A compromised service gets a bounded job

Request
what the service must accept

→

Process
the identity and capabilities it has

→

Evidence
what the team can prove

Security starts by reducing the actions available after compromise.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

Required action or forbidden action?

The API must read app code and a private database.
Which control should stop it changing its code?

> A read-only application root plus a narrow writable runtime area. Keep the required read path; deny the code mutation and test the denial.

A control is useful only when it preserves the required path and blocks a named abuse case.

Four levers shrink blast radius

Identity
non-root UID; drop capabilities

Filesystem
read-only root; narrow tmpfs

Reachability
publish only the edge

Resources
CPU, memory, PIDs bounded

Identity, filesystem, reachability, and resources constrain different failure modes.

Delivery is a chain of trust boundaries

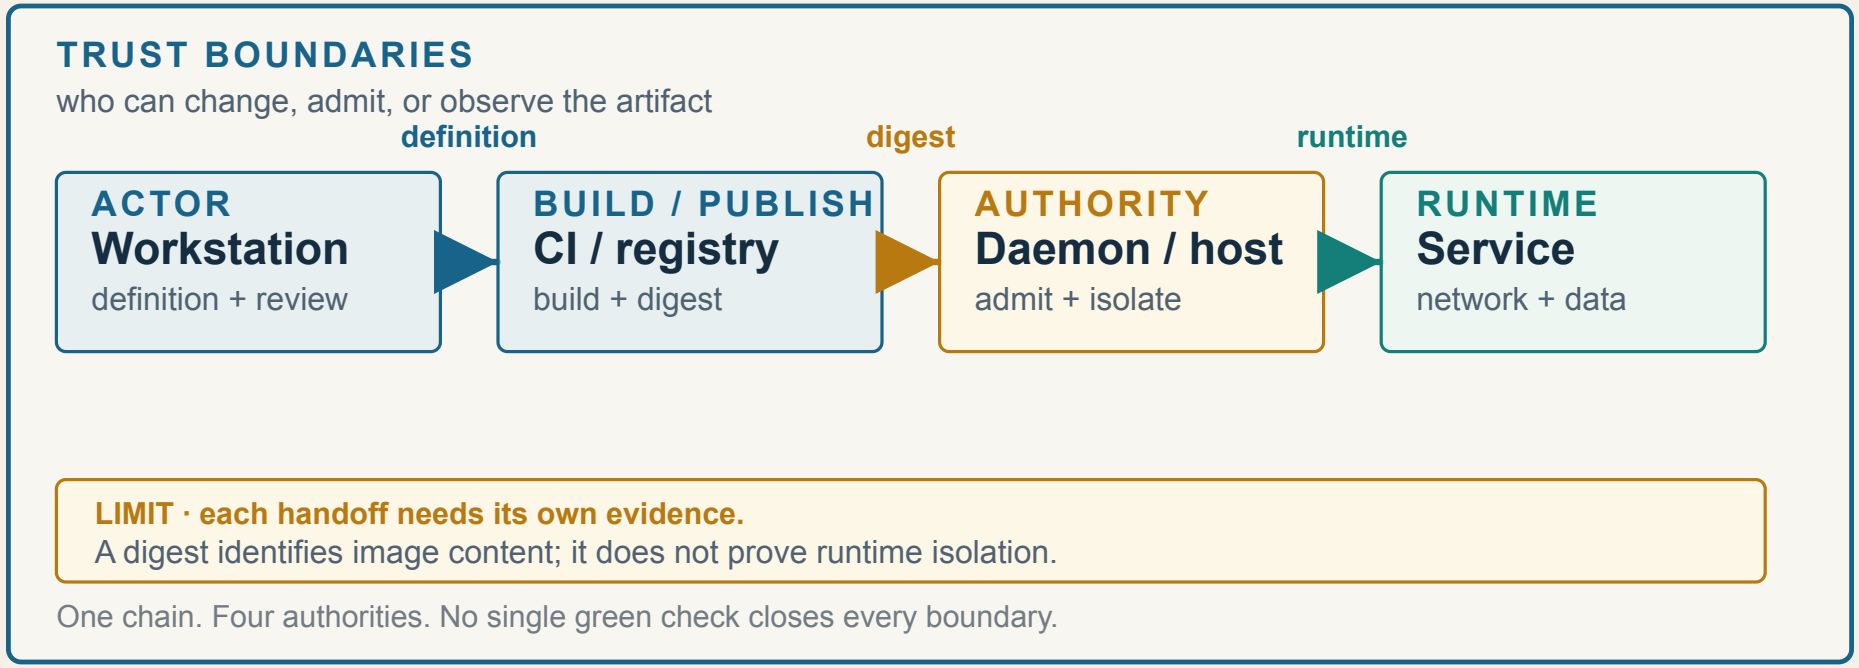

Every handoff changes who can modify, admit, or observe the artifact.

