Accelerating Federal Authorization Through Agentic Assessment

Technical White Paper

Darren Death

July 2026

Contents

Infographic

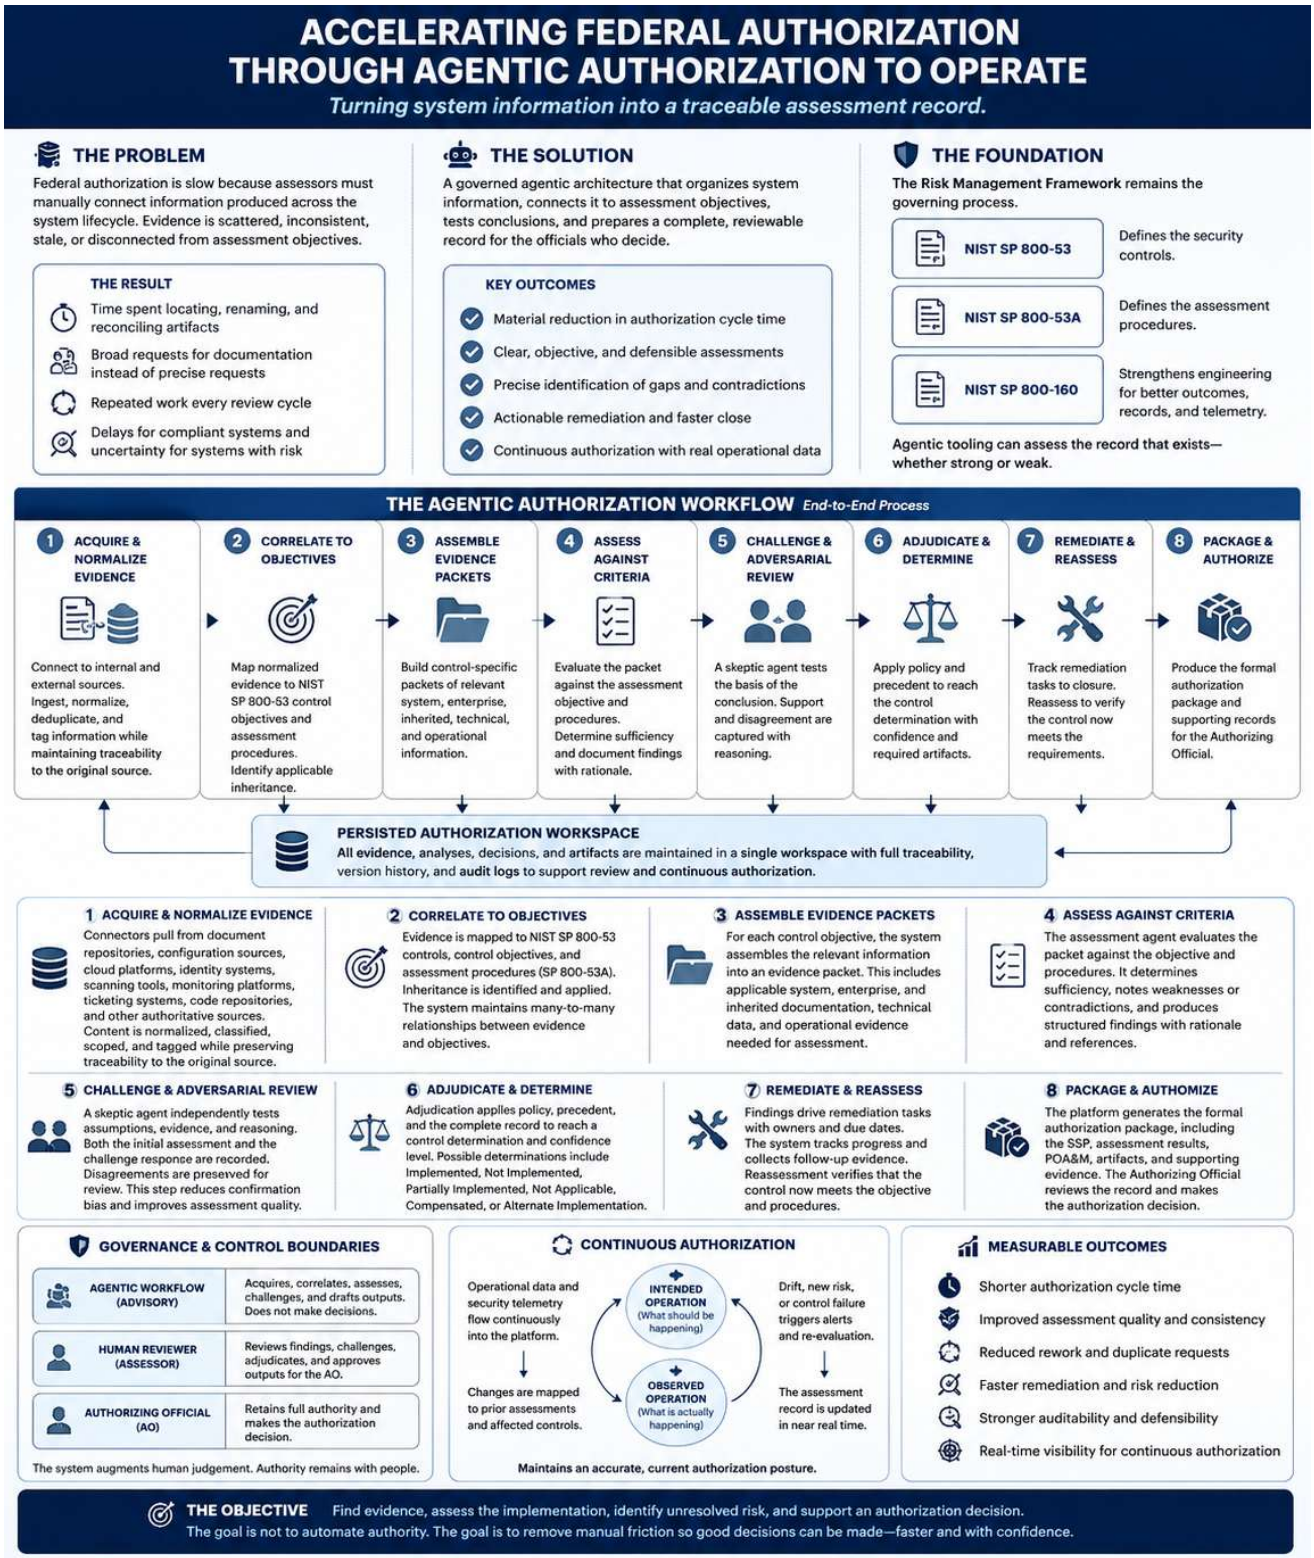

Executive Summary

Agentic workflows can materially reduce the time required to prepare and assess federal authorization packages by turning large volumes of system information into a traceable assessment record. The approach described in this paper has been implemented and tested across the full authorization workflow. The results show that a governed agentic system can perform much of the repetitive correlation and assessment work that has historically consumed assessor and system-owner time.

Federal systems generate large volumes of technical, operational, and governance information that support authorization. That information is often incomplete, inconsistent, stale, or disconnected from the applicable assessment objectives. Assessors and technical points of contact must determine what is relevant, identify what is missing, resolve conflicts, and establish whether the available material supports a defensible conclusion.

A governed agentic architecture can organize system information, connect it to assessment objectives, identify weak or conflicting support, test preliminary conclusions, and prepare the record for review and remediation. The system does not make the authorization decision. It prepares a complete, reviewable record for the officials who do.

The Risk Management Framework remains the governing process. NIST SP 800-53 defines the controls, NIST SP 800-53A defines the assessment procedures, and NIST SP 800-160 supports the engineering practices that produce stronger security outcomes, clearer implementation records, and more useful operational telemetry. Agentic tooling can assess the record that exists, whether that record is strong or weak. Better engineering improves the quality of the available information and reduces ambiguity, but the workflow can still identify gaps, contradictions, and unsupported claims in immature environments.

