

Project NANDA FAQ

What is Project NANDA?

Project NANDA is building the foundational infrastructure for an Internet of AI Agents i.e. the system where trillions of AI agents can collaborate, communicate, and transact across organizational boundaries. Created at MIT, NANDA provides the index, protocols, and tools needed to enable this decentralized, protocol-neutral ecosystem.

NANDA addresses the core challenge: how can billions of AI agents discover each other, verify capabilities, and coordinate tasks without creating bottlenecks or security vulnerabilities. The project develops both the technical infrastructure (index, protocols, SDKs) and the governance frameworks needed for a responsible, open Internet of Agents.

Development Stages

Stage 1: Foundations of Agentic Web

Foundation/Onboarding/Index (Index, Cross-platforms)

- Index Infrastructure: NANDA Index for agent discovery and identity
- Cross-Platform Communication: Protocol bridges between A2A, MCP, HTTPS
- Agent Onboarding: SDKs and tools for easy agent deployment
- Interoperability: Standards for agents to work across different platforms

Stage 2: Agentic Commerce

Knowledge pricing, edge AI

- Knowledge Pricing: Mechanisms for agents to value and exchange information
- Edge AI Integration: Distributed intelligence at network edges
- Economic Protocols: Payment and incentive systems for agent services
- Resource Markets: Platforms for trading compute, data, and capabilities

Stage 3: Society of Agents

Large Population Models, co-learning, Agents across data silos (privacy)

- Large Population Models (LPMs): Collective intelligence from agent populations
- Collaborative Learning: Agents learning together while preserving privacy
- Cross-Silo Coordination: Agents working across organizational data boundaries
- Distributed AI: Split learning and inference across agent networks

Goals

Short term goals

Research:

- High-quality algorithmic research and perspective papers (see current research at NANDA Papers Repository)
- Creating partnerships between companies, startups, and academics

Impact:

- Awareness: Educating about the need for an open, neutral Internet of AI Agents
- Responsible Governance: Promoting ethical use and governance frameworks for AI agents

Lab to market (Tech Transfer and Venture Enablement):

- Reference Implementations: Open-source tools and community development
- Ecosystem Building: Connecting startups, corporations, and investors for robust solutions

Long terms goals

Research

- Academic Infrastructure: Establishing journals, conferences, and NSF programs for agentic web research
- Algorithmic Advancement: Developing algorithms to enable the society of AI agents

Impact

- Open Framework: Creating equitable systems for agentic commerce
- Global Standards: Establishing protocols for responsible agent interaction

Lab to Market

- **Government Engagement:** Stimulating public sector adoption and policy development
- **Corporate Unlocking:** Enabling enterprises to safely deploy agent ecosystems

Is Project NANDA needed for AI agents?

As AI agents proliferate, it's helpful to have infrastructure beyond basic message passing. A framework like Project NANDA focuses on three layers:

1. **Open discovery and interoperability** so agents can find and address each other across platforms
2. **Economic primitives** so agents can price, meter, and exchange services
3. **Privacy-aware collaboration** so groups of agents can learn together without exposing sensitive data

Without these layers, ecosystems risk fragmenting into isolated silos that limit scale and shared intelligence.

What would NANDA add beyond today's communication standards?

Existing agent standards are strong once two parties are already connected. NANDA aims to complement them by providing:

- **Global discovery & addressing** across platforms and runtimes
- **Cross-protocol interoperability**, so agents built in different stacks can still interact
- **Verifiable trust signals** (e.g., attestable capabilities and provenance)
- **Privacy-preserving lookups** to protect sensitive relationships and intent
- **Economic building blocks** (quoting, pricing, metering, settlement)
- **Higher-level collaboration** such as knowledge markets and population-scale learning

Potential benefits and challenges

Benefits

- **Reduces vendor lock-in** via neutral discovery and standards
- **Enables cross-organizational collaboration** with privacy safeguards
- **Supports scaling** to millions/billions of agents

- **Opens new business models** where agents trade skills, data, and compute
- **Encourages collective intelligence** through structured, privacy-aware learning

Challenges

- **Achieving sufficient adoption** and network effects
- **Governance and interoperability** across diverse stakeholders
- **Added complexity and potential latency** from verification and cryptography
- **Abuse resistance** (spam, sybil, and malicious agents)
- **Balancing openness** with compliance, safety, and auditability

Overall, NANDA is positioned as the connective tissue that helps agent ecosystems interoperate, transact, and learn responsibly at scale, across silos.

