable Foundation for Decentralized Systems](#) ([/articles/the-adaptive-index-a-scalable-foundation-for-decentralized-systems](#)).

SECONDARY TECHNICAL

- [Anchor-Governed Hierarchical Nesting: Recursive Semantic Containers at Unlimited Depth](#) ([/articles/adaptive-indexing/anchor-nesting](#)).
- [Entropy-Triggered Index Splitting: Deterministic Partitioning Under Mutation Load](#) ([/articles/adaptive-indexing/entropy-splitting](#)).
- [Dormant Index Merging: Recursive Consolidation of Low-Entropy Subindices](#) ([/articles/adaptive-indexing/dormant-merging](#)).
- [Elastic Anchor Group Management: Governance That Scales With Criticality](#) ([/articles/adaptive-indexing/elastic-anchors](#)).
- [Trust-Weighted Quorum Voting: Consensus Where Weight Reflects Earned Trust](#) ([/articles/adaptive-indexing/trust-weighted-voting](#)).
- [Asynchronous Consensus Coordination: Offline Vote Completion With Reconciliation](#) ([/articles/adaptive-indexing/async-consensus](#)).
- [Best-Match Alias Querying: Longest-Match Resolution With Stepwise Delegation](#) ([/articles/adaptive-indexing/best-match-aliases](#)).

- [Action-Typed Aliases: Behavioral Intent Embedded in the Namespace \(/articles/adaptive-indexing/action-typed-aliases\)](/articles/adaptive-indexing/action-typed-aliases).
- [UID Persistence Through Alias Mutation: Stable Identity Across Structural Change \(/articles/adaptive-indexing/uid-persistence\)](/articles/adaptive-indexing/uid-persistence).
- [Lineage-Preserving Structural Mutation: Cryptographic History Through Every Change \(/articles/adaptive-indexing/lineage-preserving-mutation\)](/articles/adaptive-indexing/lineage-preserving-mutation).
- [Proximity-Based Routing With Trust Scoring: Dynamic Path Selection in Decentralized Networks \(/articles/adaptive-indexing/proximity-routing\)](/articles/adaptive-indexing/proximity-routing)
- [Dynamic Device Hash for Pseudonymous Authentication: Volatile Identity Without Stored Credentials \(/articles/adaptive-indexing/device-hash-auth\)](/articles/adaptive-indexing/device-hash-auth).
- [On-Demand Adaptive Caching: Cache Instances That Follow Usage, Not Configuration \(/articles/adaptive-indexing/adaptive-caching\)](/articles/adaptive-indexing/adaptive-caching)
- [Predictive Cache Prefetching: Forecasting Models That Proactively Instantiate Caches \(/articles/adaptive-indexing/predictive-prefetching\)](/articles/adaptive-indexing/predictive-prefetching).
- [Contextual Access Enforcement: Policy Graphs Evaluated With Real-Time Telemetry \(/articles/adaptive-indexing/contextual-access\)](/articles/adaptive-indexing/contextual-access).
- [Mutation Router With Contextual Signals: Policy-Aware Propagation Path Selection \(/articles/adaptive-indexing/mutation-routing\)](/articles/adaptive-indexing/mutation-routing)
- [Impact Simulation During Mutation Staging: Pre-Execution Analysis of Proposed Changes \(/articles/adaptive-indexing/impact-simulation\)](/articles/adaptive-indexing/impact-simulation).
- [DNS Bidirectional Fallback: Hybrid Resolution With Legacy DNS Compatibility \(/articles/adaptive-indexing/dns-fallback\)](/articles/adaptive-indexing/dns-fallback)
- [Asset Versioning as First-Class Metadata: Version Entries Under UIDs With Lineage Tracking \(/articles/adaptive-indexing/asset-versioning\)](/articles/adaptive-indexing/asset-versioning).
- [Telemetry-Driven Topology Mutation: Autonomous Network Reconfiguration From Operational Data \(/articles/adaptive-indexing/telemetry-topology\)](/articles/adaptive-indexing/telemetry-topology).
- [The Index Is the Territory: The Navigable Substrate Beneath Both Axes \(/articles/adaptive-indexing/the-index-is-the-territory\)](/articles/adaptive-indexing/the-index-is-the-territory)

APPLICATIONS · GENERAL

- [Decentralized AI Agent and Model Federation Without a Central Registry: Adaptive Indexing for Cross-Organization Discovery and Addressing \(/articles/adaptive-indexing/decentralized-ai-federation\)](/articles/adaptive-indexing/decentralized-ai-federation).
- [Payload-Aware Edge Caching and Live Retransmission: Replacing Address-Based CDN Heuristics With Adaptive Indexing \(/articles/adaptive-indexing/cdn-and-live-media\)](/articles/adaptive-indexing/cdn-and-live-media)
- [How to Retrofit Adaptive Indexing onto Legacy Decentralized Systems \(Web3, Fediverse, DAOs\) \(/articles/adaptive-indexing/applying-to-legacy-systems\)](/articles/adaptive-indexing/applying-to-legacy-systems).

