ON SESSIONS, STEP-UP AND THE IDENTITIES THAT ARE NOT PEOPLE

Authentication is a moment; trust is a duration

The strongest sign-in in the world says something true about one instant. Almost every consequential action happens later, on a session nobody has re-examined since.

IN BRIEF

Identity programmes concentrate almost all their effort on the door: stronger factors, phishing-resistant credentials, conditional access at sign-in. The work is correct and insufficient, because the adversary's preferred entry is not the door — it is a token, a cookie or a key that was issued to a legitimate session and outlives the conditions that justified it. This paper argues for treating a session as a continuously re-evaluated claim rather than a fact established once, describes what re-evaluation needs in order not to be theatre, and deals with the identities that vastly outnumber your staff and cannot be asked for a second factor.

1 · What the door does not cover

An organisation deploys hardware-backed, phishing-resistant authentication across the estate. It is a genuine improvement, and it changes the adversary's plan rather than ending it. What follows a successful sign-in is a bearer artefact — a session cookie, an access token, a refresh token, a cached credential — and possession of that artefact is, to almost every system that receives it, indistinguishable from the authentication that produced it.

So the attack moves. It becomes theft of the artefact from a browser profile, an agent's memory, a log file, a build system's environment, or a support tool with delegated access. None of these confronts the authentication strength at all. The credential was never contested; it was *reused*.

The scale of the exposure is set by session lifetime, and lifetimes are long for reasons that have nothing to do with security: re-authentication is disruptive, users complain, and every prompt is a small tax on the day. The result is a control whose strength is measured at one instant and whose validity is asserted for the following eight hours, or thirty days, with **no mechanism to notice that the conditions changed**.

FIGURE 1 · WHERE THE STRENGTH IS, AND WHERE THE RISK IS

THE DOOR

hardware-backed
phishing-resistant
~2 seconds

THE SESSION

8 hours to 30 days · a bearer artefact on a device whose posture changed at hour three, on a network it did not start on, held by a process nobody is watching. Every consequential action happens in here.

Effort is spent almost entirely on the left box. The adversary works almost entirely in the right one — not because the left is impregnable, but because the right is cheaper.

2 · A session as a standing claim

The alternative framing is modest: a session is not a fact but a claim that continues to hold while its supporting conditions do. The conditions are the ones that were true at issue — this device, this posture, this network, this identity behaving in this range — and each is observable afterwards, by a platform that is already collecting them for other reasons.

This is where the identity plane stops being a separate product. Whether a device's disk encryption was disabled at 14:20, whether the process holding the token is the browser or something that read its profile, whether this identity has ever previously touched this system — these are endpoint, network and behavioural signals. An identity system that cannot see them can only re-evaluate on the things it can see itself: a new address, a new user agent, a new country. That is a thin basis, and it produces the two failure modes everyone recognises — prompts that do not correlate with risk, and no prompt at all in the case that mattered.

Bind the artefact to something it cannot travel with. A token cryptographically bound to the device that requested it is worth far less when copied elsewhere. This is the single highest-leverage change available, and it converts theft from a complete compromise into a partial one.

Re-evaluate on change, not on a timer. A fixed lifetime asks the wrong question. What matters is whether anything supporting the claim has changed — and if nothing has, the prompt is pure tax; if something has, the timer is arbitrary and probably too late.

Step up against the action, not the session. Reading a document and rotating a production secret are different asks. Requiring proof proportional to what is being attempted concentrates the friction where it buys something and removes it everywhere else.

Make revocation immediate and complete. Terminating a session must invalidate the artefacts it issued, everywhere, in seconds — not wait for a token to expire on its own schedule. A revocation that takes an hour to become true is a response the incident will outrun.

Friction is a budget, not a virtue. Every prompt spends user patience, and a system that spends it randomly trains people to approve without reading — which is precisely the behaviour that fatigue attacks harvest. Continuous evaluation is worth doing partly because it lets you prompt *less*, and mean it more.

3 · The identities that are not people

In most estates, non-human identities — service accounts, workload identities, API keys, agents, integration credentials — outnumber staff by an order of magnitude. They hold broader entitlements, are provisioned by whoever needed them, rarely rotate, and cannot be asked for a second factor. It is a strange fact of the industry that the identity programme is usually named after the smaller population.