How to participate

- LinkedIn: Follow Project NANDA company page (<https://www.linkedin.com/company/projectnanda/>)
- Weekly Webinars: Join regular community calls and technical discussions
- Discord: Participate in ongoing conversations and working groups

Additional Channels:

- Events: [Lu.Ma/nanda](#)
- YouTube: [@ProjectNANDA](#)
- Newsletter: [Sign up form](#)
- Writing Group: Onboarding process (<https://tinyurl.com/nandawritingonboarding>)

Current Status

Research: See comprehensive research at NANDA Papers Repository

Impact:

- **Community Events:** Regular webinars and conferences documented on YouTube channel
- **Industry Engagement:** Growing awareness of Internet of Agents concept across tech sector

Lab2Market:

- Corporate Coalition: Multiple companies participating in development and standardization
 - Open Source: Active development of reference implementations and SDKs
 - Startup Ecosystem: Supporting ventures building on NANDA infrastructure
-

What is the NANDA Index?

The NANDA Index is a lightweight, decentralized system that functions as the "phone book" for the Internet of AI Agents. Just as DNS maps website names to IP addresses, the NANDA Index maps agent names to their locations and capabilities, enabling billions of AI agents to discover and connect with each other across different platforms and protocols.

Think of it as a universal directory where any AI agent—whether it's a translation assistant, a financial advisor, or a healthcare coordinator—can register itself and be found by other agents or users who need its services.

What is the NANDA Index in the context of the larger Project NANDA?

The NANDA Index is the foundational component of Project NANDA's three-stage vision for building the Internet of AI Agents. Using a city analogy: if A2A protocols are the streets agents use to communicate, Project NANDA builds the entire city, and the NANDA Index serves as the address system.

Project NANDA progresses through three technical phases:

- **Phase 1 - Foundations of Agentic Web:** NANDA Index provides agent identity and discovery infrastructure, plus cross-platform communication protocols and interoperability standards
- **Phase 2 - Agentic Commerce:** Knowledge pricing mechanisms, edge AI integration, economic protocols, and resource markets for agent services
- **Phase 3 - Society of Agents:** Large Population Models (LPMs), collaborative learning systems, cross-silo coordination, and distributed AI networks

The NANDA Index enables all three stages by providing the essential discovery infrastructure that lets agents find and identify each other across the ecosystem. However, it is only one part of the bigger goal of Project NANDA.

How does the NANDA Index work?

The NANDA Index operates through a three-layer architecture:

- 1. **Index Layer (Anchor Tier):** Stores minimal, essential information (≤ 120 bytes per record) including agent IDs, metadata URLs, and routing pointers
- 2. **AgentFacts Layer (Metadata Distribution Tier):** Contains detailed, verifiable information about each agent's capabilities, endpoints, and credentials
- 3. **Dynamic Resolution Layer (Adaptive Routing Tier):** Handles real-time endpoint discovery and load balancing

When you want to find an agent, the process works like this: AgentName \rightarrow NANDA Index \rightarrow AgentAddr \rightarrow AgentFacts \rightarrow Agent Endpoint

What makes the NANDA Index different from existing solutions?

The NANDA Index addresses critical limitations of current internet infrastructure:

| Problem | Traditional Approach | NANDA Index Solution |
|-------------|--|---|
| Scale | DNS handles millions of static records | Designed for billions of dynamic agents |
| Speed | Minute-to-hour updates | Sub-second global resolution |
| Trust | Only proves domain ownership | Cryptographically signed capabilities |
| Privacy | Exposes lookup patterns | Privacy-preserving resolution paths |
| Flexibility | Fixed endpoints | Dynamic, adaptive routing |

Why do we need an index for AI agents?

As AI agents become more autonomous and specialized, we're moving toward a world with billions to trillions of agents that need to:

- Discover each other across organizational boundaries
- Verify capabilities and trustworthiness
- Coordinate complex tasks requiring multiple specialists
- Adapt to changing conditions in real-time

Without a proper index system, this creates an $N \times N$ connectivity problem—every agent would need direct connections to all others. The NANDA Index transforms this into a simpler $2N$ problem by providing a shared discovery mechanism.

What problems does the NANDA Index solve?

The NANDA Index addresses five critical challenges:

1. **Index Bottleneck:** Reduces write operations by 10,000× compared to DNS by separating static identity from dynamic metadata
2. **Trust Gap:** Enables cryptographically signed capability proofs and instant revocation
3. **Privacy Exposure:** Provides anonymous lookup paths that don't reveal who's searching for what
4. **Routing Limitations:** Supports agents that move frequently or require geo-based load balancing
5. **Governance Complexity:** Creates transparent, auditable logs for compliance and accountability

Why not just use DNS or existing systems?

