

Governed Discovery for the Multi-Agent Enterprise

How ServiceNow AI Control Tower and AGNTCY Agent Directory Service Deliver Federated, Governed Multi-Agent Workflows at Scale

Chris Henson

Senior Staff Software Engineer, ServiceNow

Karthik Ramachandran

Senior Director of Product Management, ServiceNow

Ola Mabadeje

Principal Product Manager (Director), Outshift by Cisco

John Parello

Principal Engineer, Outshift by Cisco

Executive Summary

Enterprise AI is entering its most consequential phase. Organizations are deploying dozens and sometimes hundreds of specialized AI agents across IT service management, network operations, security, HR, and customer experience. These agents come from different vendors, run on different platforms, and evolve on different release cycles. Without a unifying architecture that delivers both centralized governance and distributed discovery based on open standards, the result is predictable: fragmented AI silos, ungoverned shadow agents, duplicated capabilities, and workflows that break the moment they need a specialist from a neighboring team.

This whitepaper presents a reference architecture and composable solution stack, jointly developed by Outshift by Cisco and ServiceNow. Built on commercially supported foundations and production grade open source components, the implementation shows how two complementary platforms i.e., ServiceNow's AI Control Tower (AICT) and AGNTCY's Linux Foundation backed open-source Agent Directory Service (ADS) both integrate to deliver a new operational model we call Governed Discovery: the ability to centrally govern every AI agent's lifecycle, risk posture, and business alignment through AICT, while simultaneously making those agents discoverable and composable across company and organizational boundaries through ADS's distributed, standards based directory network.

The implications are significant. This architecture eliminates the false choice between central governance and autonomy, delivering both through a layered integration that respects existing ServiceNow investments while opening the door to multi-vendor, cross-organizational agent ecosystems built entirely on open standards.

The Enterprise Multi-Agent Challenge

The trajectory is clear. Gartner projects that by 2028, enterprises using AI governance platforms will achieve 30% higher customer trust ratings and 25% better regulatory compliance scores. But governance alone does not solve the problem. The real challenge facing CIOs, CTOs, and enterprise architects today is threefold.

Fragmented visibility. For example, when ITSM teams need to deploy ServiceNow AI Agents, and Network operations teams need to execute another vendor's Agentic topology analyzers and digital twin simulators, while Security teams need to operate third-party threat models, there is no single pane of glass that covers the full agent landscape, and there is no standardized schema that describes what each agent can do.

Siloed discovery. When a cross functional change request requires AI agents to perform impact analysis, compliance checking, and configuration validation, the workflow stalls. This stall is not because the right agents don't exist, but because the multi-agent orchestration layer has no mechanism to find them across departmental, vendor or company boundaries. Discovery is manual, brittle, and dependent on tribal knowledge.

Governance without interoperability. Centralized governance platforms track and approve

agents, but their registries are mostly platform native. A governance approved agent is invisible to orchestrators outside the native platform. Real-time agentic workflows operating across vendor ecosystems are unable to trigger real-time capability based agent discoveries, relying only on the governance metadata.

The composable solution stack described here directly addresses all three gaps by coupling AICT's governance control plane with ADS's federated discovery data plane through a standards-based integration pipeline.

Architecture: Two Layers, One Ecosystem

The functional implementation separates concerns into two complementary layers that interlock through standards based integrations. Understanding each layer is essential for architects evaluating this approach.

ServiceNow AI Control Tower: The Governance Control plane

AICT provides the centralized command center for managing, governing, and measuring AI agents across the enterprise. Built on ServiceNow's Configuration Management Database (CMDB) and Common Service Data Model (CSDM), it maps every AI asset which includes models, agents, datasets, prompts etc. to the business services they support. For enterprise IT, AICT delivers the following:

Enterprise-wide visibility across departments and vendors from a single dashboard. In the integrated solution, a Central IT Ops admin viewed all ITSM and Networking Ops agents, their governance states, asset types and departmental ownership, all in one consolidated list.