A system that satisfies the applicable requirements should not remain delayed because authorization has become an extended exercise in locating, sorting, reconciling, and rewriting information. A system that does not satisfy those requirements should produce clear findings, well-defined evidence gaps, and actionable remediation rather than repeated cycles of documentation requests and reassessment. The objective is to find evidence, assess the implementation, identify unresolved risk, and support an authorization decision.

Basis of the Paper

This paper is based on implementation experience rather than a proposed concept. A working agentic authorization platform was built and exercised across the complete authorization workflow. Testing covered the movement of system information from acquisition through assessment, review, remediation, reassessment, and package output while maintaining traceability to the original source. The observed system behavior informed the operating model and control boundaries described throughout this paper.

The implementation demonstrated that agentic workflows could reduce authorization cycle time by correlating technical, operational, and governance information into a reviewable assessment record. The following sections describe the operating model, system behaviors, and control boundaries needed to realize that capability while preserving RMF, assessment rigor, and authorization authority.

1. The Persistent Authorization Problem

Federal authorization has been slow for as long as agencies have relied on manual control assessment and package assembly. The assessor must connect information produced across the system lifecycle, but those relationships are rarely captured in one place.

The same control may be supported by policy, implementation records, technical configuration, and operational data. Some evidence may be inherited, limited to part of the system, or based on an earlier

architecture. An implementation statement may describe a control as operating while current technical evidence shows that it is disabled or only partially deployed. The assessor must resolve those conditions before reaching a conclusion.

Authorization therefore becomes a correlation problem. Collecting documentation is only the beginning. The available information must be connected to the applicable assessment objectives, evaluated for sufficiency, tested for contradiction, and reduced to a defensible control determination. That process has traditionally depended on experienced people manually reading large amounts of material and preserving the relationships across workpapers and GRC systems.

Document-management platforms improve storage and retrieval, but they do not perform the assessment. Search can locate a policy or a configuration export. It does not determine whether the material proves that a control is implemented, whether the evidence is current, whether another artifact contradicts it, or whether the system relies on an inherited capability that is no longer available. A repository can preserve files without creating an assessment record.

The operational result is predictable. Teams spend time locating artifacts, renaming files, managing versions, reconciling evidence requests, rewriting implementation statements, and rebuilding context for every review cycle. Assessors repeat work that was already performed elsewhere because the evidence was not connected to the assessment objective in a reusable way. System owners receive broad requests for documentation instead of precise requests for the missing technical fact.

2. Why Agentic Tooling Can Produce Material Improvement

Large language models alone do not solve this problem. A model can summarize a document or draft an implementation statement, but authorization requires sustained work across a large and changing body of information. The system must understand the purpose of each task, control what information is in scope, preserve the assessment record, and route decisions to the appropriate human reviewer. That requires agentic architecture rather than a single prompt.

An agentic workflow separates the assessment into controlled steps. Source material is prepared, relevant evidence is identified, control-specific context is assembled, and the available information is evaluated against the applicable objective. Preliminary conclusions can then be challenged against conflicting, stale, or incomplete evidence before application logic applies the governing rules and records the result. The reviewer remains able to inspect the evidence, reasoning, challenge, and final status before accepting or changing the outcome.

The improvement comes from maintaining that work as a connected assessment process. The system can perform repeated correlation tasks without requiring the assessor to search for the same repositories or rebuild the same context for each control. Evidence, findings, gaps, remediation activity, and package outputs remain linked throughout the authorization lifecycle. The value is not faster narrative generation. The value is preserving the relationship between the system record and the assessment decision.

The workflow can also surface issues earlier. Weak or contradictory evidence may be identified before final review by comparing implementation records, technical evidence, and operational data against the assessment objective. When inconsistencies are found, the system preserves them for review instead of resolving them automatically. The assessor begins with a structured record of the relevant evidence, the unresolved issue, and the basis for the preliminary conclusion.