DNS was designed in 1983 for static web infrastructure, not dynamic AI agents. Key limitations include:

- Update frequency: DNS propagation takes minutes to hours; agents need sub-second updates
- Trust model: DNS only proves domain ownership, not agent capabilities or behavior
- Privacy: DNS lookups expose access patterns to network observers
- Metadata richness: DNS records are too simple for complex agent capabilities
- Scale: Current systems can't handle billions of updates per hour

Why not just keep agent names in a file on a website?

While storing agent names in a simple website file might seem straightforward, this approach creates critical problems:

Discovery Gaps: Agents on uncrawled websites will never be found. Agent discovery requires real-time, comprehensive coverage that static file listings cannot provide.

Stale Information: When agents update capabilities or endpoints, they won't be rediscovered until websites are re-crawled, creating windows with outdated information.

Re-centralization Risk: A centralized list inevitably leads to gatekeeping, rent-seeking, biased prioritization, and opaque censorship rules by whoever controls the list.

The NANDA Index solves these problems through its decentralized "quilt" architecture, where multiple registries operate independently while maintaining global interoperability—preventing any single entity from controlling agent discovery.

What is the "Quilt" architecture?

The NANDA Index uses a "quilt-like" architecture that accommodates different types of agent registrations:

| Registration Type | Example | Control Level |
|---------------------|--------------------|---|
| NANDA Native | @agentx | Direct registration in NANDA |
| Government | @US:shop | Location-specific domains |
| Enterprise (Routed) | @company | Access only through company registry |
| Enterprise (Direct) | @company:shop | Visible on NANDA but company-administered |
| Web3 (Routed) | @DID:company | Access through Web3 marketplace |
| Web3 (Direct) | @DID:company:agent | DID-authenticated, NANDA-visible |

This quilt approach allows organizations to maintain control over their agents while enabling global interoperability.

How does the lean index approach work?

The NANDA Index keeps records extremely lightweight (≤120 bytes) by storing only:

- Agent ID and human-readable name
- URLs pointing to detailed metadata (AgentFacts)
- Time-to-live (TTL) values
- Cryptographic signatures

All detailed information is stored separately in AgentFacts documents, which can be updated independently without touching the index. This separation reduces index write overhead by approximately 10,000× while enabling rapid updates.

What are AgentFacts?

AgentFacts are detailed, cryptographically signed JSON-LD documents that contain:

- Capabilities: What the agent can do (translation, analysis, etc.)
- Endpoints: Where to connect (with load balancing options)
- Credentials: Third-party verified certifications
- Performance metrics: Latency, availability, throughput
- Security requirements: Authentication methods and policies

Think of AgentFacts as a verifiable "resume" for each agent that can be hosted either by the agent itself or by neutral third parties for privacy.

How does privacy-preserving resolution work?

The NANDA Index supports dual-path resolution:

- PrimaryFactsURL: Direct access to agent-hosted metadata
- PrivateFactsURL: Anonymous access through third-party hosts (IPFS, neutral proxies)

This allows requesters to discover agent capabilities without revealing their identity or intentions to the agent's infrastructure—critical for competitive or sensitive scenarios.

What types of entities can be registered in the NANDA Index?

The NANDA Index supports a broad ecosystem beyond just AI agents:

- Resources: Compute power, data sources
- Tools: Web scrapers, labelers, specialized software
- MCP servers: Anthropic's Model Context Protocol services
- NLWeb servers: Microsoft's natural language web interfaces
- AI Agents: Individual and enterprise-created agents
- AgentStores: Enterprise AI agent marketplaces
- IoT devices: Internet of Things endpoints
- Web3 agents: Blockchain-based autonomous agents
- Users: Even human agents can be registered

How does NANDA Index relate to existing agent protocols?

The NANDA Index is protocol-agnostic and works with existing standards:

- A2A Protocol: Handles communication after discovery
- MCP (Model Context Protocol): Anthropic's agent context sharing
- NLWeb: Microsoft's conversational web interfaces
- HTTPS: Traditional web protocols

The NANDA Adapter automatically translates between these protocols, enabling seamless cross-platform communication.

What are the current reference implementations?

Project NANDA provides several working components:

- **NANDA Index** (<https://index.projectnanda.org>): The core indexing layer
- **AgentFacts** (<https://list39.org>): Structured agent metadata
- **NANDA Agent Framework** (<https://github.com/projnanda/adapter>): A customizable improvement logic for your agents, and easily get registered into NANDA registry
- **Agent Picker SDK**: (<https://github.com/projnanda/adapter>) Python tools for automated agent deployment
- **Join39 Agent Platform** (<https://join39.org>): Create personal AI agents, upload MCP servers, and join the decentralized network