How to Secure Endpoints in Hybrid Work Environments

Picture a Tuesday morning at any mid-size U.S. company. A sales rep logs into Salesforce from a hotel lobby in Chicago on a personal laptop, no VPN. A developer pushes a commit from a home machine four months behind on OS patches. A finance analyst pastes a revenue spreadsheet into an AI tool that nobody in IT approved. Before 10 AM, you have three real endpoint security gaps. None of them triggered an alert.

That’s hybrid work in 2026. And it’s not going away.

IBM’s 2025 Cost of a Data Breach Report put the average breach cost for U.S. organizations at $10.22 million, a record high. Verizon’s 2025 Data Breach Investigations Report, drawing on over 22,000 security incidents, found ransomware in 44% of all confirmed breaches and documented a roughly eightfold jump in VPN-targeted exploits year-over-year.

Read those two data points together: the infrastructure hybrid work depends on is the exact infrastructure attackers are prioritizing right now.

Flexible work has become a baseline employee expectation across professional roles. Organizations are not going to reverse this. So securing endpoints in hybrid work environments means building visibility and response capability that follows your users, not waiting for them to come back to the office network.

Securing endpoints in hybrid work isn’t about locking everything down. It’s about building controls that work wherever your users actually are.

5 Biggest Endpoint Security Risks in Hybrid Work (At a Glance)

Before looking at specific controls, it helps to understand where the documented risk concentrations are. These five categories come directly from Verizon’s 2025 DBIR and IBM’s 2025 breach research. They are not survey opinions, but confirmed incident patterns across thousands of real breaches.

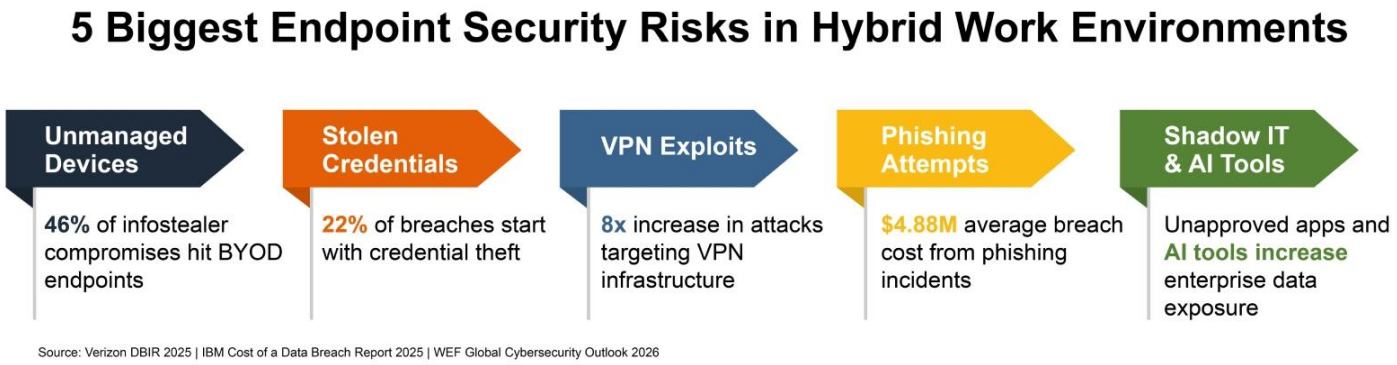

Notice the pattern. Attackers aren’t going after hardened corporate servers. They’re going after the soft edges: personal devices, credentials used on public networks, VPN appliances running

unpatched firmware. Each of these is a direct consequence of how hybrid work is structured.