They do, however, have a compensating property: their behaviour is far more predictable than a person's. A batch job runs at the same time, from the same place, against the same resources, forever. Deviation is therefore a much sharper signal for a workload than for a human — and the continuous-evaluation machinery built for sessions applies to them more cleanly than to the staff it was designed for.

WHAT MAKES THEM HARD	WHAT MAKES THEM TRACTABLE
No second factor · long-lived secrets in configuration and images · entitlements granted once for a task and never narrowed · ownership recorded in a person who has left · creation outside any provisioning process.	Narrow, stable behaviour · a knowable set of legitimate callers and destinations · a short-lived credential is usually feasible where a human factor is not · unused entitlements are visible and safe to remove.

The practical programme follows from the right column: inventory them as identities rather than as secrets; attach an owner that is a team rather than a person; measure granted against exercised entitlements and remove the difference; replace long-lived keys with short-lived, attested credentials wherever the platform allows; and baseline each one's behaviour, because **a service account doing something new is a stronger signal than a person doing something new**, and almost nobody alerts on it.

FIGURE 2 · A STOLEN TOKEN, CONTINUOUSLY EVALUATED

08:31 **sign-in, hardware-backed** · device attested, posture clean — the claim is issued

14:07 **a non-browser process reads the browser profile** · endpoint signal, no identity signal at all

14:09 **the token presents from an unattested device** · binding fails — the claim's condition is broken

14:09 **session revoked everywhere; the original device is untouched** · the legitimate user notices nothing

The decisive signal at 14:07 belongs to the endpoint plane and means nothing there on its own. It is only actionable because the identity plane can read it — which is an architectural property, not a policy setting.

4 · Where this gets uncomfortable

Continuous evaluation is a surveillance capability pointed at your own staff, and describing it otherwise would be dishonest. The mitigations are the ordinary ones and they have to be real: evaluate against the session's own baseline rather than a personal profile, keep the signals in a class with its own retention, make the criteria legible to the people subject to them, and never let session telemetry become an input to performance management. An identity plane that quietly becomes a productivity monitor will lose the consent it depends on, and it will deserve to.

There is a second discomfort. Aggressive revocation makes availability a security decision: a false positive logs out a surgeon, a trader, or an on-call engineer at the worst moment. The bounds from Paper 05 apply here in full — revocation is reversible and delegable; disabling an account is not — and the recovery path has to be as well designed as the enforcement, or the first bad revocation will end the programme.

5 · Questions worth asking, whoever you are buying from

01 Are session artefacts bound to a device, and what happens when one presents from elsewhere?

02 Which endpoint and network signals can revoke a live session, and how quickly does revocation take effect?

03 Is step-up tied to the sensitivity of the action, or only to the age of the session?

04 How many non-human identities exist in my estate, and can your platform tell me without a project?

05 Can you show me granted versus exercised entitlements for a service account over 90 days?

06 Do you baseline workload behaviour, and what does a service account doing something new produce?

07 What is the recovery path for a wrongly revoked session, and how long does it take at 03:00?

08 What session telemetry do you retain about staff, for how long, and who can query it?

6 · Where this sits in the series

Identity is where this series' earlier arguments meet a plane that is usually procured separately. The signal that matters at 14:07 is an endpoint observation; the decision at 14:09 is an identity enforcement; the judgement that no change record explains it is legitimacy; the bound that permits revocation but not account deletion is Paper 05's. Four planes, one decision, in two seconds.

Assembled from separate products, that decision is an integration diagram with a human in the middle, and it takes minutes at best. The argument of this series is not that any one of these capabilities is novel. It is that the seams between them are where the time goes, and the seams are an architectural choice.

> *You do not have an authentication problem. You have thousands of sessions, most of them held by something that is not a person, all of them still true because nobody looked again.*

Helix Counter

ABOUT THIS PAPER

Paper 07 in a series on the architecture of active defense, published by **Helix Counter Research**. Papers 01–06 cover enforcement at the edge, explainable alerts, legitimacy, compiling findings into rules, bounded autonomy, and evidence that survives. The series is deliberately vendor-neutral, and the questions in §5 are meant to be asked of us as readily as of anyone else.

Helix Counter runs identity as a plane of the same platform that observes the endpoint and the network, so a session can be re-evaluated on signals the identity system could not otherwise see. helixcounter.com