Name the actor, asset, authority

ACTORS

Developer
writes definition

CI
builds artifact

Daemon
creates runtime

ASSETS

Token
registry access

Socket
host authority

Data
application state

A threat model becomes testable when each boundary names an actor, asset, and authority.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

A tag can point somewhere else

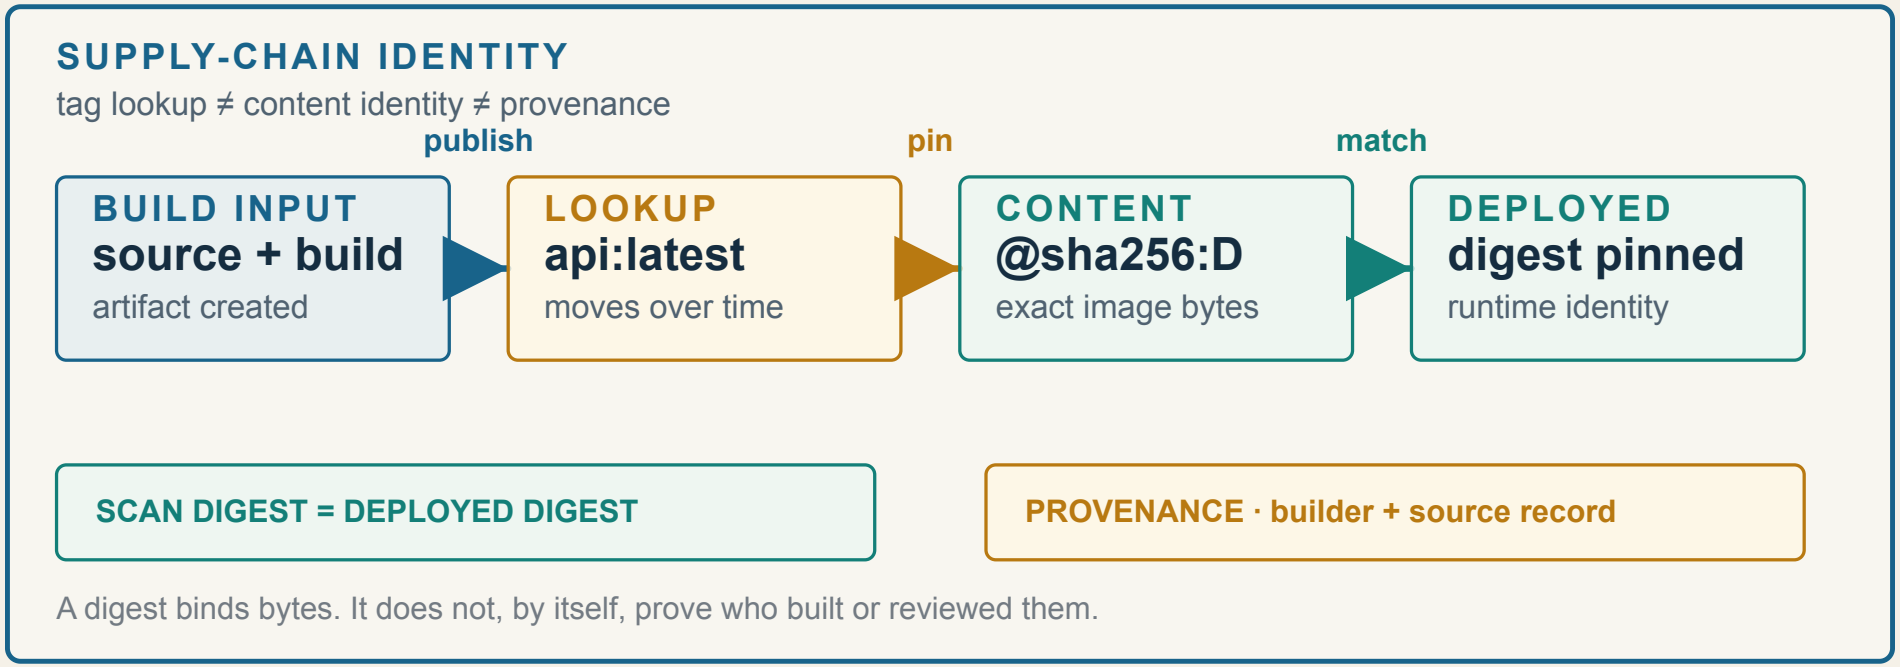

A mutable tag names a lookup; a digest binds the evidence to image content.

One row carries a claim to a test

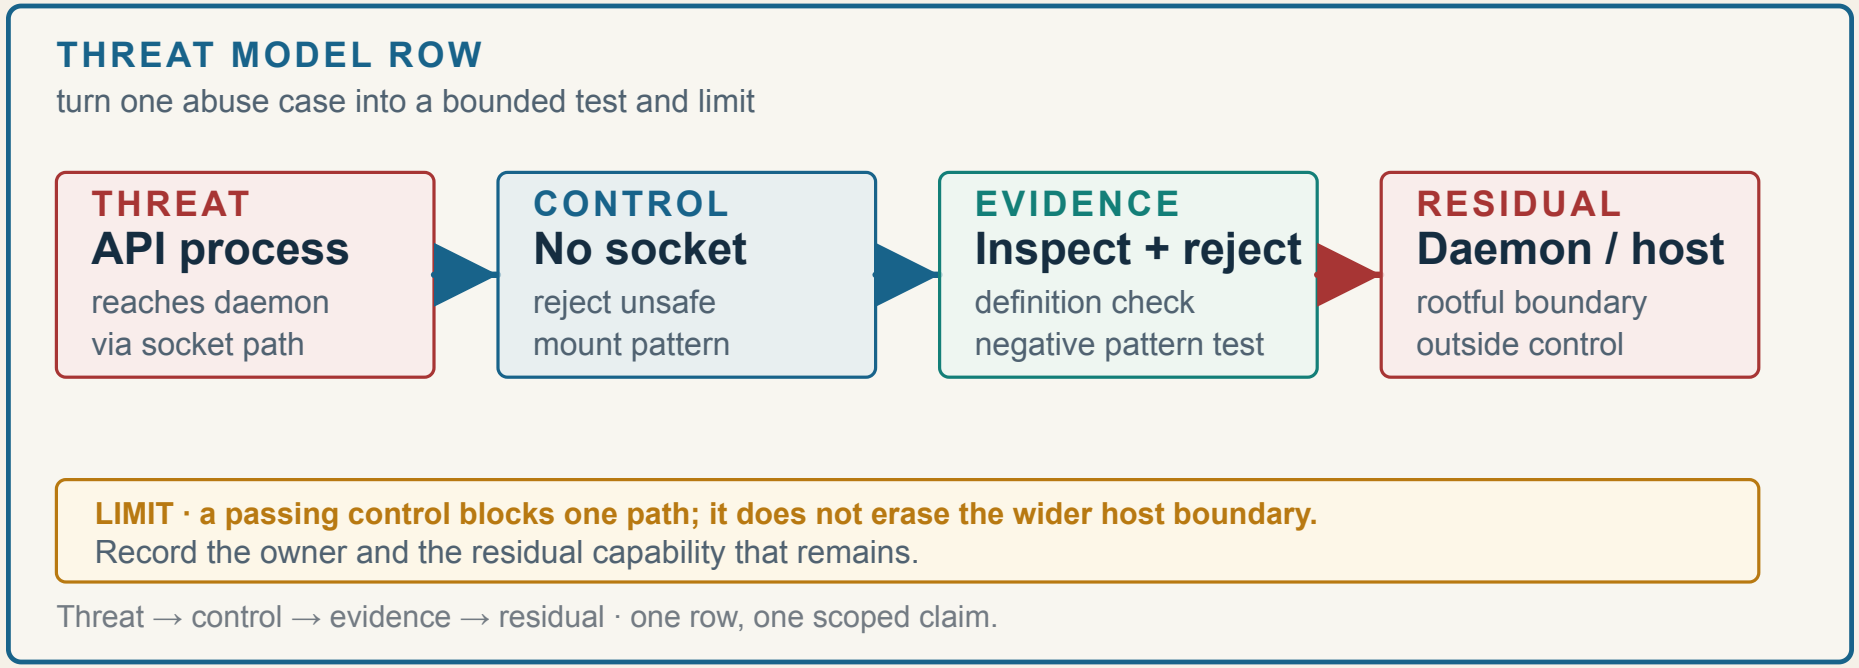

A threat-model row is complete only when control, test, owner, and residual risk are explicit.