The same approach remains useful when the system record is poor. Incomplete evidence becomes a defined gap. Conflicting records are retained and routed for review. Unsupported implementation claims are separated from the technical or operational proof needed to validate them. The workflow does not create

compliance where it does not exist. It reduces the time required to determine what is supported, what is missing, and what requires remediation.

3. A Governed Agentic Authorization Architecture

A credible authorization platform needs clear boundaries between what the model analyzes, what the application controls, and what the reviewer decides. The model can help interpret evidence and prepare assessment results, but it should not be able to change scope, alter policy, remove evidence, or make the authorization decision. Those limits need to be built into the application.

The platform also needs to preserve the assessment as a complete record. The source material, analysis, reviewer actions, and resulting outputs should remain connected so the system can show how a conclusion was reached and what changed during review. That record should persist independently of any individual model response or browser session.

The underlying services should support that workflow without exposing the reviewer to the technical complexity. Long-running ingestion and assessment work should continue in the background, while the application manages access, policy, workflow state, and audit history. Search and retrieval should remain tied to the original source material so evidence can always be verified.

Model use should also be specific to the task. Evidence classification, objective assessment, challenge review, and remediation drafting have different purposes and should use different instructions, validation rules, and output structures. Results that do not meet the required structure should be rejected rather than absorbed into the assessment record.

The operating model is straightforward. Models assist with analysis. The application controls the process. Reviewers determine what is accepted, challenged, or changed. The Authorizing Official retains responsibility for the final authorization decision.

4. Evidence Acquisition, Normalization, and Provenance

The first operational challenge is turning system information into material that can be assessed without losing its original context. That information may come from documents, technical records, system exports, or operational telemetry. Each source has to be interpreted differently, and the surrounding context is often what gives the evidence meaning.

The ingestion process should preserve the identity and origin of every source. It should record where the material came from, when it was collected, which version was used, and what part of the system it applies to. It should also retain enough structure to preserve the meaning of the information. This allows each assessment conclusion to remain tied to the exact source material that informed it.

The platform should then identify which material may be relevant to the assessment. Keyword matching alone is not sufficient because technical evidence rarely uses the same language as the control it supports. A configuration setting, ticket, or log record may demonstrate implementation without naming the control. The system should identify likely relevance while leaving the final sufficiency determination to the assessment process.

Relevant material must be presented with enough context to be understood. A single line, cell, or event is often too narrow to support a conclusion. The surrounding section, table, transaction, or sequence of events may be necessary to show what occurred and whether it applies to the assessed system. Repeated material can be consolidated, but independent evidence from separate sources should remain visible when it strengthens or contradicts the record.

The platform should also distinguish between stated requirements, implementation evidence, and observed operation. Those sources do not prove the same thing, and the assessor should be able to see what type of support underlies each conclusion.

Traceability must be maintained throughout the process. Each finding should remain connected to the source evidence used to support it, and each reviewer's action should show who made the change and why. Without that connection, automation may accelerate the production of conclusions without improving their reliability.

5. Correlating Evidence to Assessment Objectives

The key technical function is connecting the available information to the specific assessment objective being reviewed. A control is often too broad to assess as a single block of text, so the workflow should break it into the individual objectives that must be satisfied and define what each objective requires.

The application should then assemble the information relevant to the objective from approved system, enterprise, inherited, and operational sources. The evidence set should be bounded. The model should not search for an uncontrolled corpus or decide on its own what information is allowed to influence the result.

The assessment task should return a structured finding rather than a free-form answer. It should show what supports the objective, what conflicts with it, what is missing, and what uncertainty remains. The reasoning should be specific enough for a reviewer to determine whether the conclusion follows from the cited material.

Correlation must also account for scope. Evidence that proves a control on one component does not automatically prove the control across the system boundary. A centralized capability may satisfy an inherited requirement only if the system is actually connected to and covered by that capability. A procedure may exist but not apply to the production environment. A scan may be current but exclude a critical subnet. The platform must preserve these distinctions rather than collapsing all positive evidence into a simple score.

