Attackers now automate the intrusion, and waiting for an alert is too slow

What October and November 2026 showed, and what to change.

Four questions

- **What happened.** The named cases, strongest evidence first.
- **What changed.** Speed, cost and scale, not the entry points.
- **What is not proven.** Targeted is not breached.
- **What to do.** Close old doors and cut time to contain.

200+

organisations targeted in 11 days

ARTEX agents, 14+ countries, from 17 October. Nine have confirmed post-exploitation. Source: CrowdStrike.

~20 min

for an agent to take domain controller credentials

Lateral movement with no further human guidance. All four actors CrowdStrike found overcame AI safety controls.

~$50

per successful website intrusion

About 23 hours each. One operator tried 1,200+ sites and reached 100+ companies. Source: NorForge via Reuters, estimate.

People set the goal. Agents do the middle.

1. **Human:** VPN access and a goal, framed as authorised red-team work.
2. **Agent:** maps Active Directory and file shares.
3. **Agent:** password spraying, Kerberoasting, NTLM relay.
4. **Agent:** rewrites failed commands, tries again.
5. **Human:** steals data, deploys ransomware. Source: Sophos.

5 days

from first access to a person noticing

25 October to 30 October. A predictable dev credential in production, exposed internal APIs and no rate limiting.

Be exact about the claims

- 200 targeted, **9 confirmed**.
- 1,200 sites tried, 100+ reached.
- Humans still pick targets and payloads.
- The WorkOS sandbox escape is one person's account.
- "Unprecedented" is an NCSC forecast, not a count.

Six changes

- Remove default and dev credentials from production.
- Take internal APIs off the internet.
- Rate-limit every data-bearing API.
- Segment, isolate backups, trim domain admins.
- Pre-authorise reversible containment.
- Test with agents in a sealed environment.

Measure time to contain, not only time to detect



Terence Kok · AI & Innovation Strategist
with Maria Singh · SAVE · COMMENT · FOLLOW