Compare threat, control, evidence, residual

COMPARE FOUR THREAT ROWS

Socket → no socket mount → inspect and reject pattern → daemon/host authority remains if that boundary is compromised.

Read-only → RO root + tmpfs → allowed endpoints and denied `/app` write → compromised API can still read allowed paths and use authorized DB access.

Secret → no value in image/logs/history/evidence → path/readiness/leakage checks → a compromised process that reads the secret can still expose it.

Mutable tag → digest pin → scanned digest matches deployment → a compromised build or registry actor can still alter future delivery inputs.

Use the control and evidence together; keep residual attacker capability visible.

Prioritize reachable authority and exposure

Threat
what can the attacker do?

Capability
which asset or path?

Control
which boundary blocks it?

Decision
test, owner, residual

Compare threats through reachability, capability, control evidence, and residual risk.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

The daemon can out-rank the service

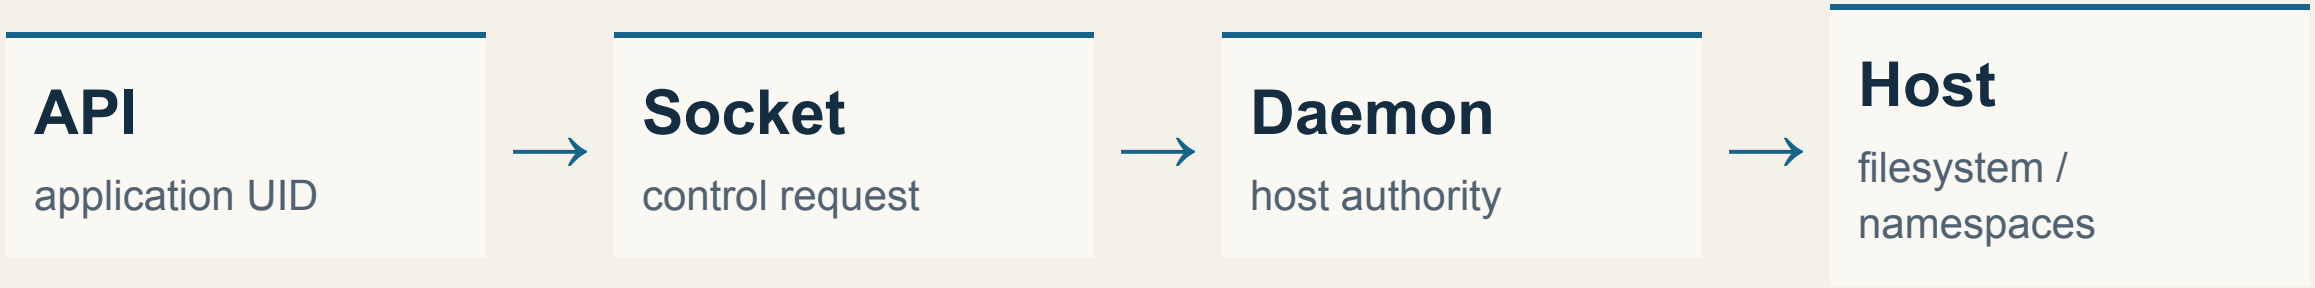

A socket path can turn an application compromise into a host-authority request.

Four shortcuts defeat isolation

Socket
daemon control path

Privileged
broad capability set

Host namespace
shared kernel view

Host mount
host-wide filesystem

Reject socket, privileged, host namespace, device, and broad host mounts before execution.

Non-root is one control, not the boundary

The API runs as UID 10001. What additional check is needed before saying the host is protected?

> Inspect socket, privileged, host namespace, device, and host-root mounts. UID evidence alone does not remove daemon or kernel authority paths.

Identity evidence is necessary and still insufficient when a stronger authority path remains.

