RESEARCH PAPER

A Security Claim Should Carry Its Own Receipts

A security claim is reusable only when its scope, source, date, responsibility, conflicts, and reviewer stay attached.

STATUS	LAST REVIEWED	RESEARCH LANES	CLAIM REGISTER
Current	**August 13, 2026**	**Evidence · Evaluation · Mechanism**	**V-02 · V-03 · V-04 · V-05 · C-05 · C-06 · C-07 · C-08 · C-11**

Current design note for evidence-profile version 0.1 and its synthetic package. The profile is a proposal and working implementation, not a Government standard or authorized assessment method.

"MFA is enabled" is not enough information for a CMMC reviewer.

For which accounts? In which system? Observed when? From which export? Does the export include local emergency access? Who checked it? What happens when a new administrator is added?

Remove those facts and the sentence is easy to copy, easy to misunderstand, and hard to test. Keep them attached and the claim can move from the contractor to an advisor, provider, assessor, or Government reviewer without losing the reason anyone believed it.

That is what we mean when we say a security claim should carry its own receipts.

Documents are outputs

SSPs, assessment workbooks, provider statements, executive summaries, and machine-readable files are all useful. They serve different readers.

The problem begins when each becomes a separately rebuilt source of truth. One fact is described in the SSP, copied into an objective worksheet, summarized in a readiness dashboard, and entered again into an assessor's package. A later change can leave four different answers behind.

Instead, keep one record behind the claim and generate the views people need. The document remains important. It no longer has to carry the whole truth alone.

What has to stay with the claim

INFORMATION ANOTHER REVIEWER NEEDS

PART	QUESTION
Claim	What is being asserted, and against which requirement or objective?
Implementation	Which people, process, technology, and configuration make it true?
Source	Which export, policy, ticket, test, interview, or observation supports it?
Scope	Which system, service, asset, account population, and time period does it cover?
Responsibility	What does the contractor do, what does a provider do, and what remains open?
Date	When was the evidence collected, when should it be checked again, and which changes make it stale?
Conflict	What evidence disagrees, is missing, or does not cover the full population?
Review	What did the software flag, what did the reviewer find, and has any official decision occurred?

Those fields do not make the claim true. They make it possible for someone else to check.

Keep the contradiction

Our public synthetic example starts with a claim that MFA covers every privileged account. An identity export reports 24 covered accounts. An exception record reveals an emergency access path outside the export's population.

A weak workflow accepts "24 of 24," rounds the claim up to passing, or buries the exception in a note. A useful workflow keeps the claim, export, population limit, exception, software warning, and human finding separately visible.

CLAIM **3.5.3**

UNCERTAIN

Multifactor authentication is enforced for synthetic privileged and network access paths in scope.

OBSERVED POPULATION

24 / 24

synthetic accounts visible to the identity-provider export

SOURCE

Synthetic Identity Provider

read only api · 2026-08-10T12:00:00Z

FRESHNESS

7 Days

Next review 2026-08-17T12:00:00Z

INHERITED RESPONSIBILITY

Shared

syn-msp-01 · valid until 2026-09-01T00:00:00Z

SHA-256 d489b55bd14ee968acb329fe6dedf82f8a821d3a52416a249da7037cf1491abd

Synthetic specimen. No customer, contractor, CUI, FCI, assessment, or production-system data. Inspect the complete walkthrough or download the evidence package.

The contradiction is not a blemish to clean from the file. It tells the reviewer where routine checking stops and judgment begins.

What software can do

Software can read exports, check required fields, compare dates and populations, find broken references, map evidence to objectives, and call out conflicts. AI can suggest a mapping or draft a possible finding and cite the records it used.

A qualified reviewer still decides whether the evidence is sufficient, whether the scope is right, whether a compensating condition matters, and what the finding should be. Authorized officials retain certification, affirmation, contracting, acceptance, and enforcement decisions.

The record should make those handoffs obvious. A downstream reader should never have to guess whether a sentence came from a source, a rule, software, an assessor, or a Government decision.

Can another tool use it?

Calling a format “vendor-neutral” proves nothing. It has to pass practical tests:

- Can two independent tools represent the same claim without losing its scope, source, or responsibility split?
- Can another tool validate the package without access to private product state?
- Can a reviewer distinguish missing evidence from evidence that shows failure?
- When a system changes, can the record identify which claims and findings must be checked again?
- Can a contractor export the record without giving the software authority to make a CMMC decision?

Deep Fathom has published a versioned schema, rule-based validator, synthetic package, and draft OSCAL mapping so these questions can be tested against working files. That is evidence of an implementation, not proof that the format is complete or that Government should adopt it.

The next test

Give ordinary review and AI-assisted review the same synthetic cases. Have qualified reviewers establish the reference findings before they see system output. Count missed evidence, unsupported conclusions, conflicts found, corrections, reviewer time, and disagreements. Publish the failures as well as the wins. Do not let any result affect a contractor.

If another tool cannot read the record, fix the format. If reviewers need facts the record does not carry, add them. If the boundary between software and human judgment is unclear, stop.

Inspect the worked evidence record, follow the claim through the software, or help vet the Government pilot.

TRACE THE CLAIM

Sources and revision history.

PRIMARY AND GOVERNING SOURCES

01 **CMMC Assessment Guide — Level 2, Version 2.13**
Department of War Chief Information Officer
↗

02 **Open Security Controls Assessment Language**
National Institute of Standards and Technology
↗

03 **OSCAL Frequently Asked Questions**
National Institute of Standards and Technology
↗

04 **Cybersecurity Maturity Model Certification Program, 32 CFR part 170**
Electronic Code of Federal Regulations
↗

05 **CMMC 20X Common Evidence Profile**
Deep Fathom
↗

CORRECTIONS AND MATERIAL REVISIONS

No material corrections recorded.

See an error or a source we missed? Send a correction. Material changes are recorded here rather than silently overwritten.