- [Why Edge Platforms Still Depend on a Central Authority \(/articles/adaptive-indexing/why-edge-platforms-depend-on-central-authority\)](/articles/adaptive-indexing/why-edge-platforms-depend-on-central-authority).
- [Supply Chain Tracking Through Governed Namespace Resolution \(/articles/adaptive-indexing/supply-chain-provenance\)](/articles/adaptive-indexing/supply-chain-provenance).
- [Social Media Platforms Without Central Namespace Authority \(/articles/adaptive-indexing/decentralized-social\)](/articles/adaptive-indexing/decentralized-social).
- [Healthcare Data Federation Through Scoped Governance \(/articles/adaptive-indexing/healthcare-data-federation\)](/articles/adaptive-indexing/healthcare-data-federation).
- [Sovereign Government Digital Identity Without a Central Registry \(/articles/adaptive-indexing/government-identity-infrastructure\)](/articles/adaptive-indexing/government-identity-infrastructure).
- [Governed Securities Identifier Resolution for Financial Market Data \(/articles/adaptive-indexing/financial-market-data\)](/articles/adaptive-indexing/financial-market-data).
- [Cross-Platform Gaming and Metaverse Namespace Governance for Portable Player Identity and Assets \(/articles/adaptive-indexing/gaming-metaverse-namespace\)](/articles/adaptive-indexing/gaming-metaverse-namespace).

APPLICATIONS · SPECIFIC

- [Cloudflare's Edge Has a Namespace Problem \(/articles/adaptive-indexing/cloudflare\)](/articles/adaptive-indexing/cloudflare).
- [DNS Is 40 Years Old and Still Running the Internet \(/articles/adaptive-indexing/dns\)](/articles/adaptive-indexing/dns).
- [ENS Solved the Wrong Half of the Naming Problem \(/articles/adaptive-indexing/ens\)](/articles/adaptive-indexing/ens).
- [Handshake Decentralized the Root. Everything Below It Is Still Ungoverned. \(/articles/adaptive-indexing/handshake\)](/articles/adaptive-indexing/handshake).
- [IPFS Solved Content Addressing. It Didn't Solve Naming, Persistence, or Governance. \(/articles/adaptive-indexing/ipfs\)](/articles/adaptive-indexing/ipfs).
- [Fastly Built the Fastest Cache Invalidation in the Industry. The Authority to Invalidate Still Lives in One Place. \(/articles/adaptive-indexing/fastly\)](/articles/adaptive-indexing/fastly).
- [Akamai Built the Internet's Delivery Infrastructure. It Was Designed for a World That Needed Central Control. \(/articles/adaptive-indexing/akamai\)](/articles/adaptive-indexing/akamai).
- [Bluesky Identified the Right Problem. The Architecture That Solves It Is the Adaptive Index. \(/articles/adaptive-indexing/bluesky\)](/articles/adaptive-indexing/bluesky).
- [Consul's Service Catalog Is Brilliant Infrastructure. It Is Still a Central Registry. \(/articles/adaptive-indexing/consul\)](/articles/adaptive-indexing/consul).
- [Istio Solved Programmable Traffic Policy. The Namespace That Routes Traffic Is Still Central. \(/articles/adaptive-indexing/istio\)](/articles/adaptive-indexing/istio).
- [Unstoppable Domains Proved NFT Ownership Works. The Namespace Governance Model Is Still Unresolved. \(/articles/adaptive-indexing/unstoppable-domains\)](/articles/adaptive-indexing/unstoppable-domains).
- [The Graph Built the Index Layer for Web3. The Index Itself Still Has a Governance Problem. \(/articles/adaptive-indexing/the-graph\)](/articles/adaptive-indexing/the-graph).

