Multi-Agent Change Verification: From Ticket to CAB in Minutes, Not Weeks

How Multi-Vendor AI Agents, a Network Digital Twin, and Open Standards Transform Enterprise Change Management Workflows

Chris Henson

Senior Staff Software Engineer, ServiceNow

Jason Leung

Director of Product Management, ServiceNow

Ola Mabadeje

Principal Product Manager (Director), Outshift by Cisco

John Parello

Principal Engineer, Outshift by Cisco

April 2026

Executive Summary

Part 1 of this series established the architectural foundation for Governed Discovery — how ServiceNow's AI Control Tower (AICT) and AGNTCY's Agent Directory Service (ADS) integrate to provide centralized governance with federated agent discovery. Part 2 answers the question that every operations leader asks next: what happens when those governed, discovered agents actually go to work?

This whitepaper walks through the end-to-end execution of a multi-agent network change verification workflow, demonstrated in a composable solution stack jointly developed by Outshift by Cisco and ServiceNow. The workflow transforms a sparse, two sentence change request into a fully documented, CAB ready evidence packet that is complete with ITSM ticket quality assessment, network and service risk analysis grounded in real network topology, a comprehensive test plan, and verified test execution against a Network Digital Twin, all executed within minutes rather than the days or weeks that manual processes typically require.

The agents involved are multi-vendor (ServiceNow AI Agents and Outshift by Cisco's Agentic Network Validator (ANV) agents), multi-skilled (each contributing specialized reasoning to unique workflow steps), and multi-purpose (operating across ITSM ticketing, network topology analysis, security policy validation, and pre-production simulation). They communicate through the A2A protocol, are described by OASF compliant metadata, and derive their execution context from authoritative sources of truth i.e., the ServiceNow Configuration Management Database for service relationships and a Network Digital Twin for real-time topology and configuration data.

The Change Management Problem

Network change management is among the most labor intensive, error prone, and consequential workflows in enterprise IT. A single firewall policy update can affect dozens of devices across multiple security zones, and the consequences of inadequate verification range from service outages to security breaches. Yet the process remains stubbornly manual across most enterprises.

A typical change request journey: the engineer receives a request, manually researches affected infrastructure, writes an implementation plan and justification, creates test cases from memory, submits for CAB review, gets rejected for incomplete documentation, revises and resubmits, and eventually executes the change with manual verification. The cycle from initial request to approved execution often stretches from days to weeks, with most time spent not on engineering judgment but on documentation, research, and administrative rework. The composable solution stack demonstrates that this entire workflow can be

orchestrated by specialized AI agents that reason dynamically on each unique change request, drawing context from authoritative enterprise data sources.

The Agent Team: Multi-Vendor, Multi-Skilled, Multi-Purpose

The change verification workflow is executed by a team of specialized agents from two vendors, coordinated by a NetOps Supervisor Agent. Each agent brings domain specific reasoning capabilities, and each was dynamically discovered through the federated ADS network described in Part 1. The team operates as follows:

NetOps Supervisor AI Agent (Orchestrator) — Coordinates end-to-end workflow, manages A2A inter agent communication, creates and updates the ITSM change request ticket, and assembles the final CAB evidence packet.

Quality Assessor Agent (ServiceNow) — Evaluates change request completeness against standards, identifies missing fields, generates AI-powered content recommendations, and scores ticket quality.

Risk Analysis Agent (ServiceNow) — Performs risk assessment using historical ticket data, evaluating change patterns, failure rates, and impact precedents from the ServiceNow CMDB.

Network Impact Agent (Outshift by Cisco ANV) — Collaborates with the Risk Analysis Agent as it queries real-time topology, computes blast radius, runs traffic simulations, and identifies unintended security policy violations, thereby providing networking context to the Risk Analysis and Test Planner agents.

Test Planner Agents (ServiceNow + Outshift by Cisco ANV) — Both collaborate to produce a comprehensive test plan: the ITSM test planner generates structure from the services perspective and best practices, while the network-side planner enriches it with topology aware test cases from actual device configurations.

Test Executor Agent (Outshift by Cisco ANV) — Executes the full test suite against the Network Digital Twin: forks a production snapshot, applies the config candidate, recomputes the model, and runs all verification tests in isolation.

> Architecture Insight: Every agent is A2A compliant, OASF described in the federated ADS, and governance approved through AICT. The Supervisor discovers specialists by skill taxonomy and verifies governance status before workflow execution as described in Part 1.

Phase 1: Change Request Input and Orchestration

The workflow begins when a NetOps engineer initiates a change request from the Networking Ops Supervisor. The change request in this case, is a CRM team's request for firewall policy updates. The engineer provides a brief description and the Supervisor Agent immediately takes over: it creates a formal ITSM change request ticket (e.g., CHG0030074), displays the full team of collaborating agents (showing each agent's metadata such as A2A compliance status and matched skills), and generates a networking perspective summary of the proposed change. The engineer has full transparency into which agents will execute which steps before the workflow proceeds.

