



AF-TR-9500-001

Quantitative Methods and Worked Calculations

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-TR-9500-001 - Quantitative Methods and Worked Calculations
Parent standard	AF 9500 - Aerospace Digital Systems, Software & Cybersecurity
Edition / language	FIRST EDITION / English master text
Family	Companions
Primary source anchors	7, 16-17, 23, 27, 30-34; informative Annex O
Suggested local owner	Systems / statistics / assurance specialist
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 Keep superseded evidence and decisions traceable.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 Use of quantitative examples

All values and cases in this companion are synthetic. They explain assumptions, calculations and limitations; they are not system limits, product evidence or mandatory default parameters. Select and approve the actual model and values before relying on a calculation.

2 Worked calculations

Status: Informative. All values are synthetic and are not mandatory thresholds, actual product results or evidence of operational approval. The method, assumptions and limitations are part of each example.

0.1 Traceability, disposition and denominator control

A baseline contains 120 requirements: 100 PASS, 8 FAIL, 10 HOLD and 2 justified N/A. All 120 have trace rows. For this example, a verified disposition means a valid PASS or FAIL result; HOLD remains unresolved. Different organizations may define another metric, but shall state its meaning.

Applicable requirements = $120 - 2 = 118$

PASS fraction = $100 / 118 = 84.7458\%$

Valid PASS-or-FAIL dispositions = $108 / 118 = 91.5254\%$

Trace-row presence = $120 / 120 = 100\%$

A complete matrix therefore coexists with eight failures and ten unresolved requirements. None of these percentages permits release where an essential condition remains unresolved. Removing HOLD items from the denominator would produce a misleading success metric unless the changed population and purpose were explicitly justified.

0.2 End-to-end latency budget

A synthetic control path has approved conservative bounds of 8 ms acquisition, 12 ms computation, 4 ms authentication, 9 ms communication and 15 ms actuation. A separate non-overlapping allowance of 5 ms covers relevant jitter and interference. Assume the stages are sequential and these bounds are valid for the same operating conditions.

Total bound = $8 + 12 + 4 + 9 + 15 + 5 = 53$ ms

Margin to a 60 ms deadline = $60 - 53 = 7$ ms

Margin / deadline = $7 / 60 = 11.6667\%$

This is a deterministic budget calculation, not a measured percentile or probability of meeting the deadline. Parallel stages, correlated delays, omitted queues or double-counted allowances require a different model. A mean benchmark cannot replace the stated conservative bounds. Adding a 9 ms security operation would exceed the deadline by 2 ms under these assumptions.

0.3 Serial-service availability

Assume a service requires both application A and identity service B continuously. Their steady-state availabilities are 0.999 and 0.9995, and their availability states are independent. No offline fallback is credited. For this deliberately simplified model:

Series availability = $0.999 \times 0.9995 = 0.9985005$

30-day period = $30 \times 24 \times 60 = 43,200$ minutes

Expected unavailability = $43,200 \times (1 - 0.9985005)$

Expected unavailability = 64.7784 minutes

A 99.95% target would permit 21.6 minutes over this period under the same accounting definition. The modeled series service does not meet that target. Shared outages, nonstationary behavior, planned downtime rules and actual recovery performance can invalidate the simple model. A contractual availability promise is not measured system evidence.

0.4 Detector precision and base-rate effects

A representative evaluation contains 100,000 observation windows, including 1,000 true adverse windows and 99,000 ordinary windows. A detector catches 90% of adverse windows and incorrectly alerts on 1% of ordinary windows. Assume labels are valid and each window uses the same declared rule.

True positives = $1,000 \times 0.90 = 900$

False negatives = $1,000 - 900 = 100$

False positives = $99,000 \times 0.01 = 990$

Precision = $900 / (900 + 990) = 47.6190\%$

Recall = $900 / (900 + 100) = 90.0000\%$

More than half of alerts in this example are false positives despite the 1% false-positive rate. Response staffing and operating effects therefore need explicit evaluation. Results cannot be transferred to another deployment without examining prevalence, workload, labels and distribution. A detector that misses a critical scenario cannot be accepted solely on aggregate precision.

0.5 Recovery time and recoverable state

A service fails at 10:12. Its latest independently verified consistent checkpoint is 10:00 and there is no usable later transaction journal. Restoration takes 80 minutes, functional and security validation takes 25, and state reconciliation takes 20. These activities are sequential. The locally approved objectives are a 120-minute return-to-service interval from failure and no more than 10 minutes of lost committed state.

Actual recovery interval = $80 + 25 + 20 = 125$ minutes

Recoverable-state gap = $10:12 - 10:00 = 12$ minutes

The time objective is exceeded by 5 minutes and the recoverable-state objective by 2 minutes. A report that counts only the 80-minute file restore would misstate the approved objective. If journals exist, their completeness and replay consistency must be demonstrated before crediting a newer recovery point.

0.6 Event ordering with bounded clock uncertainty

Event A is recorded at 100.000 s with an absolute clock-error bound of 0.004 s. Event B is recorded at 100.006 s with a bound of 0.005 s. Assume the bounds cover the relevant capture times and there is no other ordering evidence.