- [Filecoin Proved Verifiable Storage. Discovery and Namespace Governance Are Still Unsolved. \(/articles/adaptive-indexing/filecoin\)](#).
- [Arweave Made Data Permanent. It Has No Governance Model for What Permanent Data Means Over Time. \(/articles/adaptive-indexing/arweave\)](#).
- [Ceramic Built Mutable Data Streams for Web3. The Governance of Those Streams Is Still Not Local. \(/articles/adaptive-indexing/ceramic\)](#).
- [Kubernetes Service Discovery Resolves Within Clusters. Cross-Cluster Namespace Is Central. \(/articles/adaptive-indexing/kubernetes\)](#).
- [Amazon Route 53 Is the Most Reliable DNS on Earth. It Is Still DNS Architecture. \(/articles/adaptive-indexing/amazon-route53\)](#).
- [HashiCorp Nomad Distributes Scheduling. The Namespace That Organizes It Is Still Central. \(/articles/adaptive-indexing/hashicorp-nomad\)](#).
- [ZooKeeper Coordinates Distributed Systems. The Coordinator Is a Single Point of Authority. \(/articles/adaptive-indexing/zookeeper\)](#).
- [etcd Stores the State of Kubernetes. The State Store Has No Scoped Governance. \(/articles/adaptive-indexing/etcd\)](#).
- [Consul KV Distributes Configuration. The Distribution Authority Is Still Central. \(/articles/adaptive-indexing/consul-kv\)](#).
- [Raft Made Consensus Understandable. It Did Not Make Consensus Scope-Aware. \(/articles/adaptive-indexing/raft-protocol\)](#).
- [Paxos Proved Consensus Is Possible. It Did Not Address Namespace Governance. \(/articles/adaptive-indexing/paxos\)](#).
- [Cosmos Tendermint Enabled Sovereign Blockchains. The Namespace Between Them Is Ungoverned. \(/articles/adaptive-indexing/cosmos-tendermint\)](#).
- [AWS Cloud Map Discovers Services. The Discovery Authority Lives in One Region's Control Plane. \(/articles/adaptive-indexing/aws-service-discovery\)](#).
- [Azure Traffic Manager Routes Globally. The Routing Authority Is Centrally Defined. \(/articles/adaptive-indexing/azure-traffic-manager\)](#).
- [GCP Service Directory Centralizes Service Registration. Registration Is Not Governance. \(/articles/adaptive-indexing/gcp-service-directory\)](#).
- [Netlify DNS Simplifies Deployment Routing. The Namespace Authority Is Still Netlify's. \(/articles/adaptive-indexing/netlify-dns\)](#).
- [Vercel's Edge Network Executes at the Boundary. Routing Authority Does Not. \(/articles/adaptive-indexing/vercel-edge\)](#).
- [Bunny CDN Delivers Content Globally. Cache Governance Is Still Central. \(/articles/adaptive-indexing/bunny-cdn\)](#).

- [KeyCDN Optimized Content Delivery. The Delivery Namespace Is Centrally Controlled. \(/articles/adaptive-indexing/keycdn\)](/articles/adaptive-indexing/keycdn).
- [Limelight Networks Built Private Infrastructure for Delivery. The Namespace Governance Is Still Central. \(/articles/adaptive-indexing/limelight\)](/articles/adaptive-indexing/limelight).
- [StackPath Combined CDN With Edge Computing. Namespace Authority Remained Central. \(/articles/adaptive-indexing/stackpath\)](/articles/adaptive-indexing/stackpath).
- [Envoy Proxy Made Service Mesh Data Planes Programmable. The Control Plane Still Governs. \(/articles/adaptive-indexing/envoy-proxy\)](/articles/adaptive-indexing/envoy-proxy)
- [NGINX Powers the Web's Reverse Proxy Layer. Its Configuration Is Statically Defined. \(/articles/adaptive-indexing/nginx\)](/articles/adaptive-indexing/nginx).
- [Traefik Discovers Services Automatically. The Discovery Namespace Is Still External. \(/articles/adaptive-indexing/traefik\)](/articles/adaptive-indexing/traefik)
- [Linkerd Simplified the Service Mesh. The Namespace It Meshes Is Still Kubernetes. \(/articles/adaptive-indexing/linkerd\)](/articles/adaptive-indexing/linkerd).
- [Namecheap Made Domain Registration Accessible. Domain Governance Remains the Registrar Model. \(/articles/adaptive-indexing/namecheap\)](/articles/adaptive-indexing/namecheap).
- [GoDaddy Registered More Domains Than Anyone. The Namespace Model Has Not Changed. \(/articles/adaptive-indexing/godaddy\)](/articles/adaptive-indexing/godaddy)
- [DNSimple Made DNS Management Developer-Friendly. The Governance Model Is Still DNS. \(/articles/adaptive-indexing/dnssimple\)](/articles/adaptive-indexing/dnssimple).
- [Datadog Observes Everything. The Namespace It Observes Has No Governed Structure. \(/articles/adaptive-indexing/datadog\)](/articles/adaptive-indexing/datadog)
- [Grafana Unified Observability Visualization. The Data Namespace It Queries Has No Governed Structure. \(/articles/adaptive-indexing/grafana\)](/articles/adaptive-indexing/grafana).
- [Prometheus Defined Cloud-Native Monitoring. Its Metric Namespace Has No Governance Layer. \(/articles/adaptive-indexing/prometheus\)](/articles/adaptive-indexing/prometheus).
- [New Relic Pioneered APM. The Telemetry Namespace It Built Is Centrally Indexed. \(/articles/adaptive-indexing/new-relic\)](/articles/adaptive-indexing/new-relic)
- [Splunk Indexes Machine Data at Scale. The Index Namespace Is Centrally Administered. \(/articles/adaptive-indexing/splunk\)](/articles/adaptive-indexing/splunk).
- [GitHub Copilot and Copilot Workspace \(/articles/adaptive-indexing/github-copilot-workspace\)](/articles/adaptive-indexing/github-copilot-workspace).
- [Tableau Pulse and Salesforce Tableau AI \(/articles/adaptive-indexing/tableau-pulse\)](/articles/adaptive-indexing/tableau-pulse).
- [Notion AI Knowledge Platform \(/articles/adaptive-indexing/notion-ai\)](/articles/adaptive-indexing/notion-ai).

[Adaptive Indexing overview → \(/adaptive-indexing\)](/adaptive-indexing)