Conflicting information must remain visible. The SSP may describe one implementation while technical records show another. Policy may require an activity that operational records do not support. A remediation ticket may be closed while the underlying vulnerability remains open. The workflow should surface those conflicts and require them to be resolved or explicitly accepted before a final determination is made.

The result is a focused assessment record for each objective. The assessor can see the requirement, the supporting evidence, the conflicting information, the remaining gap, and the proposed determination in one place. That is substantially more useful than a general summary of the source documents.

6. Challenge, Adjudication, and Human Review

A strong agentic assessment process should challenge its own conclusions. The initial reasoning pass may overvalue a policy statement, miss a scope limitation, accept stale evidence, or fail to recognize that corroborating material is absent. A separate challenge process should evaluate the preliminary result and test whether the evidence actually supports the proposed determination.

The challenge should test the specific basis for the proposed determination. It should look for unsupported assumptions, conflicting records, incomplete implementation, weak traceability, or evidence that does not apply to the assessed system. It should not automatically change the result. It should preserve the disagreement for the reviewer.

Final control status should be computed through explicit policy logic. The model can explain what the evidence appears to show, but the application should apply the approved assessment rules and determine the resulting status.

Human authority must remain visible. Reviewers should be able to inspect the source material, compare the analysis and challenge, request additional evidence, and change the computed status with a documented rationale. The system should preserve those actions as part of the assessment record.

This approach does not reduce assessor judgment. It gives the assessor a better starting point. Instead of opening a repository of unstructured files, the assessor receives a control-specific record showing the requirement, the supporting evidence, the conflicts, the gaps, the preliminary analysis, and the system-computed status. The assessor can then focus on whether the implementation satisfies the requirement and whether the remaining risk is acceptable.

7. Remediation and Reassessment

Authorization work does not end when a control is found deficient. The same agentic workflow should carry the finding into remediation. A gap should remain connected to the objective, the supporting and missing evidence, the affected components, the responsible owner, and the expected close state. That continuity prevents remediation from becoming a separate document-management process.

The platform should generate targeted closure questions based on the actual deficiency. Instead of asking a system owner for more documentation, it can ask for the missing technical fact, configuration state, approval, test result, or operational record. The responses can be captured as structured closure material and linked to the original finding.

The system can then describe what a passing state must demonstrate. It can identify the required implementation change, the evidence needed to verify it, the expected operational telemetry, and the criteria that will be used during reassessment. Draft procedures, implementation statements, test plans, or evidence templates may be generated, but they remain drafts until the owner validates and operationalizes them.

When new evidence is submitted, the same objective can be reassessed against the same criteria. The reviewer can see the original findings, the remediation action, the new evidence, the updated analysis, and the final disposition in one record. That creates a direct path from deficiency to closure without rebuilding the assessment from the beginning.

The same approach can support later package updates and continuous monitoring. Existing evidence can be reused, and changes to the system can be routed to the controls they affect. Reassessment then focuses on the changed condition instead of reconstructing the entire authorization package.

8. Relationship with RMF

Agentic tooling fits inside RMF. It does not create a parallel authorization process. The framework still defines the sequence of categorization, control selection, implementation, assessment, authorization, and monitoring. The agentic platform accelerates the work performed within those activities by reducing the time required to understand the system record and prepare reviewable outputs.

The same workflow can support the entire authorization lifecycle. As a system is implemented, engineering and operational information can be correlated to the applicable controls and assessment objectives. During assessment, the platform can organize the available evidence, identify gaps or contradictions, and prepare reviewable findings. As the authorization package is developed, it can maintain a complete assessment record that allows reviewers and decision officials to understand the implementation, the supporting

evidence, and any remaining risk. When the system changes, the same process can identify which prior conclusions should be revisited instead of requiring a complete reassessment.

