

Forum:[Hack Forums](#)

Thread: [Critical LangSmith Vulnerability CVE-2026-25750 Allows Silent Account Takeover](#)

JoseMar - March 18, 2026, 10:33 am

Critical LangSmith Vulnerability CVE-2026-25750
Allows Silent Account Takeover

Alright so... I just read about few days ago, and I had to post it here!

Not your usual "user clicked phishing link and got owned" story.

This is cleaner. And honestly, kind of scarier. ?

We're talking about a critical bug in LangSmith, yeah, the AI debugging / monitoring platform from LangChain.

The one that processes insane amounts of data daily.

And somehow... it trusted user input a little too much.

Like, way too much.?

What Actually Broke Here

So the issue is tied to this baseUrl thing inside LangSmith Studio.

Basically devs could point the frontend to different backend APIs.

Sounds flexible, cool feature, right?

Yeah... until you realize it didn't validate where that URL was pointing.

At all.???

So the app just goes, "oh you gave me a URL? sure bro" and sends requests there.

Including session data.

I mean... I guess someone thought "who would abuse this?" but like... come on.

No Phishing, No Interaction, Just... Done

This is the part that got me.

You don't need to enter credentials.

You don't even need to think.

User is logged in --> visits a malicious page --> some hidden JS fires --> loads a crafted LangSmith URL --> boom.

Session token gets sent to attacker-controlled server.

That's it.???

No fake login page, no "verify your account", nothing.

Honestly, feels like the kind of bug that slips through because it's "just a config thing".

Until it isn't.

Why This One Hits Harder Than Usual

Normally, account takeover is bad, yeah yeah.

But here?

We're talking AI observability data.

Which means logs. Traces. Raw outputs.

Sometimes even internal database responses.

So if someone hijacks your account, they're not just seeing dashboards.

They could see:

- proprietary prompts
- internal logic
- maybe even customer data depending on what's logged

And that's the part people underestimate, logs are messy.

Devs dump everything in there "just for debugging".

I've done it too, not proud of it.

The 5 Minute Window Thing

Apparently the Hacker gets like ~5 minutes before the token expires.

At first I thought, "eh, that's tight".

But then again... scripts don't panic.

5 minutes is forever if everything is automated.

Plenty of time to pull data, maybe pivot, maybe lock you out.

So yeah, not as “limited” as it sounds.?

Patch Is Out, But Still...

LangChain fixed it by forcing an allowed origins list.

So now only trusted domains can be used as base URLs.

Which, I mean... that’s exactly what should’ve been there from day one.

Cloud users are already safe.

Self-hosted guys though, you gotta update manually or you’re basically leaving the door open.

And let’s be real, how many people actually update immediately?

Yeah. ? I don’t know...

Not every breach starts with phishing.

Sometimes it’s just... trust in the wrong place.

A single parameter.

A small assumption.

And suddenly your entire AI stack is exposed.

Anyway, curious what you guys think??