



[Home](#) [Licensing](#) [Patents](#) [Articles](#)

Proximity-Based Routing With Trust Scoring: Dynamic Path Selection in Decentralized Networks

by [Nick Clark](#) | Published March 27, 2026 | [PDF](#)

The adaptive index does not route through fixed paths. When multiple anchors can serve a resolution or mutation request, the system selects among them based on a composite of network proximity and trust score. Closer anchors with higher trust are preferred, but a nearby anchor with low trust may be bypassed in favor of a more distant anchor with better governance credentials. The result is routing that optimizes for both latency and safety simultaneously.

What It Is

Proximity-based routing with trust scoring is the path selection mechanism used when a resolution or mutation request can be served by more than one anchor. The routing function computes a composite score from two inputs: the network distance to each candidate anchor and the trust score each anchor has earned within the relevant governance scope. The anchor with the highest composite score handles the request.

Network proximity may be measured by latency, hop count, or geographic distance depending on the deployment context. Trust scores are maintained per-scope as described in the trust-weighted voting mechanism. The composite function and its weighting are configurable per scope, allowing different scopes to prioritize latency or trust differently.

Why It Matters

Pure proximity routing sends requests to the nearest available node regardless of that node's trustworthiness. This is efficient but unsafe: a compromised node at a network edge can intercept and manipulate traffic from nearby clients. Pure trust routing sends requests to the most trusted node regardless of distance, sacrificing latency for safety.

Combining proximity and trust produces routing that is both efficient and safe. The system naturally avoids untrusted nearby nodes and unnecessarily distant trusted nodes, converging on paths that satisfy both operational and governance requirements.

How It Works Structurally

When a resolution request enters a scope, the requesting node evaluates available anchors using a scoring function that combines measured proximity with the anchor's trust score. The scoring function is defined in the scope's routing policy and may weight proximity and trust differently for different action types: read operations might favor proximity, while mutation operations might favor trust.

The routing decision is re-evaluated periodically as proximity measurements update and trust scores evolve. A path that was optimal yesterday may not be optimal today if an anchor's trust has degraded or if network conditions have changed. This dynamic re-evaluation ensures routing remains optimal over time without manual reconfiguration.

What It Enables

Proximity-trust routing enables the adaptive index to operate safely in adversarial environments without sacrificing performance. Edge deployments route through the nearest trustworthy node rather than the nearest node. Multi-region deployments balance latency against governance quality automatically. Networks under partial compromise route around untrusted nodes without operator intervention.

This mechanism transforms routing from a pure infrastructure concern into a governance-aware decision, embedding safety into the data path itself.

[Adaptive Indexing All 21 steps →](#)

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

Patent

[US 19/326,036](#) · published

Primary Technical Disclosure

[◦ The Adaptive Index: A Scalable Foundation for Decentralized Systems](#)

Secondary Technical

[◦ Anchor-Governed Hierarchical Nesting: Recursive Semantic Containers at Unlimited Depth](#)◦ [Entropy-Triggered Index Splitting: Deterministic Partitioning Under Mutation Load](#)◦ [Dormant Index Merging: Recursive Consolidation of Low-Entropy Subindices](#)◦ [Elastic Anchor Group Management: Governance That Scales With Criticality](#)◦ [Trust-Weighted Quorum Voting: Consensus Where Weight Reflects Earned Trust](#)◦ [Asynchronous Consensus Coordination: Offline Vote Completion With Reconciliation](#)◦ [Best-Match Alias Querying: Longest-Match Resolution With Stepwise Delegation](#)◦ [Action-Typed Aliases: Behavioral Intent Embedded in the Namespace](#)◦ [UID Persistence Through Alias Mutation: Stable Identity Across Structural Change](#)◦ [Lineage-Preserving Structural Mutation: Cryptographic History Through Every Change](#)● [Proximity-Based Routing With Trust Scoring: Dynamic Path Selection in Decentralized Networks](#)◦ [Dynamic Device Hash for Pseudonymous Authentication: Volatile Identity Without Stored Credentials](#)◦ [On-Demand Adaptive Caching: Cache Instances That Follow Usage, Not Configuration](#)◦ [Predictive Cache Prefetching: Forecasting Models That Proactively Instantiate Caches](#)◦ [Contextual Access Enforcement: Policy Graphs Evaluated With Real-Time Telemetry](#)◦ [Mutation Router With Contextual Signals: Policy-Aware Propagation Path Selection](#)◦ [Impact Simulation During Mutation Staging: Pre-Execution Analysis of Proposed Changes](#)◦ [DNS Bidirectional Fallback: Hybrid Resolution With Legacy DNS Compatibility](#)◦ [Asset Versioning as First-Class Metadata: Version Entries Under UIDs With Lineage Tracking](#)◦ [Telemetry-Driven Topology Mutation: Autonomous Network Reconfiguration From Operational Data](#)

Applications (General)

[◦ Applying Adaptive Indexes to Legacy Decentralized Systems](#)◦ [Why Edge Platforms Still Depend on a Central Authority](#)◦ [Supply Chain Tracking Through Governed Namespace Resolution](#)◦ [Social Media Platforms Without Central Namespace Authority](#)◦ [Healthcare Data Federation Through Scoped Governance](#)◦ [Government Identity Infrastructure at Scale](#)◦ [Financial Market Data With Governed Resolution](#)◦ [Gaming and Metaverse Namespace Governance](#)

Applications (Specific)

