Independent Submission
Internet-Draft
Intended status: Informational
Expires: 24 September 2026

C. Hood
Nomotic, Inc.

23 March 2026

AGTP Agent Discovery and Name Service
draft-hood-agtp-discovery-00

Abstract

The Agent Transfer Protocol (AGTP) enables agents to communicate once they know each other's canonical identifiers. It does not define how agents find each other. This document specifies the AGTP Agent Discovery and Name Service (ANS): a protocol for dynamic agent discovery using the AGTP DISCOVER method and a governed Agent Name Service that returns ranked sets of Agent Manifest Documents matching a discovery query. ANS servers act as Scope-Enforcement Points for discovery queries and enforce behavioral trust score thresholds, trust tier requirements, and governance zone constraints. This document also defines the DISCOVER method, the Discovery Query language, and the Agent Name Service registration and lookup protocol.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 24 September 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (https://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction

1.1. The Discovery Gap

A deployed AGTP ecosystem contains agents registered across many organizations, governance zones, and trust tiers. An orchestrator agent that needs to delegate a task to a capable peer faces a fundamental problem: it knows what it needs (an agent that can audit smart contracts, translate legal documents, or process medical imaging data) but not who can provide it. Without a standardized discovery mechanism, agents resort to out-of-band registries, hardcoded identifiers, or proprietary marketplace APIs -- all of which are inconsistent with AGTP's governance model.

This document fills that gap. The AGTP Agent Name Service (ANS) is a governed, queryable registry of agent capabilities. The DISCOVER method is AGTP's native query interface to ANS servers. Together they enable dynamic agent federation without requiring prior knowledge of specific agent identifiers.

1.2. Design Principles

Agent-native. Discovery returns Agent Manifest Documents, not web pages or proprietary records. The result of a DISCOVER query is structurally identical to the result of resolving a known agent URI.

Governed. ANS servers are AGTP-aware infrastructure. They enforce trust tier requirements, behavioral trust score floors, and governance zone constraints before returning any result. An ANS server is a Scope-Enforcement Point for discovery traffic.

Ranked. Discovery results are ranked by behavioral trust score, trust tier, and declared capability depth. The highest-quality agents matching the query appear first.

Brokered. ANS servers MAY negotiate authority scope on behalf of the requesting agent, informing it of what scope is required to interact with each discovered agent before commitment.

2. Terminology

The key words "*MUST*", "*MUST NOT*", "*REQUIRED*", "*SHALL*", "*SHALL NOT*", "*SHOULD*", "*SHOULD NOT*", "*RECOMMENDED*", "*NOT RECOMMENDED*", "*MAY*", and "*OPTIONAL*" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals.

DISCOVER: An AGTP Tier 1 method that queries an ANS server for

agents matching a capability description, trust requirement, and governance context. Returns a ranked list of Agent Manifest Documents.

ANS (Agent Name Service): An AGTP-aware server implementing the DISCOVER method and maintaining an index of registered agents and their capabilities. Acts as a SEP for discovery queries.

Behavioral Trust Score: A numeric value in the range 0.0 to 1.0 assigned to an agent at packaging time by a governance platform's pre-packaging verification pipeline. The score reflects the degree of observed behavioral alignment between the agent's declared capabilities, its archetype, and its verified conduct during behavioral testing. A score of 1.0 indicates full alignment; a score of 0.0 indicates disqualifying behavioral discrepancy. The score is embedded in the agent's .agent or .nomo package at issuance, covered by the package integrity hash, and surfaced in the Agent Manifest Document. It cannot be agent-asserted or modified without invalidating the package. ANS servers *MUST* use the behavioral trust score from the verified Agent Manifest Document when ranking DISCOVER results.

Discovery Query: A structured expression of the requesting agent's capability need, expressed as a natural language intent, structured capability criteria, or both.

Agent Manifest Document: As defined in [AGTP] Section 6.4. The canonical representation of an agent's identity and capabilities, returned as a discovery result.

