

Threat scenarios and the controls they require

Endpoint compromise, replay and disrupted links require different controls. These design options have distinct assumptions and evidence limits.



WHAT THE ADVERSARY DOES	WHAT THE ATTACK EXPLOITS	CONTROL TO CONSIDER	WHAT IT STILL WILL NOT DO
A trusted node is turned against you The node is genuinely yours, correctly issued, and sits exactly where it is supposed to sit.	Identity alone may not reveal altered software. A valid identity can remain usable after compromise unless additional controls intervene.	Proposed: evidence before issuance Require accepted hardware evidence before issuing a credential under an agreed policy. RFC 9334 · RFC 9711 · RFC 9999	Issuance checks a particular point in time. Exposure also depends on expiry, revocation, reappraisal and how existing sessions end. Refused renewal does not cancel a valid leaf.
A forward sensor is captured and returned It is recovered, its software replaced, and put back into service still wearing your badge.	A credential may survive a software change. A recovered device needs checks appropriate to its hardware and the suspected compromise.	Measure supported device software DICE and PSA support suitable embedded platforms; not every existing sensor qualifies. TCG DICE · RFC 9783 PSA token	Extracted keys can enable impersonation. Identity, key protection and software evidence must be considered together. Withdrawal cannot reach an isolated receiver immediately.
Reports arrive from a sensor that does not exist Traffic is injected that no real device produced, naming a call sign that exists only on paper.	An unverified source label is not identity. Bind the reporting principal to authenticated credentials; do not trust its label alone.	Bind source identity to authentication A verified channel identifies its endpoint. A gateway needs separate source bindings. keys per unit, never per batch	Per-device keys support individual attribution. A shared batch key weakens that distinction. Gateway attestation alone does not authenticate each physical sensor behind the gateway.
Yesterday's traffic is played back as today's Captured traffic is re-sent, or a session is diverted so the proof describes a different link.	A signature alone does not establish freshness. Reused evidence or reports need explicit replay protection and acceptance rules.	Bind evidence and enforce freshness Bind the evidence to the intended key/session. Use a challenge or defined time-based policy. RFC 9266 channel binding · nonce	Session binding is not observation time. Protected counters need persistent state and reset/rollback rules. They can support ordering, but cannot deliver reports without a link.
The equipment is subverted before it deploys The adversary reaches the kit, or the list of what the kit should look like, before it is ever issued.	Reference values are part of the trust boundary. A corrupted approval list can make unwanted software satisfy an otherwise correct check.	Protect the approved reference values Authenticate policy and software approvals; control updates and withdrawal of old values. CoRIM draft · CoSWID: RFC 9393	Review the authority behind the reference list. Check its integrity through a trusted route. Approval of a software image does not establish that it is free from exploitable defects.
The link is jammed instead of the equipment Disruption can prevent fresh evidence checks and delay delivery of operational information.	Online verification depends on connectivity. Define degraded operation in advance rather than silently dropping assurance requirements.	Design option: carry attestation results The passport model can reduce dependence on an online verifier at the point of use. RFC 9334 passport model	Validity periods and cached policy need limits. Disconnected withdrawal and later catch-up need separate design and testing. This is not a demonstrated anti-jamming capability.
The sensor itself is deceived GPS spoofing, decoys, or patterns built to make automatic recognition report the wrong object.	Untampered software can process false inputs. A valid measurement does not establish that the sensor interpreted the world correctly.	Attestation alone does not resolve this Independent sensing and checks may help. Protecting software supports those controls. this is question 4 · see sheet 2	Other sensing or security measures may detect particular deceptions. Software assurance does not make the observations accurate or guarantee that cross-checks find every error.

Recorded evidence and proposed controls
An independent client has appraised a protected workload and completed a call.
Issuer gating and disconnected attestation remain separate design work.
Per-sensor provenance needs evidence for the complete source-to-report path.

Limits of this evidence
It does not prove sensor accuracy, universal tamper resistance or operation under jamming. Each control must be tested on its intended deployment path.
The controls above are not a claim of a complete fielded defence.