Reject the definition before the daemon sees it

Parse
known fields

→

Reject
unsafe authority

→

Review
owner + residual

→

Admit
approved runtime

Admission is a control point: unsafe authority should fail before dynamic execution.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

A bounded admission transcript

ILLUSTRATIVE REFERENCE TRANSCRIPT · contexts labelled; execution deferred to approved TD05 host

```
$ bash demos/admission-reference/run_reference.sh
fixture: cm5-admission-local
input: approved local definition (teaching fixture)
reject: socket mount, privileged, host namespace, host root mount
decision: HOLD for owner review
execution: reference display only
```

A transcript can teach the decision shape without claiming that a host or provider was executed.

What belongs in the admission record?

Choose the minimum record for a rejected definition: raw log, field, reason, owner, residual?

> Record the definition identity, rejected field, reason, reviewer/owner, and residual or next action. Keep evidence narrow and reproducible.

An admission decision is evidence when its input, reason, owner, and next action are preserved.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

Least privilege has separate knobs

UID
10001 for API

→

Caps
drop ALL; justify adds

→

FS
RO root + narrow tmpfs

→

Limits
CPU / memory / PIDs

Least privilege is a set of independent runtime constraints with independent observations.

The database is not the API

API / PROXY

UID
10001 / non-root

Root
read-only

Tmpfs
only required paths

POSTGRESQL

UID
non-root steady state

Data
named writable volume

Residual
exception documented

A security profile must preserve the service contract while documenting deliberate exceptions.

Bound the process and its writable paths

/app
read required; deny mutation

/tmp
narrow tmpfs only

PIDs
bounded process count

Devices
none by default

A narrow writable path and resource bounds turn a policy sentence into observable runtime facts.

Break · return with one residual risk

10 min

Stand up, then return with one unanswered evidence question.

We resume in ten minutes with denied-write diagnosis and secret paths.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

Nonzero is a symptom, not a diagnosis

OBSERVED

write failed
status ≠ 0

one observation

DIAGNOSE

Read-only
rootfs policy

UID
ownership / permission

Path
missing mount / directory

A failed write needs its error, identity, and mount context before it becomes evidence.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

Allowed HTTP, denied app write

ILLUSTRATIVE REFERENCE TRANSCRIPT · contexts labelled; execution deferred to approved TD05 host

```
host$ curl -fsS http://127.0.0.1:8080/healthz
{"status":"ok"}   # allowed service behavior
container$ touch /app/cm5-write-probe
touch: cannot touch '/app/cm5-write-probe': Read-only file system
# reference transcript · inspect UID, mount, and rootfs before conclusion
```

A useful hardening check preserves an allowed request and explains a denied mutation.

Read-only, UID, or path?

The probe fails with `Permission denied` and `/app` is absent from the mount list. Which claim is safe?

> No read-only conclusion yet. The path may be missing and the UID may lack permission; inspect the path and mount context before attributing cause.

The error text and mount context determine the claim; failure alone does not.

