

CASE STUDY

US Federal Agency Red Team Runs Assume-Breach Tests Across a Hundreds-Site Estate

A US federal agency protects hundreds of sites nationwide, all in one Active Directory domain, with a three-person security engineering team. Its detection engineers had run Picus Breach and Attack Simulation for over a year. Then the four-person red team took on a different question: once an attacker has a foothold in one site, how far can they get?

In summer 2026 the red team began gray-box, assume-breach assessments with Picus Autonomous Penetration Testing. The runs showed no successful lateral movement from the initial foothold, but they did find an exploitable path to credentials: on some hosts, password hashes could be extracted directly from the operating system.

Industry:

Federal Government



Products:

- Autonomous Penetration Testing
- Breach and Attack Simulation (BAS)

About the company

The organization is a US federal agency supporting hundreds of sites nationwide from a single Active Directory domain. A staff of three runs its validation program alongside a four-person red team, on a fully on-premises Picus deployment.

Challenges and Results:

- **Hundreds of sites in one domain, three people to cover them**
→ Autonomous pentesting that scales across hundreds of sites
- **Post-compromise paths untested from a real foothold**
→ Attack paths validated by execution from a real foothold, site by site
- **Credential exposure unknown from inside a site**
→ Password hashes successfully extracted, with evidence
- **No evidence of how far an attacker could move from a foothold**
→ Runs showed no successful lateral movement
- **Pentest findings hard to act on across hundreds of sites**
→ A clear attack map from every autonomous pentesting run
- **Detection engineering needed realistic, repeatable testing**
→ BAS in regular use for over a year, feeding SIEM dashboards

Problem:

Hundreds of Sites, One Domain, a Staff of Three

The agency's estate is one Active Directory domain, with each site organized as an organizational unit rather than a separate domain, so a foothold at any one site sits inside the same directory as every other. Detection engineering team had validated controls with BAS for more than a year. The red team needed to see what happens after initial access, site by site, without the people to do it by hand across the whole estate.

"We're hoping to do this all automatically because otherwise... we only have a staff of three to manage it," said the Senior Security Engineer.

The team focused on two priorities:

- Test what an attacker could do from a foothold inside a site, under an assume-breach model.
- Scale testing across hundreds of sites so the red team can focus more on work that needs creativity and deeper thinking.

Solution:

Autonomous Pentesting, One Site at a Time

Rather than adding another tool, the agency extended the Picus Platform it already ran to the red team, with **Picus Autonomous Penetration Testing**.

It works like a real attacker, mapping the network in scope, executing real-world exploitation techniques, and chaining them step by step to show what an attacker can actually reach. Teams define the scope for each run, so the test stays controlled and safe to run in production. Each run ends with an attack map of every step completed, giving teams validated findings they can act on. By handling the repeatable offensive work, it lets red teams scale their coverage.



"We're getting ready to... bring our red team on to use the tool a lot more. Not just our detection engineering team." The Senior Security Engineer

Outcome:

Evidence of What an Attacker Can Do, and What Stops Them

The autonomous pentesting runs gave the red team a clear picture of what an attacker could execute from inside a site. Lateral movement was attempted from the initial foothold, and no path succeeded. Credential access was a different story: on some hosts, the runs successfully extracted password hashes straight from the operating system.

On the control side, detection engineering continues to run BAS under both standard-user and admin credentials. The simulations show whether the agency's security controls block or detect attacker behavior, and the results feed the custom SIEM dashboards the team built for internal reporting.

Both teams now work from one platform and one body of evidence: the red team shows which actions an attacker can execute from inside a site, and detection engineering shows how the agency's controls respond to that behavior. That shared view replaces assumptions about each site with results.



"Two out of 29. Maybe 10%, right? That tells you something," the Senior Security Engineer said about detection coverage across sites. *Picus* turns ***"what could an attacker do here?"*** and ***"would our controls stop it?"*** into evidence, site by site.



Picus Security, a leading exposure validation company, **proves what attackers can exploit and what your defenses stop**, then closes real gaps with ready-to-deploy fixes and re-validates to confirm, at the machine speed today's AI threats demand. The Picus Platform spans [Breach and Attack Simulation](#), [Autonomous Penetration Testing](#) and [Exposure Validation](#), unified by [Picus Swarm](#), a swarm of AI agents that runs the whole validation loop continuously with human oversight. With 75+ integrations, it reaches across on-prem, hybrid cloud, and endpoint environments. Trusted by many Fortune 500 enterprises, Picus was named a 2026 Gartner Peer Insights Customers' Choice for Adversarial Exposure Validation and is recognized as an Innovation Leader in the Frost Radar for Automated Security Validation.