



AF-HB-9500-002

Model Assurance and AI-Assisted Engineering Handbook

Space / Aviation / Defense

FIRST EDITION

AF 9500 SUPPORTING DOCUMENT

Document identity and use

FIRST EDITION MANUSCRIPT - Prepared for Aerospace Federation review

Field	Controlled identity
Document	AF-HB-9500-002 - Model Assurance and AI-Assisted Engineering Handbook
Parent standard	AF 9500 - Aerospace Digital Systems, Software & Cybersecurity
Edition / language	FIRST EDITION / English master text
Family	Manuals
Primary source anchors	14-16, 20, 27-28, 36; Annexes F, K-L
Suggested local owner	Model and assisted-engineering owner
Authority	Publisher approval and effective date are not assigned. Local use requires authorized adoption and completed parameters.

Use with the main standard

This document implements the identified AF 9500 clauses and annexes. Read the complete applicable requirements in the main standard. The source anchors are navigation references; a template does not replace technical judgement, evidence, permission or eligible approval.

AF-ML-9500-001 lists the complete supporting set. AF-CM-9500-001 maps the main clauses and annexes to working documents. AF-TP-9500-001 establishes the common local parameter baseline; add the relevant discipline profiles.

Record control

Identify the object, configuration, record owner and dates. Reference attachments by stable identifier and version. Keep superseded evidence and decisions traceable. Use the approved local retention, access and authentication arrangements. No technical limit, numerical limit, interval or acceptance value is assigned by a blank template.

Implementation model: Translate the stated responsibilities and controls into the actual organization. Informative examples remain explanatory; use approved programme-specific values for operational decisions.

1 Distinguish the model from the surrounding function

A trained model is only one part of an operational system. Define the decision, data pipeline, preprocessing, thresholds, interfaces, supervision and fallback. Use AF-FR-9500-020 and AF-TP-9500-014 to state intended domain, prohibited use and consequence. Do not infer aerospace suitability from a generic benchmark.

1.1 Identity and provenance

Preserve model architecture, weights, training configuration, data lineage, dependencies and deployed settings. A model name or version label may not identify actual weights or preprocessing. Maintain AF-RG-9500-012 and connect model changes to configuration and release records.

1.2 Data and independent evaluation

Record source, rights, labeling, transformations, exclusions and populations. Split correlated entities and time periods appropriately for the evaluation purpose. Prevent tuning from silently consuming the independent assessment set. Synthetic data can support development but requires an argument for representativeness when used as evidence.

2 Evaluate consequential behavior

Examine critical subpopulations, rare conditions, shift, uncertainty and corrupted inputs. Use clear denominators and distinguish measured behavior from untested assumptions. AF-TR-9500-001 illustrates why aggregate accuracy and zero observed failures can mislead when interpreted without their models.

2.1 Safeguards and supervision

Demonstrate detection or bounding of out-of-domain behavior and a verified fallback. Human oversight requires information, competence, time and actual authority to intervene. Contact delay and workload can make nominal supervision ineffective.

2.2 Change and drift

Control retraining, weights, thresholds, preprocessing and dependencies. Define drift signals, owners and restricted states. Online adaptation shall stay within an approved and monitored strategy. Reevaluate changed behavior and preserve historical claims and deployed effectivity.

3 Control AI-assisted engineering

Use AF-FR-9500-021 and AF-TP-9500-015 for generative assistants and agents. Define permitted uses, data exposure, tool privileges and accountable reviewers. Generated requirements, code, tests and reports remain unverified until the necessary checks are complete.

3.1 Independent evidence

Check outputs against authoritative requirements and observations. Generated tests can share the implementation's mistaken assumption; separate evidence is needed where that risk matters. Review rights of use, hidden dependencies, fabricated references and unsupported confidence claims.

3.2 Untrusted instructions and bounded actions

Treat retrieved content and external tool results as information whose authority is determined by the controlled workflow. They shall not grant access or override requirements. Define approval points, action limits, logging, interruption and recovery from partial execution. Test prohibited actions in a controlled environment.

4 Release only the supported use

Use AF-PR-9500-028 to connect exact model or workflow identity to valid evidence and restrictions. A model confidence value, generated explanation or agent success message is not technical approval. Keep actual decision authority, source evidence and changed assumptions visible through operation and retirement.

Required evidence: Reproducible model/data identity, independent consequential evaluation, verified safeguards and bounded assisted or agentic workflows.