Scope Negotiation: The ANS's declaration of what Authority-Scope the requesting agent will need to interact with each discovered agent.

3. The DISCOVER Method

3.1. Method Definition

DISCOVER is a Tier 1 core AGTP method defined in this document and registered in the IANA AGTP Method Registry. It is the protocol-level interface for querying an ANS server.

Purpose: Query an ANS server for agents matching a capability description and governance constraints. Returns a ranked list of Agent Manifest Documents. The requesting agent *MUST* carry discovery:query in its Authority-Scope header to invoke DISCOVER. Category: ACQUIRE. Idempotent: Yes.

Parameter	Required	Description
intent	*SHOULD*	Natural language description of the needed capability
capability_domains	*MAY*	Structured capability criteria matching DESCRIBE domain names
trust_tier_min	*MAY*	Minimum Trust Tier to include in results (1, 2, or 3)
behavioral_trust_min	*MAY*	Minimum behavioral trust score (0.0-1.0)
governance_zone	*MAY*	Restrict results to a specific governance zone
org_domain	*MAY*	Restrict results to agents from a specific org domain
limit	*MAY*	Maximum number of results to return (default: 10)
scope_negotiate	*MAY*	If true, include required Authority-Scope for each result

Table 1: DISCOVER Parameters

3.2. DISCOVER Request Example

```
AGTP/1.0 DISCOVER
Agent-ID: agtp://agtp.acme.tld/agents/orchestrator
Principal-ID: usr-chris-hood
Authority-Scope: discovery:query agents:delegate
Session-ID: sess-discovery-01
Task-ID: task-disc-001
Content-Type: application/agtp+json

{
  "method": "DISCOVER",
  "task_id": "task-disc-001",
  "parameters": {
    "intent": "Audit smart contracts for security vulnerabilities",
    "capability_domains": ["methods", "tools"],
    "trust_tier_min": 1,
    "behavioral_trust_min": 0.85,
    "governance_zone": "zone:external-partner",
    "limit": 5,
    "scope_negotiate": true
  }
}
```

3.3. DISCOVER Request Example: Domain-Constrained Agent Search

The following example illustrates a targeted DISCOVER query locating agents capable of auditing Solidity smart contracts within a specific governance zone and trust tier threshold. The requesting orchestrator carries discovery:query and agents:delegate scope, signaling that it intends to act on the results by delegating work.

```
AGTP/1.0 DISCOVER
Agent-ID: agtp://agtp.finco.tld/agents/orchestrator
Principal-ID: usr-compliance-team
Authority-Scope: discovery:query agents:delegate audit:*
Session-ID: sess-compliance-q2
Task-ID: task-disc-solidity-01
Content-Type: application/agtp+json

{
  "method": "DISCOVER",
  "task_id": "task-disc-solidity-01",
  "parameters": {
    "intent": "Audit Solidity smart contracts for reentrancy, overflow, and access contro
l misconfiguration",
    "capability_domains": ["methods", "tools", "zones"],
    "trust_tier_min": 2,
    "behavioral_trust_min": 0.88,
    "governance_zone": "zone:finance",
    "limit": 3,
    "scope_negotiate": true
  }
}
```

