



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

## **Bluesky Identified the Right Problem. The Architecture That Solves It Is the Adaptive Index.**

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

Bluesky's AT Protocol separates handles from DIDs, federates data hosting through PDS operators, and enables portable identity, correctly diagnosing why both DNS and blockchain-based naming fail at scale. This article examines the structural dependency that remains: the PLC directory, a single global registry for all did:plc identities that Bluesky has acknowledged as a centralization problem. The path forward requires distributing directory authority into the identity namespace itself through anchor-governed scoped resolution.

---

Bluesky built something worth examining carefully. AT Protocol is a serious attempt to design decentralized social infrastructure from first principles: portable identity, user-controlled data, open federation, cryptographic self-authentication. The team understood that the failure modes of earlier decentralized protocols were not philosophical but architectural, and they set out to address them directly.

They got a lot right. The separation of handle (human-readable, updatable) from DID (persistent, cryptographic) is a sound design. The PDS federation model is real: any host can run a Personal Data Server and any user can migrate their account between hosts. The data repository model means users own their content in a portable, verifiable form. These are genuine structural improvements over both centralized social platforms and earlier decentralized protocols.

The problem is the PLC directory. Bluesky knows it. Their own documentation says so.

## What the PLC directory is and why it matters

Every AT Protocol identity is a DID: a persistent identifier that resolves to a DID document containing the user's signing keys, hosting service, and handle mapping. The most common DID method in the network is did:plc, the Public Ledger of Credentials, developed by Bluesky because no existing DID method met their requirements for consistency, availability, recoverability, and cryptographic security.

The way did:plc works: a central directory server collects and validates operations, and maintains a transparent log of operations for each DID. Bluesky's own GitHub repository for the DID PLC method is direct about this: "A central directory server collects and validates operations." When a client resolves a did:plc identifier, it queries plc.directory — a single global service — to retrieve the current DID document and operation log.

Everything in AT Protocol that depends on identity depends on the PLC directory. Handle resolution requires a DNS lookup followed by a PLC directory query. Identity caches are recommended to have a maximum TTL of 24 hours, after which they must re-validate through the directory. Account migrations, key rotations, and handle updates all propagate through the PLC directory. There are over twelve million registered DIDs in the network as of late 2024. All of them are anchored in a single directory operated by Bluesky Social PBC.

Wikipedia's article on AT Protocol states what the Bluesky team has acknowledged: the did:plc method "has been noted as being a single point of failure within the protocol, as its current implementation relies on a single registry hosted by Bluesky Social, with no system to independently verify the document's current state."

## What Bluesky is trying to do about it

Bluesky's Fall 2025 protocol check-in is candid about where things stand and where they are going. They are working to establish a Swiss association to operate the PLC directory and own its assets, with the explicit goal of making governance independent of Bluesky PBC. They have stated that "we do not expect a single global directory to be the final technical architecture for the system." The direction is clearly toward something more distributed, but the path from here to there has not been fully specified.

The check-in also frames their broader goal as "hard decentralization": making the network less dependent on Bluesky PBC's existence. They are working on this through open PDS federation, independent infrastructure, and governance separation of the PLC system.

The intent is right. The question is what architecture actually achieves it.

## Why a consortium directory does not fully solve it

The intermediate step Bluesky has described for PLC is a permissioned consortium: multiple parties operating the directory under shared governance rather than a single organization. This is a meaningful improvement in resilience and governance accountability. It is not the same as distributing authority into the identity namespace itself.

A consortium directory is still a single global directory. Identity updates still propagate through it. Resolution still requires querying it. The authority over what a DID document means, and who can change it, still flows through a central system — now with multiple operators rather than one, but still structured as a global registry whose approval is required for any identity mutation.

The structural alternative is: each segment of the identity namespace governed by the nodes responsible for it, under locally held policy, with mutations proposed and validated through local consensus. Identity resolution traverses the hierarchy stepwise without querying a central registry. When a user migrates between hosts, the migration is a mutation proposal evaluated by the anchors governing their identity scope, not a global transaction that must be recorded in a single directory. When a key is rotated, the rotation is validated locally and its lineage preserved cryptographically — not through a central operation log.

This is what AT Protocol's own description of the problem points toward. Their requirements for the identity system were: strongly consistent, highly available, recoverable, and cryptographically secure with fast and cheap propagation of updates. An anchor-governed adaptive index satisfies all four. Consistency is provided by local quorum. Availability is provided by the redundancy of the anchor group governing each scope. Recoverability is provided by cryptographically committed lineage. Propagation is local, not global.

## The deeper alignment

What Bluesky got right is the framing: identity portability requires that users, not platforms, be the authority over their own identity. The PLC directory was built to enable that. The gap is that the directory itself is now the platform — a different one from Twitter or Meta, with better intentions and more accountable governance, but still a single point that must be trusted for identity authority to function.

The adaptive index does not replace AT Protocol's identity model. It is the infrastructure layer that would let the PLC directory's function be distributed into the namespace itself. Each PDS operator, each community, each region of the network could govern its own identity scope under locally held policy. The global operation log that the PLC directory currently maintains would become a traversable lineage across anchor-governed scopes, verifiable without querying any single registry.

Bluesky has done the hard work of building an open, federated, portable identity protocol. The last dependency — the central directory — is the one they have already named as a problem. The architecture that eliminates it is the adaptive index.

[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](mailto:nick@qu3ry.net)
- 72 28 14 36 01



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