A actual-time interval = [99.996, 100.004] seconds

B actual-time interval = [100.001, 100.011] seconds

B minus A interval = [-0.003, +0.015] seconds

The intervals overlap. B may have occurred up to 3 ms before A or up to 15 ms after it, so their order is not established by these timestamps alone. These are hard bounds, not standard deviations; combining them by root-sum-square would introduce an unsupported statistical interpretation. Sequence or causal evidence may resolve the order separately.

0.7 Zero observed failures and a statistical bound

In 300 independent Bernoulli trials from a fixed representative population, no failures are observed. Assume constant failure probability and a valid detection rule. The exact one-sided 95% upper confidence bound for that trial failure probability is:

$$\text{Upper bound} = 1 - 0.05^{(1 / 300)} = 0.0099361$$

The bound is approximately 0.9936%, not zero. It is a frequentist confidence bound for the stated sampling model, not a probability that software is safe or a universal operational failure rate. Correlated tests, systematic defects, incomplete requirements or nonrepresentative inputs invalidate that interpretation. Additional deterministic and scenario-based assurance remains necessary.

0.8 Exposure accounting without hidden populations

Ten identified assets remain exposed to a particular unresolved condition for 3 days, and another 20 remain exposed for 7 days. Assume all intervals use the same start/stop rule and do not overlap for the same asset. A separate set of 4 assets has unknown state and is not included in the computed total.

$$\text{Confirmed exposure} = 10 \times 3 + 20 \times 7 = 170 \text{ asset-days}$$

The report shall also state four assets with unknown exposure. Unknown is not zero. Asset-days can describe backlog trends but cannot make a short catastrophic exposure equivalent to a long minor one. Criticality, actual reachability, compensating controls and deadline breaches remain separate decision inputs.

0.9 Log-storage capacity

Assume a sustained 500 events per second, an average stored payload of 600 bytes per event, 30 days retention, two full copies and a 25% additional capacity allowance. Decimal GB and TB are used. Compression, indexes, metadata and variation are excluded unless covered by the stated allowance.

$$\text{Daily payload} = 500 \times 600 \times 86,400 = 25.92 \text{ GB}$$

$$\text{30-day payload} = 25.92 \times 30 = 777.6 \text{ GB}$$

$$\text{Two copies plus allowance} = 777.6 \times 2 \times 1.25$$

$$\text{Planned capacity} = 1,944 \text{ GB} = 1.944 \text{ TB}$$

This calculation does not prove adequate peak ingestion, search capacity or retention during an incident. Measurements must establish actual event distributions, indexing costs, redundancy and failure behavior. Losing a collector or exhausting storage requires a defined response rather than a silent gap.

0.10 Contact-window update feasibility

A remote target receives a 120 MiB package over a 2 Mbit/s link with 60% effective payload efficiency. One MiB is 1,048,576 bytes; one Mbit is 1,000,000 bits. Validation, activation and confirmation require 90, 60 and 120 seconds sequentially after transfer. Assume continuous usable contact and no retransmission beyond the efficiency allowance.

$$\text{Transfer} = 120 \times 1,048,576 \times 8 / (2,000,000 \times 0.60)$$

$$\text{Transfer} = 838.8608 \text{ seconds}$$

$$\text{Complete sequence} = 838.8608 + 90 + 60 + 120$$

$$\text{Complete sequence} = 1,108.8608 \text{ seconds} = 18.4810 \text{ minutes}$$

A 20-minute window leaves 91.1392 seconds. This nominal margin is not an operational approval: link uncertainty, power, interruption recovery, required reserves and permitted mission phase must also be assessed. Download completion alone does not confirm activation.

0.11 Fleet confirmation and trust migration

A trust update targets 8,000 assets. Of these, 7,400 have confirmed the new active baseline, 200 received the package but have not confirmed activation, and 400 have not been contacted. Assume these populations are mutually exclusive and exhaustive.

$$\text{Package delivery fraction} = 7,600 / 8,000 = 95.0\%$$

Confirmed active fraction = $7,400 / 8,000 = 92.5\%$

Unconfirmed or unreachable = $600 / 8,000 = 7.5\%$

Revoking the old trust path may strand some of the remaining assets. The authority must assess their actual state, fallback and exposure before transition. The 95% delivery figure shall not be presented as confirmed migration. A dormant or unknown target remains part of the affected population until dispositioned.

O.12 Aggregate model accuracy and a critical subpopulation

A synthetic evaluation contains 9,950 ordinary cases and 50 consequential rare cases. The model classifies 9,900 ordinary cases and 30 rare cases correctly. Assume valid labels and a fixed evaluation rule.

Aggregate accuracy = $(9,900 + 30) / 10,000 = 99.30\%$

Ordinary-case accuracy = $9,900 / 9,950 = 99.4975\%$

Rare-case accuracy = $30 / 50 = 60.00\%$

Twenty of the 50 consequential cases fail. The high aggregate accuracy cannot compensate for that result. Confidence intervals, representativeness, safeguards and scenario-specific criteria are also needed. The example does not assign an acceptable performance threshold for any aerospace function.