


Nine questions to ask your provider about on-chain risk

Every provider has to make the same engineering decisions, and they either enable your compliance team or limit it. You won't see any of them in a demo. These are the questions that reveal what a provider chose.

1 Do you compute exposure during a screening, or in advance?

A stored result is only as current as the last time it was worked out, and stolen funds start moving within seconds.

Red flag: a screening that returns in a few milliseconds, which usually means the result was stored rather than computed.



2 What uptime did you deliver last year, and can you prove it?

99.9% uptime is nearly nine hours a year with no screening. 99.99% is 53 minutes. Ask for the figure delivered, not the figure in the contract.

Red flag: an answer that only quotes the SLA.



3 Tell me about the last time your system handled a volume spike. What did you have to do?

A spike in volume usually brings a spike in risk, and it rarely comes with notice.

Red flag: "Just let us know in advance."



4 How many blockchains do you fully cover, and what does "fully" mean?

"Covered" can mean full transaction history with the entities behind the addresses, or it can mean a sanctions list check and nothing else. Ask what you can actually do on each chain.


Red flag: a single headline number and no breakdown of what "covered" includes.



5 Show me how a screening automatically follows funds across blockchains.

When funds move from one blockchain to another, nothing links the two sides automatically. If your provider hasn't done that work, the trail stops there.

Red flag: a result that shows exposure to a bridge and stops there.



6 If a sanctioned entity creates a new address tomorrow, will your screening know it is theirs?

A sanctions list only ever holds a fraction of the addresses a designated entity controls, and it can create new ones at any time.

Red flag: an answer that covers lists in detail and never mentions entities.



Continue to next page for more

7

How reliable is the data and intelligence behind your risk rules?

A rule against mixer exposure only means something if someone has labeled which addresses are mixers, so ask what your configuration is applied to, not just what you can configure.

Red flag: a sophisticated answer about the rules and a vague one about the data.



8

How do you keep risk accurate over time without flooding my team with notifications?

A score is only true at the moment it was calculated. An address that scored low at onboarding can become high risk months later.

Red flag: every change becomes an alert.



9

Have you tested a full failover of your primary region, and can you show me the results?

If everything runs in one place and that place fails, you cannot screen. Ask whether they have tested it, not just planned for it.

Red flag: no successfully tested failover of a primary region.



Add your
notes here

Ask us first. **Built for compliance** sets out each of the nine decisions and the choices Elliptic made. Get your digital copy.