Phase 2: Intelligent Quality Assessment

The ServiceNow Quality Assessor AI Agent evaluates the ticket against CAB submission standards, immediately identifying five incomplete critical fields: implementation plan, test plan, risk analysis, backout plan, and justification. The agent reasons about what each field requires and generates context aware content recommendations. The ticket receives an initial score of "Poor" with a clear explanation. The engineer reviews the suggestions, accepting recommendations for implementation plan, backout plan, and justification, but deliberately deferring test plan and risk analysis to specialist agents that will generate those from multi-agent collaboration using shared context from authoritative data sources i.e. live network topology and the CMDB. The Quality Assessor automatically updates accepted fields in the ITSM ticket.

> ROI Insight: Incomplete change requests are the single largest source of CAB rejection and rework. The Quality Assessor catches deficiencies at creation time, eliminating multi day feedback loops. Engineers spend minutes reviewing AI suggestions instead of hours writing documentation.

Phase 3: Risk Analysis and Network Impact Assessment

This phase demonstrates one of the most powerful capabilities of the multi-vendor agent architecture: two agents from different vendors reasoning collaboratively on the same change request, each drawing from its own authoritative data source.

The ServiceNow Risk Analysis AI Agent performs preliminary risk assessment by analyzing historical change tickets in the CMDB and identifying patterns from similar past changes, their success/failure rates, and common failure modes.

The Network Impact Agent (Outshift ANV) then performs network infrastructure level analysis via direct tool calls to the Network Digital Twin by querying live topology, identifying 9 affected devices across 2 security zones, computing blast radius, running traffic simulations, and flagging 2 unintended permit rules as security violations. It also performs failure resilience analysis, assessing network behavior if individual components fail after the change.

The combination is greater than the sum of its parts. The ITSM agent provides historical and organizational risk context while the network agent provides real-time infrastructure impact assessment. Neither could produce a complete risk picture alone. The Risk Analysis Agent updates the ITSM ticket's risk field with the combined analysis.

> ROI Insight: Each unintended permit rule discovered before production is a potential security incident prevented. Traditional change processes rarely include pre-production impact assessment. The Network Digital Twin makes this a standard step, not an exceptional one.

Phase 4: Collaborative Test Planning

Test planning follows the same multi-vendor collaborative pattern. ServiceNow Test Planner AI Agent creates a draft test plan from change management best practices as well as context from Phase 3 i.e., risk analysis context. The ANV Test Planner then refines the test plan, basing its execution on provided context, the change intent and by querying the Digital Twin for current topology and device configurations, thereby grounding every test case in actual infrastructure state and desired intent. This is a critical step where the agent reasons through the intent and context to curate a few high impact relevant tests. This approach contrasts with current automated testing practices where a myriad of tests cases are performed with minimal context, on the configuration candidate, and a good percentage of these pre-determined tests have no relevance to the change request intent. This current practice leads to wasting expensive resources and time on irrelevant testing that fails to catch production network impacting change-induced failures.

The multi-agent collaboration produces 8 test cases in three categories: 2 positive validation tests (confirming intended firewall rules work), 3 negative security tests (verifying unauthorized traffic is blocked and blocker violations are addressed), and 3 regression tests (ensuring existing policies are unaffected). Each case includes specific device references, expected outcomes, and pass/fail criteria from the real topology. The plan is written to the ITSM ticket automatically by the ServiceNow Test Planner AI Agent.

ROI Insight: Intent based and topology grounded test plans catch failures that template based plans miss. The 3 negative security tests exist specifically because the Network Impact Agent discovered unintended permits in Phase 3 — a plan without that context would have missed them.

Phase 5: Quality Re-Assessment

With risk analysis and test plan fields now populated by the multi-agent collaborative process via specialist agents, the Quality Assessor performs a second evaluation. All five required fields are present and substantive. The quality score improves from “Poor” to “Excellent,” automatically written to the ticket’s work notes by the NetOps Supervisor AI Agent. This closed loop pattern of ticket assessment, enrichment and re-assessment runs within a single automated workflow. The engineer never manually updates a quality checklist, never copies results between systems, and never wonders whether the ticket is CAB ready.

Phase 6: Pre-Production Test Execution via Network Digital Twin

This is the phase that has no manual equivalent in most enterprise change processes. The Test Executor Agent (Outshift ANV) takes the 8 test cases from Phase 4 and executes them against the Network Digital Twin through a deterministic four step sequence:

Step 1 — Fork production snapshot. The agent creates an isolated copy of the current production network state in the Digital Twin, ensuring that test execution cannot affect a live knowledge graph representation of the infrastructure.

Step 2 — Pull and apply configuration candidate. The engineer provides a PR link (e.g., PR #202) containing the proposed firewall rule changes. The agent pulls the configuration candidate from the repository and applies it to the forked snapshot.

Step 3 — Recompute network model. The Digital Twin recomputes the complete network model with the proposed changes applied, updating routing tables, security zone boundaries, and policy evaluation paths.