The ANS server returns a ranked result set. The top result carries Trust Tier 1 with a behavioral trust score of 0.97, exposing the required_scope field indicating what the orchestrator must carry before a DELEGATE to this agent will be accepted:

```
{
  "status": 200,
  "task_id": "task-disc-solidity-01",
  "result": {
    "query_id": "qry-sc-audit-3c9f",
    "total_matches": 7,
    "returned": 3,
    "results": [
      {
        "rank": 1,
        "manifest_uri": "agtp://agtp.auditor.tld/agents/solidity-auditor",
        "canonical_id": "4f7e1a2b9c3d...",
        "agent_label": "solidity-auditor",
        "org_domain": "auditor.tld",
        "trust_tier": 1,
        "behavioral_trust_score": 0.97,
        "supported_methods": ["QUERY", "ANALYZE", "VALIDATE", "REPORT"],
        "capability_match_score": 0.96,
        "governance_zone": "zone:finance",
        "required_scope": "audit:read code:analyze",
        "job_description": "Static and dynamic analysis of Solidity smart
                            contracts. Detects reentrancy, overflow, and
                            access control vulnerabilities."
      }
    ],
    "ans_signature": {
      "algorithm": "ES256",
      "key_id": "ans-key-finance-01",
      "value": "[base64-encoded-signature]"
    }
  }
}
```

The requesting agent verifies the ans_signature before trusting any result. It then issues a DESCRIBE to agtp://agtp.auditor.tld/agents/solidity-auditor to confirm method support before committing to DELEGATE.

3.4. DISCOVER Response

The DISCOVER response body contains a ranked list of Agent Manifest Documents with optional scope negotiation fields:

```
{
  "status": 200,
  "task_id": "task-disc-001",
  "result": {
    "query_id": "qry-7f3a9c",
    "total_matches": 23,
    "returned": 5,
    "results": [
      {
        "rank": 1,
        "manifest_uri": "agtp://agtp.auditor.tld/agents/contract-auditor",
        "canonical_id": "3a9f2c1d8b7e4a6f...",
        "agent_label": "contract-auditor",
        "org_domain": "auditor.tld",
        "trust_tier": 1,
        "behavioral_trust_score": 0.97,
        "supported_methods": ["QUERY", "VALIDATE", "REPORT"],
        "capability_match_score": 0.94,
        "required_scope": "audit:read code:analyze",
        "job_description": "Static and dynamic analysis of Solidity smart contracts."
      }
    ],
    "ans_signature": {
      "algorithm": "ES256",
      "key_id": "ans-key-public-01",
      "value": "[base64-encoded-signature]"
    }
  }
}
```

The ans_signature covers the full result set including all manifests and the query_id. The requesting agent *MUST* verify the signature before trusting any discovery result. An unsigned or invalid DISCOVER response *MUST* be rejected.

3.5. Ranking Algorithm

ANS servers *MUST* rank results by a composite score computed as:

```
rank_score = (trust_tier_weight * normalized_trust_tier)
           + (behavioral_trust_weight * behavioral_trust_score)
           + (capability_match_weight * capability_match_score)
```

Default weights: trust_tier_weight = 0.3, behavioral_trust_weight = 0.4, capability_match_weight = 0.3. ANS servers *MAY* publish alternative weight configurations for domain-specific deployments.

The capability_match_score is computed by the ANS server as a normalized relevance score between the query's intent and capability_domains and the candidate agent's declared capabilities. The computation method is implementation-defined but *MUST* be deterministic for identical inputs.

4. Agent Name Service

4.1. ANS Architecture

An ANS server is an AGTP endpoint that:

1. Maintains an indexed registry of agent capabilities derived from registered agents' DESCRIBE responses and Agent Manifest Documents.

2. Enforces Trust Tier and behavioral trust score thresholds on all queries before returning results.

3. Signs all DISCOVER responses with the ANS's governance key.

4. Acts as a SEP for discovery traffic: agents without discovery:query scope *MUST* be rejected with 451 Scope Violation.

ANS servers are themselves AGTP agents: they have canonical Agent-IDs, Birth Certificates, and Agent Manifest Documents. Their capabilities include discovery:query and discovery:register as Authority-Scope domains.

Multiple ANS servers MAY be deployed for different governance zones, organizations, or trust tier populations. Cross-ANS federation is described in Section 5.

4.2. ANS Registration

An agent registers with an ANS server through a governed process tied to AGTP activation:

1. On completion of the AGTP ACTIVATE transaction, the governance platform automatically submits the agent's Agent Manifest Document and DESCRIBE capability data to the designated ANS server(s) for the agent's governance zone.