Embedded GRC workflows enforcing compliance before any agent enters production. Automated fairness audits, model drift detection, explainability, and regulatory mapping to ISO/IEC 42001, the EU AI Act, GDPR, and HIPAA are built in.

Business-outcome alignment through dynamic dashboards connecting AI performance to productivity, revenue, and customer satisfaction KPIs. This answers the questions of "what agents exist" as well as "are they delivering value?"

Action Fabric for real-time A2A and MCP-based agent communication, with integrations from Cisco, Accenture, Adobe, Box, Google Cloud, IBM, Microsoft, Zoom, and others.

AGNTCY Agent Directory Service: The Federated Discovery Data Plane

ADS occupies a fundamentally different layer. It is a distributed, open-source directory for discovering, indexing, and verifying agent capabilities across heterogeneous multi-agent systems. Built on the [Open Agentic Schema Framework](#) (OASF), ADS delivers:

Peer-to-peer federation using a Kademlia-based DHT connecting ADS instances across departments, organizations, and companies. In our joint reference architecture, the ITSM

ADS and NetOps ADS operated as federated peers, each with its own agent records, enabling cross-directory discovery. This is horizontal scaling: add a node, federate it, and the entire network gains access to its agents.

Hierarchical skill taxonomies enabling multi dimensional capability queries. Every record includes structured OASF skill paths (e.g., infrastructure.network.topology-analysis). A supervisor or orchestrator agent describes the skill it needs, and the DHT resolves matching records across the network.

Content-addressed, immutable records identified by cryptographic CIDs with Sigstore-based provenance signing. Every agent's metadata is tamper-evident and globally unique i.e., verifiable before it is trusted.

Open, IETF-track standards with a vendor-neutral reference implementation in Go, deployable via Helm on any Kubernetes cluster with gRPC and REST APIs. ADS is a neutral substrate that works alongside any governance platform.

The Integration in Action: Functional Implementation Walkthrough

Built around an enterprise scenario, we demonstrated three critical workflows validating the Governed Discovery model.

Workflow 1: Centralized Governance Visibility

We start from a Central IT Ops administrator's perspective in the AICT dashboard. A single consolidated view shows every AI asset across ITSM Ops and Networking Ops with the asset name, department, type, and governance state. This is the governance baseline: before any agent is discoverable in the ADS federated network, it must exist in AICT with an explicit governance posture.

ROI Insight: Centralized visibility eliminates shadow AI. Ungoverned deployments are consistently the top source of compliance violations and security incidents. AICT makes every agent visible and accountable.

Workflow 2: Search-Fetch-Load: – The AICT-to-ADS Sync Pipeline

When an enterprise's ServiceNow subscription delivers new AI agents that have been approved in AICT but are not yet registered in the departmental ADS, the ITSM Ops admin or the ADS can dynamically trigger a three stage pipeline:

Search. The pipeline queries AICT's Table API for governance-approved agents absent from ADS, filtering by department, governance state, and asset type and only returning agents that have cleared AICT's compliance and risk workflows.

Fetch. For each discovered agent, AICT metadata is transformed into OASF compliant

directory records. This maps CMDB fields (CI name, operational status, business service relationships) to OASF surfaces (skills arrays, domain taxonomies, A2A Agent Cards, governance_metadata modules). The system ensures the same AICT record always produces the same globally unique identifier.

Load. Transformed OASF records are published into the ITSM ADS via gRPC. The three new AI agents in our reference implementation are immediately discoverable across all federated peer directories, growing the directory from 5 agent records to 8.

ROI Insight: Search-Fetch-Load automates what is typically a manual, error prone registration process. Every governance decision in AICT automatically propagates to the discovery layer – zero duplicate data entry, zero metadata drift, zero stale records. Annotation of records at the departmental level adds uniqueness and relevance that improves quality of workflow outcomes.