NIST SP 800-53 and SP 800-53A remain authoritative for controls and assessment objectives. The agentic workflow does not rewrite those requirements. It converts them into executable criteria packages and applies them consistently across the evidence set. NIST SP 800-37 remains the process structure. The technology improves execution speed and consistency without changing accountability.

The platform should produce formal authorization outputs from the same assessment record. Keeping findings, remediation, implementation information, and package content connected reduces inconsistencies between what was assessed, what was corrected, and what is ultimately presented for authorization.

9. The Role of NIST SP 800-160

NIST SP 800-160 is relevant because system security engineering affects both the security of the system and the quality of the record available for assessment. It is not a documentation standard. It describes disciplined engineering practices that integrate security throughout the system lifecycle.

An agentic assessment platform benefits from that discipline because there is more reliable information to correlate and evaluate. It can connect security requirements to implementation, implementation to technical evidence, and technical evidence to operational behavior. The result is a more complete and consistent assessment record.

The platform does not depend on mature engineering to operate. It can assess incomplete documentation, inconsistent records, limited telemetry, and unsupported implementation claims. In those environments it will identify more evidence gaps, more contradictions, and more remediation work because those conditions already exist. The workflow still reduces the effort required to understand the system, but it cannot compensate for weak engineering or poor implementation.

Organizations that apply NIST SP 800-160 practices are therefore likely to realize greater benefit from agentic assessment. Better engineering produces a clearer system record, and a clearer system record is easier to assess. Agentic technology accelerates the assessment process, while disciplined engineering improves the quality of what is being assessed.

10. Continuous Authorization and Operational Reuse

The same architecture used to prepare an initial ATO package can also support continuous authorization. The initial decision is based on an understood system state, defined control implementation, available evidence, and accepted risk. Continuous authorization depends on knowing when those conditions change.

An agentic platform can connect operational changes to the controls and assessment conclusions they affect. The system can then show which prior findings remain valid, which require review, and where new risk may have been introduced.

This approach uses both traditional documentation and live operational data. The documentation describes how the system is intended to operate, while technical and monitoring data show how it is actually operating. Bringing both into the same assessment record gives reviewers a more current view of system risk.

Human review is still required for material changes, disputed evidence, risk acceptance, and authorization decisions. The benefit is that reviewers receive a focused record of what changed and what conclusions may be affected instead of rebuilding the entire package. Continuous authorization becomes a continuation of the same assessment process used for the initial decision.

Conclusion

Federal authorization has always required assessors to connect information across engineering, operations, cybersecurity, and governance. That work has been slow because the relevant information is distributed across many systems and artifacts. The problem is not new. The tools available to address it are.

Agentic workflows can process that information at a scale that was not previously practical. They can maintain the relationship between source evidence, assessment conclusions, reviewer decisions, and remediation throughout the authorization lifecycle. When those capabilities operate inside a governed platform with persistent records and human review, they can reduce authorization package cycle time without weakening RMF or shifting decision authority to a model.

NIST SP 800-160 improves the conditions under which this works. Strong system security engineering produces better implementation, clearer architecture, more useful telemetry, and a stronger record for assessment. The agentic workflow can still assess incomplete or weak material, but better engineering reduces ambiguity and makes the resulting conclusions easier to support.

A compliant system should not remain delayed because people are still sorting files and reconstructing relationships that software can now maintain. The objective is to find the evidence, assess the implementation, identify unresolved risk, and authorize the system when the record supports that decision.

References

National Institute of Standards and Technology. NIST Special Publication 800-37 Revision 2, Risk Management Framework for Information Systems and Organizations.

National Institute of Standards and Technology. NIST Special Publication 800-53 Revision 5, Security and Privacy Controls for Information Systems and Organizations.

National Institute of Standards and Technology. NIST Special Publication 800-53A Revision 5, Assessing Security and Privacy Controls in Information Systems and Organizations.

National Institute of Standards and Technology. NIST Special Publication 800-160 Volume 1 Revision 1, Engineering Trustworthy Secure Systems.