

What is an API key, and what happens if one is stolen?

Short answer: An API key is a secret string that identifies and authorizes a program — rather than a person — when it calls a service. A stolen key lets an attacker make requests as your application, often with no user interaction, no MFA, and no obvious sign anything is wrong.

Why they're attractive targets

- **No human in the loop.** Keys are built for automated, unattended use, so there's no MFA prompt to interfere.
- **Often broadly scoped.** Keys are frequently issued with more permission than the task requires.
- **Long-lived.** Many are never rotated, so a key stolen today may still work months later.
- **Stored in predictable places.** Configuration files, environment variables, and credential stores — all locations automation can enumerate.

Typical consequences

Depending on what the key authorizes: reading or exporting customer data, sending messages as your organization, spending money against a cloud account, or pivoting into other connected systems.

Because the requests are properly authenticated, they usually appear in logs as normal application traffic. Detection commonly happens well after the fact, via a bill or an audit.

Reducing exposure

Scope keys narrowly, rotate them regularly, and keep them out of source control. Those are sound practices, but they limit blast radius rather than prevent theft.

Cyber Crucible targets the theft itself. Credential and configuration locations are protected, and a program reading them is evaluated in context: fake data or denial for a legitimate overreach, rejection or suspension for an unknown or compromised process. The key is not handed over either way.