2. The ANS server verifies the agent's lifecycle state is Active.

3. The ANS server indexes the agent's capabilities, trust tier, and behavioral trust score.

4. The ANS server updates its result index within 60 seconds.

Manual registration is not supported. ANS registration is a consequence of AGTP activation, not an independent step. This ensures that only Active, registered agents appear in discovery results.

4.3. ANS Deregistration

When an agent's lifecycle state transitions to Suspended, Revoked, or Deprecated, the governance platform *MUST* notify the relevant ANS servers within 60 seconds. The ANS server *MUST* remove the agent from its result index before the next DISCOVER request is processed.

ANS servers *MUST* treat deregistration as urgent: a Revoked agent that continues to appear in discovery results is a governance failure. ANS servers *MUST* log deregistration events.

4.4. Index Freshness

ANS servers *MUST* refresh capability data for indexed agents at a configurable interval (recommended: 24 hours). Refresh is triggered by issuing DESCRIBE requests to each indexed agent's canonical URI. If an agent fails to respond to a DESCRIBE request after three consecutive attempts, the ANS server *MUST* remove it from the index and notify the governance platform.

5. Cross-ANS Federation

5.1. The Federation Problem

An agent in one governance zone may need to discover agents in a different zone. Organization-scoped ANS servers do not have visibility into other organizations' agents.

5.2. Federation Protocol

AGTP supports a federated discovery model in which ANS servers can query peer ANS servers on behalf of requesting agents. A federated DISCOVER query:

1. The requesting agent sends a DISCOVER request to its local ANS server.

2. If the local ANS server's index does not contain sufficient results (fewer than limit results above the requested thresholds), it MAY forward the query to federated peer ANS servers it trusts.

3. Federated queries *MUST* carry the original requesting agent's Agent-ID and scope requirements. The forwarding ANS server *MUST NOT* expand the scope or lower the trust requirements in the forwarded query.

4. Peer ANS servers apply their own governance policies to the forwarded query and return signed result sets.

5. The local ANS server merges, re-ranks, and re-signs the combined result set before returning it to the requesting agent.

6. The merged result includes a federation_path field listing the ANS servers that contributed results.

5.3. Trust in Federation

Federated ANS results *MUST* be treated with the same trust level as the federated ANS server's own Trust Tier claim. An ANS server *MUST* declare the trust tier of its indexed agents accurately; misrepresenting agent trust tiers in federated results is a governance violation.

6. Dynamic Routing and Failover

6.1. Capability-Based Routing

ANS servers *MAY* support dynamic routing: when a requesting agent's preferred target agent is unavailable (lifecycle state Suspended, not responding, or rate-limited), the ANS server *MAY* return an equivalent agent from its index as a routing alternative.

Dynamic routing responses *MUST* indicate: - routing_reason: the reason the primary agent was not returned - equivalent_confidence: a score (0.0-1.0) indicating how closely the alternative matches the primary agent's capabilities

The requesting agent retains the right to reject the alternative.

6.2. Agent Availability Signals

ANS servers *SHOULD* monitor indexed agents' availability by periodically checking the /status sub-resource defined in [AGTP] Section 6.5.4. Agents with lifecycle_state: Suspended or active_session_count at capacity *SHOULD* be marked as unavailable in routing responses.

7. Security Considerations

7.1. Discovery Result Injection

Threat: A malicious ANS server returns forged Agent Manifest Documents pointing to attacker-controlled agents.

Mitigation: All DISCOVER responses *MUST* be signed by the ANS server's governance key. The requesting agent *MUST* verify the ans_signature field before trusting any result. The ANS server's governance key *MUST* be resolvable via the ANS server's own Agent Manifest Document, creating a verifiable trust chain.

7.2. Discovery Enumeration

Threat: A malicious agent uses repeated DISCOVER queries to enumerate all agents in a governance zone, gaining knowledge of the organization's agent architecture.