Workflow 3: Local Department Annotation

Local annotation of canonical records. When the NetOps ADS fetches a newly approved version of an agent record e.g., Test Executor agent from the AICT, the NetOps engineer is able to annotate the OASF record e.g., define a Change approval board (CAB) constraint, specific to a NetOps workflow. For example, the NetOps engineer can specify workflow annotations on that agent to ensure that change requests that include a CAB requirement always require the NetOps Test Executor AI Agent, due to the tools it uses for task execution – CAB approved testing tools. The NetOps ADS instance stores the annotations exclusively to the local ADS instance. When another department's ADS synchronizes canonical Test Executor agent records, it receives the immutable AICT fetched record without any cross-department CAB metadata.

Runtime discovery with workflow relevance. During workflow runtime discovery, when a multi-agent supervisor or orchestrator queries a department's ADS for an agent by skill, the discovery response includes (1) the canonical OASF fields (name, version, skills, locators, core extensions) and (2) the department-specific CAB annotations. The consuming supervisor or orchestrator can parse these annotations to configure its invocation behavior before calling the agent. This query time behavior means that discovery is always cognizant of departmental agent record uniqueness and workflow relevance, beyond the generally available records fetched from the AICT.

ROI Insight: Without local department annotations, enterprises face a binary choice: one global record that satisfies no department fully, or separate agent artifacts per department that multiply governance overhead. Imagine 50 shared agents fetched into two separate federated departmental ADSs. Without the annotation feature, each department either has exact copies of 50 generic agents, or each will need a distinct artifact that must be independently versioned, signed, registered, approved observed and audited. At two departments and 50 shared agents, that's 100 agent artifacts instead of 50 i.e. 100 AICT Configuration Items to maintain. This compounds the operational cost at the governance and workflow execution layers as agent counts and ADS counts grow.

Workflow 4: Dynamic Multi-Agent Discovery for Change Verification

The final workflow demonstrates the full power of Governed Discovery. A NetOps engineer receives a change request which specifies a firewall rule modification from the CRM team. The NetOps engineer inputs a corresponding natural language query/prompt into the Multi-agent change verification application called the [Outshift Agentic Network Validator \(ANV\)](#), that launches a Change Verification workflow. A NetOps Supervisor Agent orchestrates the process in the following manner:

Skill Decomposition. The Supervisor analyzes the change request/natural language user prompt and determines the required OASF skill paths: network topology analysis, impact simulation, ticket quality assessment checking, change risk assessment, and configuration validation.

Federated Directory Query. The Supervisor issues a semantic search across the NetOps ADS, which queries all federated peers (including the ITSM ADS). The DHT routes capability queries to matching records regardless of which directory hosts them.

Governance-Aware Selection. Discovered agents are filtered by their relevance to the change request query and governance metadata module. Only agents with governance_state = "approved" are eligible. AICT's governance decisions are enforced at the point of agent selection, not just at registration.

Skills-based Team Assembly. A2A Agent Card data from OASF modules determines communication endpoints. The output is a fully assembled specialist team with agents selected based on their skills and relevance score that matches the user prompt query/change request. A workflow plan is generated, mapping each verification step to a specific agent, its matched skill, its source ADS, and its expected output type.

ROI Insight: Dynamic, skill-based team assembly replaces static agent configurations. New agents approved in AICT and synced to ADS are automatically available for future workflows – no code changes, no workflow redesign, no manual routing tables. The system evolves as your agent ecosystem evolves.

Business Value and Strategic Differentiation

For enterprise leaders, the strategic value extends well beyond the technical integration. AICT + ADS unlocks five categories of business impact.

1. Horizontal Scaling and Autonomy Without Governance Fragmentation

Deploy department level or cross-company ADS nodes that federate through the DHT while AICT maintains a single governance layer across all of them. Add a business unit, geography, or partner organization and then federate a new ADS node, sync from AICT and annotate records to improve relevance for departmental workflows. The network expands without increasing governance complexity. For CIOs, this is constant cost scaling of AI agent oversight rather than linear cost growth.