Step 4 — Execute all verification tests. All 8 test cases run against the recomputed model. The implementation showed a 100% pass rate (8/8), covering all three test categories: positive validation, negative security, and regression.

Every tool call is timestamped and logged, providing a complete audit trail of the pre-production verification process. The results include a pass/fail status for each test case, actual versus expected outcomes, and specific device level details for any findings.

> ROI Insight: The Network Digital Twin transforms test execution from theoretical to verifiable. Changes are proven safe before they touch production. This provides a level of pre-production rigor that fundamentally reduces change-induced incident rates.

Phase 7: Final Aggregation and CAB Evidence Packet

The Supervisor Agent aggregates every output from the multi-agent workflow into a single, structured CAB evidence packet: the quality assessment (Excellent), the complete risk analysis (with NDT validation and blocker findings), the 8-case test plan (categorized by type), and the NDT execution results (8/8 passed). The packet is assembled as a PDF document and attached directly to the change request ticket's work notes, alongside detailed summaries of all four major workflow outputs.

The NetOps engineer reviews the complete evidence packet on the ticket, verifies all documentation is present, and submits for CAB approval. The entire journey from a two sentence change description to a fully documented, tested, and verified CAB submission is completed in a single workflow session without the engineer leaving the single User Interface.

> ROI Insight: CAB evidence compilation as currently done today i.e. copying results from disparate tools, formatting docs, chasing test evidence etc., is eliminated entirely. Automated AI agent assisted aggregation produces more comprehensive, more consistent packets than manual compilation ever achieves.

Strategic Value for the Enterprise

Multi-Vendor Agents as a Force Multiplier

The composable solution stack demonstrates that multi-vendor agent teams are not a compromise i.e., they are an advantage. ServiceNow AI Agents bring deep ITSM expertise: ticket quality standards, historical change patterns, compliance requirements. Outshift's

ANV agents bring infrastructure intelligence: real-time topology awareness, security policy simulation, pre-production Digital Twin testing. No single vendor delivers both. The A2A protocol and OASF schema make this interoperability seamless without custom integration code.

Deterministic Workflows with Dynamic Reasoning

The workflow is deterministic i.e., every change request follows structured phases as designed by the Networking Ops team, but the reasoning within each phase is agentic and dynamic. The test plan for a firewall policy change looks nothing like one for a routing table update, because each agent reasons against the change intent, the specific topology, risk profile, and service dependencies of that unique change. This is the distinction between automation scripts (predefined steps) and agentic intelligence (reasoning about each step in the context of the specific problem).

Authoritative Context, Not Hallucinated Analysis

Every analytical output is grounded in authoritative source of truth data. Risk analysis draws from CMDB historical tickets. Network impact queries live topology from the Digital Twin. Test cases reference actual device configurations and zone boundaries. Test execution runs against a forked production snapshot. No agent generates analysis from parametric knowledge alone and every claim is traceable to a data source the enterprise controls and trusts.

From Weeks/Days to Minutes: The Change Management ROI

The aggregate time savings are transformative. Quality assessment and documentation: hours become minutes. Risk analysis across multiple systems: a single automated phase drawing from two authoritative sources. Intent based and topology grounded Test planning: a relevant, prioritized and categorized test suite. Pre-production verification (often skipped entirely): a standard, automated step. CAB evidence compilation: instant aggregation replacing hours of copy & paste.

For enterprises processing hundreds of network changes monthly, the cumulative impact on cycle time, first pass CAB approval rates, and change related incident reduction is substantial.

Conclusion

The composable solution stack implemented by Outshift by Cisco and ServiceNow demonstrates that multi-agent change verification is not a theoretical future but an

operational reality. Multi-vendor agents collaborating through open standards (A2A, MCP, OASF) and drawing context from authoritative enterprise data sources (CMDB, Network Digital Twin) can transform the most documentation heavy, error prone workflows in enterprise IT into a streamlined, verifiable, and auditable process.

Combined with the Governed Discovery architecture from Part 1, the complete picture emerges: AICT governs the agent lifecycle. ADS provides federated discovery so the right specialists are always available. And the multi-agent workflow engine delivers results that no single-vendor solution, no manual process, and no static automation can match.

The enterprise change management teams that adopt this architecture will not just move faster — they will move safer, with every change verified against real infrastructure, every risk grounded in real data, and every CAB submission backed by evidence generated, tested, and assembled by a governed team of specialist AI agents.

About This Whitepaper Part 2 of 2. Based on a composable solution stack jointly developed by Outshift by Cisco and ServiceNow. AGNTCY is an open-source Linux Foundation project (agntcy.org). ServiceNow AI Control Tower and ServiceNow AI Agents are commercially available on the ServiceNow AI Platform. For architecture details, integration patterns, and deployment guidance, contact the respective teams.

outshift by cisco

outshift.com

@ outshiftbycisco

CREA-1531

V.2 – April 2026