Mitigation: ANS servers *MUST* enforce rate limiting per requesting Agent-ID. ANS servers *MAY* require discovery:query scope to include agents from restricted governance zones. ANS servers *SHOULD* log high-frequency DISCOVER queries as anomaly signals.

7.3. Behavioral Trust Score Manipulation

Threat: An agent artificially inflates its behavioral trust score to appear higher in discovery results.

Mitigation: Behavioral trust scores are computed by the governance platform's pre-packaging verification pipeline and embedded in the agent's .agent or .nomo package at issuance time. The score is part of the signed manifest and cannot be modified without invalidating the package integrity hash. ANS servers *MUST* use the trust score from the verified Agent Manifest Document, not from any agent-asserted field.

7.4. Scope Negotiation Abuse

Threat: An ANS server inflates the required_scope field in discovery results, causing requesting agents to request broader scope than necessary.

Mitigation: Requesting agents *MUST NOT* blindly request the scope declared in discovery results. The required_scope field is informational; agents *MUST* evaluate the scope against their own authorization and request only what is needed for the specific task. Governance platforms *SHOULD* flag scope requests that significantly exceed the scope of prior tasks in the same session.

8. IANA Considerations

8.1. DISCOVER Method Registration

This document requests registration of the DISCOVER method in the IANA AGTP Method Registry established by [AGTP] Section 9.2:

Method	Status	Reference
DISCOVER	Permanent	This document, Section 3.1

Table 2: DISCOVER Method Registry Entry

8.2. New Authority-Scope Domains

This document uses the following Authority-Scope domains defined in [AGTP] Appendix A:

Domain	Used For
discovery:query	Sending DISCOVER requests to ANS servers
discovery:register	ANS registration operations (governance platform only)

Table 3

No new domains are requested; these are from the existing reserved set.

9. References

9.1. Normative References

[AGTP] Hood, C., "Agent Transfer Protocol (AGTP)", Work in Progress, Internet-Draft, draft-hood-independent-agtp-02, 2026, <https://datatracker.ietf.org/doc/html/draft-hood-independent-agtp-02>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/rfc/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

[RFC8555] Barnes, R., Hoffman-Andrews, J., McCarney, D., and J. Kasten, "Automatic Certificate Management Environment (ACME)", RFC 8555, DOI 10.17487/RFC8555, March 2019, <https://www.rfc-editor.org/rfc/rfc8555>.

9.2. Informative References

[AGTP-CERT] Hood, C., "AGTP Agent Certificate Extension", Work in Progress, Internet-Draft, draft-hood-agtp-agent-cert-00, 2026, <https://datatracker.ietf.org/doc/html/draft-hood-agtp-agent-cert-00>.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/rfc/rfc1034>.

[RFC7519] Jones, M., Bradley, J., and N. Sakimura, "JSON Web Token (JWT)", RFC 7519, DOI 10.17487/RFC7519, May 2015, <https://www.rfc-editor.org/rfc/rfc7519>.

Appendix A. ANS Deployment Considerations

Small organizations may not operate dedicated ANS servers. The AGTP governance platform *MAY* provide built-in ANS functionality as part of the agent registry. In this case, the governance platform’s registry endpoint serves both registration (ACTIVATE) and discovery (DISCOVER) queries.

Large organizations with hundreds or thousands of agents *SHOULD* deploy dedicated ANS infrastructure with appropriate indexing and caching. The ANS index is a read-heavy workload; standard caching and replication patterns apply.

Cross-organization federation requires bilateral trust establishment between ANS operators. The protocol for establishing ANS federation trust is out of scope for this document but follows the same patterns as AGTP Trust Tier 1 verification: DNS ownership challenge and mutual certificate exchange.

Author's Address

Chris Hood
Nomotic, Inc.
Email: chris@nomotic.ai
URI: https://nomotic.ai