[◦ Cloudflare's Edge Has a Namespace Problem](#)◦ [DNS Is 40 Years Old and Still Running the Internet](#)◦ [ENS Solved the Wrong Half of the Naming Problem](#)◦ [Handshake Decentralized the Root. Everything Below It Is Still Ungoverned.](#)◦ [IPFS Solved Content Addressing. It Didn't Solve Naming, Persistence, or Governance.](#)◦ [Fastly Built the Fastest Cache Invalidation in the Industry. The Authority to Invalidate Still Lives in One Place.](#)◦ [Akamai Built the Internet's Delivery Infrastructure. It Was Designed for a World That Needed Central Control.](#)◦ [Bluesky Identified the Right Problem. The Architecture That Solves It Is the Adaptive Index.](#)◦ [Consul's Service Catalog Is Brilliant Infrastructure. It Is Still a Central Registry.](#)◦ [Istio Solved Programmable Traffic Policy. The Namespace That Routes Traffic Is Still Central.](#)◦ [Unstoppable Domains Proved NFT Ownership Works. The Namespace Governance Model Is Still Unresolved.](#)◦ [The Graph Built the Index Layer for Web3. The Index Itself Still Has a Governance Problem.](#)◦ [Filecoin Proved Verifiable Storage. Discovery and Namespace Governance Are Still Unsolved.](#)◦ [Arweave Made Data Permanent. It Has No Governance Model for What Permanent Data Means Over Time.](#)◦ [Ceramic Built Mutable Data Streams for Web3. The Governance of Those Streams Is Still Not Local.](#)◦ [Kubernetes Service Discovery Resolves Within Clusters. Cross-Cluster Namespace Is Central.](#)◦ [Amazon Route 53 Is the Most Reliable DNS on Earth. It Is Still DNS Architecture.](#)◦ [HashiCorp Nomad Distributes Scheduling. The Namespace That Organizes It Is Still Central.](#)◦ [ZooKeeper Coordinates Distributed Systems. The Coordinator Is a Single Point of Authority.](#)◦ [etcd Stores the State of Kubernetes. The State Store Has No Scoped](#)

[Governance.](#)[Consul KV Distributes Configuration. The Distribution Authority Is Still Central.](#)[Raft Made Consensus Understandable. It Did Not Make Consensus Scope-Aware.](#)[Paxos Proved Consensus Is Possible. It Did Not Address Namespace Governance.](#)[Cosmos Tendermint Enabled Sovereign Blockchains. The Namespace Between Them Is Ungoverned.](#)[AWS Cloud Map Discovers Services. The Discovery Authority Lives in One Region's Control Plane.](#)[Azure Traffic Manager Routes Globally. The Routing Authority Is Centrally Defined.](#)[GCP Service Directory Centralizes Service Registration. Registration Is Not Governance.](#)[Netlify DNS Simplifies Deployment Routing. The Namespace Authority Is Still Netlify's.](#)[Vercel's Edge Network Executes at the Boundary. Routing Authority Does Not.](#)[Bunny CDN Delivers Content Globally. Cache Governance Is Still Central.](#)[KeyCDN Optimized Content Delivery. The Delivery Namespace Is Centrally Controlled.](#)[Limelight Networks Built Private Infrastructure for Delivery. The Namespace Governance Is Still Central.](#)[StackPath Combined CDN With Edge Computing. Namespace Authority Remained Central.](#)[Envoy Proxy Made Service Mesh Data Planes Programmable. The Control Plane Still Governs.](#)[NGINX Powers the Web's Reverse Proxy Layer. Its Configuration Is Statically Defined.](#)[Traefik Discovers Services Automatically. The Discovery Namespace Is Still External.](#)[Linkerd Simplified the Service Mesh. The Namespace It Meshes Is Still Kubernetes.](#)[Namecheap Made Domain Registration Accessible. Domain Governance Remains the Registrar Model.](#)[GoDaddy Registered More Domains Than Anyone. The Namespace Model Has Not Changed.](#)[DNSimple Made DNS Management Developer-Friendly. The Governance Model Is Still DNS.](#)[Datadog Observes Everything. The Namespace It Observes Has No Governed Structure.](#)[Grafana Unified Observability Visualization. The Data Namespace It Queries Has No Governed Structure.](#)[Prometheus Defined Cloud-Native Monitoring. Its Metric Namespace Has No Governance Layer.](#)[New Relic Pioneered APM. The Telemetry Namespace It Built Is Centrally Indexed.](#)[Splunk Indexes Machine Data at Scale. The Index Namespace Is Centrally Administered.](#)

[Adaptive Indexing overview](#) →

AQ

deterministic

autonomy

Legal

Subject to one or more pending U.S. and international patent applications, see [Patents](#) for the current list and status. No license, express or implied, is granted. Any use requires a separate written agreement—see [Licensing](#). Patent applications referenced on this site are pending. Claim scope, if any, is subject to examination and may issue in altered form or not at all. See [Legal](#) for terms and conditions.

Adaptive Query™ is a trademark of Nicholas Clark. U.S. federal registration is pending. federal registration. AQ™, AQ Inside™, Adaptive Index™, Adaptive Network™, Semantic Agent™, @AQ™, AQID™, and Adaptive Coin™ are used as trademarks in connection with the Adaptive Query platform and brand. Other names may be trademarks of their respective owners.

Platform operated by Adaptive Query LLC, which provides patent and trademark licensing services. Copyright © 2025-2026 Nicholas Clark. All rights reserved.

Last updated: 2026-03-03



- [Inventive Steps](#)
- [Licensing](#)
- [Patents](#)
- [Articles](#)
- [Legal](#)
- [Opportunities](#)
- [Sitemap](#)



-
- nick@qu3ry.net
- 72 28 14 36 01



[Invented by Nick Clark](#) | Founding Investors: Devin Wilkie