A secret can leak before the request

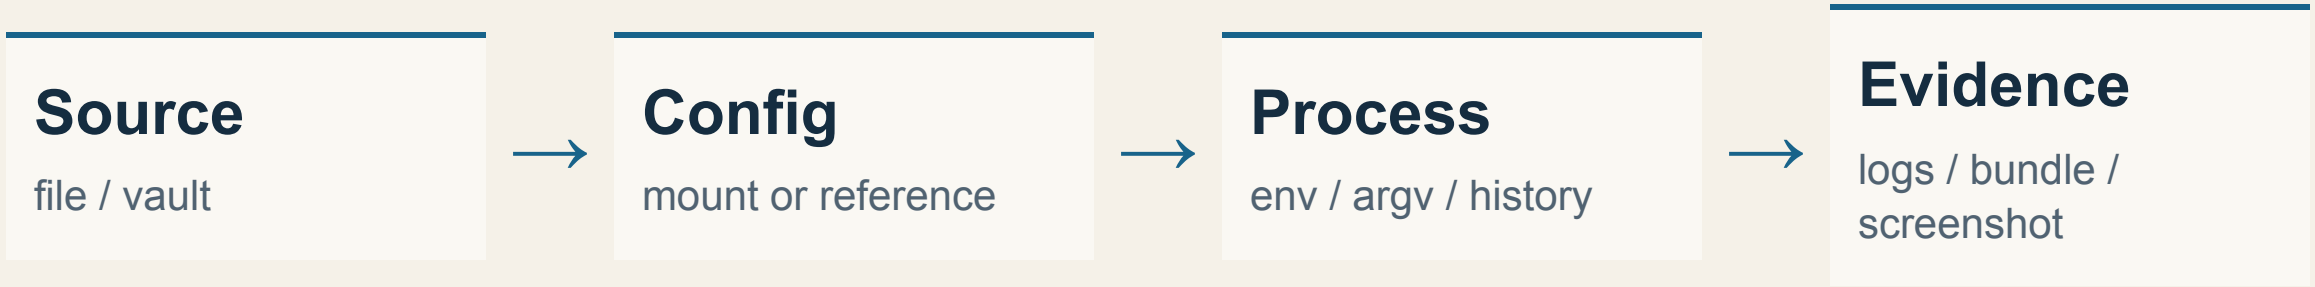

Secret safety includes configuration, process surfaces, and the evidence record.

Path visible; value absent

ALLOWED EVIDENCE

APP_DB_PASSWORD_FILE=
/run/secrets/db_password

path / contract only

FORBIDDEN EVIDENCE

password=••••
env / history / image

value or copy

The contract can expose a secret path while keeping the secret value out of layers, env, history, and evidence.

Least privilege includes the evidence trail

Mount
narrow file path

Permission
only intended UID

Logs
no value / no echo

Evidence
redacted and scoped

A secret control is incomplete if the runtime is safe but the evidence package copies the value.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

An SBOM is an inventory, not a safety verdict

An SBOM identifies components; the decision still needs exposure, exploitability, and runtime evidence.

Scan the exact content you deploy

EVIDENCE CHAIN

tar checksum
→ image ID
→ SBOM
→ scan JSON

BOUNDARY

Input
approved export

Scope
no daemon socket

Cleanup
known temp only

A reproducible scan records its input, digest, generated inventory, findings, and bounded cleanup.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

Zero CVEs is not zero risk

An image has zero high and critical findings.
Can the team call it safe?

> No. Check socket/host authority, public exposure, secrets, mutable identity, runtime controls, and residual risk; scanner coverage and database age also matter.

A clean scan is one input to a risk decision, never the decision itself.

A finding needs a decision row

Finding
ID + component + version

Exposure
reachable? used?

Fix
version / mitigation

Owner
deadline + review

Severity starts triage; exposure, exploitability, fix, owner, and residual complete it.

Severity is not exposure

FINDING

medium component

scanner classification

TRIAGE

used? reachable? fix? owner?

decision context

Triage asks how the component is used and exposed before selecting a disposition.

Every disposition leaves a residual

A finding is fixed in the next image. What still belongs in the record?

> The deployed digest, validation result, owner/date, scanner/database context, and any residual uncertainty or exposure left by the change.

A fix closes one finding; the record still carries identity, validation, and residual uncertainty.

A security change must preserve the contract

Before
health + readiness + data

→

Change
one control at a time

→

After
positive + negative probes

→

Rollback
known artifact / reason

Hardening evidence is a before/after comparison that keeps required behavior and explains denied behavior.

Residual risk is part of approval

Pass
positive contract holds

Pass
negative abuse path blocked

Record
digest + evidence

Residual
owner + review date

Approval means the tested claim is bounded, traceable, and assigned for what remains.

Before TD5: agree the boundary

ALLOWED HOST

Scope
approved Ubuntu lab

Input
approved image export

Role
student pair

STOP CONDITIONS

Secret
stop + report

Socket
reject socket mount

Privilege
reject definition

TD05 begins with an approved host, bounded input, explicit roles, and stop conditions.

Author & Course Lead — Badr Tajini · ESILV · 2026–2027

Evidence packet, then TD05

PACKAGING CHECKLIST · file names are a bounded template, not generated evidence

```
THREAT-MODEL.md    · six abuse cases + owner/residual
sbom.cdx.json      · exact digest linkage
scan.json          · scanner/database context
triage.md          · finding decisions + deadlines
hardening.diff     · one control at a time
positive-negative.md· endpoint + denied-write evidence
```

TD05 starts with a threat model and ends with evidence a reviewer can trace to the tested artifact.

Contingency · recover the claim

10 min

Use for questions, diagnosis, or an approved-host preflight issue.

Recap: which bounded claim will you carry into TD05, and what residual remains?