2. Multi-Vendor Interoperability on Open Standards

OASF provides a vendor-neutral schema describing any agent's capabilities, A2A protocols, and governance posture. ADS's OCI-compliant storage and Sigstore provenance mean records are distributable and verifiable by any standards compliant system. The implementation demonstrated ServiceNow AI Agents and Cisco network analysis tools coexisting in the same federated directory, discoverable through shared taxonomies, and composable into unified workflows.

3. Cross-Organization and Cross-Company Federation

The DHT model federates across partner organizations, managed service providers, and industry consortia and not just internal departments. Every record carries governance metadata, enabling consuming organizations to make trust decisions based on governance posture, provenance signatures, and organizational identity. This opens the door to governed agent marketplaces for enterprises with complex partner ecosystems.

4. Compliance that Travels with the Agent

Traditional governance enforces compliance at a single checkpoint that becomes invisible downstream. AICT + ADS embeds governance state directly into portable OASF records. Supervisor Agents inspect each candidate's governance status by reading cryptographically signed metadata without requiring live callbacks. For CISOs, governance decisions are verifiably enforced at every consumption point across the federated network.

5. Dynamic Workflow Composition Reduces Operational Toil

Skill-based composition eliminates hardcoded agent references, manual routing updates, and workflow redesign when capabilities change. For SREs, this means reduced MTTR,

fewer change-related incidents from missing verification steps, and significantly lower operational toil in maintaining multi-agent orchestrations.

Observability and Continuous Improvement

A governed, federated architecture demands observability across both layers. AICT provides business-level dashboards connecting agent performance to enterprise KPIs, lifecycle health, and GRC posture. The [AGNTCY Observe SDK](#) complements this with OTel-compliant multi-agent telemetry: handoff tracing across A2A interactions, fork/join tracking for parallel execution, cross-process distributed traces, and a Metrics Computation Engine evaluating correctness, consistency, and robustness. Together, these deliver end-to-end visibility from boardroom KPIs to individual agent interactions.

Architecture Principles for Adoption

Govern centrally, discover federally. AICT is the single governance authority. ADS nodes deploy per department or boundary, federated through the DHT, synchronized via dynamic Search-Fetch-Load.

Embed governance in the record. OASF governance metadata modules carry governance state and policy identifiers so consuming systems enforce compliance at the point of use, not just the point of origin.

Adopt OASF as the canonical schema. Whether agents come from ServiceNow, Cisco, Microsoft, or internal teams, OASF ensures discoverability across any federated ADS network.

Design for dynamic composition. Supervisor agents query ADS by skill taxonomy at runtime – never hardcode agent identifiers – so workflows automatically absorb new capabilities.

Instrument both layers. Combine AICT business dashboards with the AGNTCY's Observe SDK's OTel tracing for complete telemetry.

Conclusion

This standard based composable solution stack implemented by Outshift by Cisco and ServiceNow demonstrates that the perceived tension between centralized governance and distributed autonomy is a false dichotomy. By integrating AICT as the governance control plane with ADS as the federated discovery data plane, enterprises gain a composable architecture that scales multi-agent deployments without sacrificing compliance, visibility, or operational control.

For the CIO, this means a single governance authority across every AI vendor in the portfolio. For the CISO, compliance decisions are cryptographically verifiable and enforced at

every consumption point. For the enterprise architect, an open-standards foundation based on OASF, A2A, MCP, OCI, Sigstore, avoids vendor lock-in. For platform engineers and SREs, self-evolving workflows dynamically compose the right agent team for every operational challenge.

The multi-agent enterprise is not a future state. It is happening now, across every industry, in every organization deploying AI at scale. The question is no longer whether to govern and federate your agent ecosystem, but how quickly you can build the architecture to do it. Governed Discovery – powered by AI Control Tower and Agent Directory Service – is that architecture.

About This Whitepaper

Based on a composable solution stack jointly developed by Outshift by Cisco and ServiceNow. AGNTCY is a Linux Foundation open-source project (agntcy.org). ServiceNow AI Control Tower is commercially available on the ServiceNow AI Platform. For architecture details, integration patterns, and deployment guidance, contact the respective teams.



outshift.com ↗

@ outshiftbycisco