



Six clocks for AI security. Start, stop, owner, target.

A CompoundTrust working set, not a regulatory requirement. Write each target before you measure.

CLOCK	STARTS WHEN	STOPS WHEN (VERIFIED)	DOES NOT STOP IT
1 Find	Earliest evidence an AI use exists: feature switched on, app access granted, first traffic seen	On the inventory with a named owner	A ticket raised; a scan nobody reviewed
2 Decide	A request to use an AI tool is received, by any channel	Requester told yes, no, or yes with conditions	A review meeting booked
3 Revoke	Decision to remove an agent's or app's access is recorded	Relevant credentials, and any existing sessions, can no longer reach the protected resources (tested)	Account disabled in a console; ticket closed
4 Stop	Stop instruction given	Running task, queue, helper agents and downstream work already triggered confirmed halted; anything that cannot be stopped recorded	The agent's screen says "stopped"; stopping does not reverse completed actions
5 Detect	Harmful or out-of-scope activity begins, including with granted permissions; also record when an alert fired	A responsible person acknowledges it	An alert fired that nobody read
6 Close	Exposed secret or serious finding identified	Fix applied and verified from outside; for a key, old key revoked and tested	New key issued; patch installed but service not restarted

CLOCK	OWNER	TARGET	TYPICAL	SLOWEST (PERIOD, N)	OPEN / UNTESTED	LAST DRILLED
Find						
Decide						
Revoke						
Stop						
Detect						
Close						

RULES

- Start at the event, not the ticket. Stop when the defined outcome is verified.
- A time is good or bad only against a target set first, from the potential harm and the controls available.
- Report typical and slowest, with the period and number of cases. Say when a sample is small.
- Keep unclosed, undetected and untested cases on the page. Unknown means you cannot yet demonstrate how the control performs.
- Speed is one part of the picture. Track coverage too: AI uses with an owner · live agents with a tested stop · clocks with a target.

SAFE DRILL · TEST AGENT, TEST CREDENTIAL, SYNTHETIC DATA, TEST OUTBOX WITH NO EXTERNAL DELIVERY

- Make one harmless request; confirm access works.
- Record the decision to revoke; start stopwatch 1; revoke.
- Repeat the harmless request until refused; note the time. Check existing sessions.
- Fresh run with fake work queued. Call stop; start stopwatch 2.
- Check running task, queue, helper agents and test outbox; note each check time and anything still running.
- Revoking access may not cancel work another service already accepted: check downstream.

Worked example (hypothetical; targets illustrative): find 11 days (target 5 days) · decide 3 weeks (target 1 week) · revoke 40 minutes (target 1 hour; 1 drill) · stop untested · detect unknown · close 19 days (target 1 day; old key left live for a script). Investigate unknowns alongside missed targets, then act in order of risk.

Reporting is separate. Entities covered by CERT-In's Directions of 28 April 2022 must report specified incidents within 6 hours of noticing them or being informed of them. That is a reporting deadline, not a containment target. Check with your legal team how it applies to you.

Context: CERT-In advisory CIAD-2026-0049 (15 Oct 2026). CERT-In Directions under section 70B(6) of the IT Act, 28 Apr 2022. cert-in.org.in