What the 2025 Breach Data Shows About Hybrid Work Threats

Three authoritative sources define the current threat baseline. Cross-referencing their findings gives a more complete picture than any single report:

Report Key Statistic What It Means for Hybrid Endpoints Verizon DBIR 2025 Ransomware in 44% of breaches; credential theft in 22% Remote workers are the primary credential theft target Verizon DBIR 2025 46% of infostealer hits targeted unmanaged/BYOD devices Personal devices are the largest single endpoint blind spot Verizon DBIR 2025 VPN-targeted exploits increased roughly 8x year-over-year Remote access infrastructure is now a primary attack vector IBM Breach Report 2025 AI automation saved organizations $1.9M per incident on average Automated detection directly reduces financial exposure IBM Breach Report 2025 Shadow AI breaches cost $670K more than average incidents Unapproved AI tools on endpoints are a documented loss vector WEF Cybersecurity Outlook 2026 72% of security leaders say AI advantages attackers over defenders Threat velocity is outpacing manual security response capabilities

Worth calling out specifically: edge device exploitation of routers, VPN appliances, and firewalls now sits alongside phishing as a primary initial access method.

Some CVEs are reaching mass exploitation status on the same day they’re publicly disclosed. For organizations still running 30-day patch cycles on remote access infrastructure, that’s a serious structural exposure.

Hybrid Infrastructure, Hidden
Risk: The Visibility Problem
CISOs Must Address

- Delayed Detection
- Increased Exposure of Sensitive Data
- Alert Fatigue and Analyst Burnout

Get the Practical Playbook to Address Hybrid Risks

Why Traditional Antivirus Software Can’t Protect a Hybrid Workforce

Legacy antivirus was purpose-built for a specific threat model: known malware, delivered as files, on endpoints sitting behind a corporate firewall. That described most enterprise threats through roughly 2015. Today it describes a shrinking minority of actual attacks.

Three structural gaps make legacy tools wrong for hybrid environments:

- **Signature matching can’t catch what it hasn’t seen.** Fileless malware and living-off-the-land (LOLBin) attacks execute entirely through legitimate system processes like PowerShell, WMI, or cmd.exe. Nothing gets written to disk, so there’s nothing to scan. Verizon’s 2025 DBIR documented LOLBin abuse across every industry vertical in the research dataset.
- **There’s no response capability built in.** Quarantining a file and firing an alert is the ceiling of what traditional antivirus can do. It can’t isolate a compromised endpoint, terminate an active malicious session, or begin forensic collection before an attacker clears evidence.

Endpoint security for hybrid work demands that response capability. On a remote laptop that stays VPN-connected for hours after initial compromise, its absence is extremely costly.

- **It sees one device in isolation.** A ransomware campaign entering through a phishing email on a home machine, using stolen credentials to access a cloud application three hours later, then moving laterally to an on-premises file server that evening. Traditional endpoint AV has no visibility into that chain. One alert on one device, if anything.

Microsoft's security research puts a concrete number on this: 80-90% of successful ransomware attacks originate on unmanaged devices. Personal machines, home computers, tablets where traditional antivirus either isn't installed, isn't current, or lacks the behavioral depth to catch modern attack techniques.

6 Endpoint Security Strategies That Work in Hybrid Work Environments

No single control secures a distributed endpoint environment.

What works is layering these six strategies so each one closes gaps the others leave open. All are grounded in CISA guidance, NIST standards, and the documented incident patterns from primary research. They are not vendor recommendations.

1. Build Every Access Decision on Zero Trust Principles

Zero trust starts from one uncomfortable premise: you can't trust that any access request is legitimate just because it's coming from inside your network, from a recognized device, or from an account with a valid password.

Every request gets evaluated against what you actually know: who's asking, from which device, from where, and whether that combination makes contextual sense.

CISA's Zero Trust Maturity Model and NIST SP 800-207 both frame this across five pillars: Identity, Devices, Networks, Applications, and Data. Each pillar demands the same things: explicit verification, least-privileged access, and continuous monitoring throughout every session. Not just at login.

For hybrid endpoint security, zero trust principles translate into four specific controls:

- Conditional access policies that verify device compliance before approving access requests
- Phishing-resistant multi-factor authentication using FIDO2 hardware keys or passkeys
- Least-privilege access limiting accounts to required resources only
- Continuous session verification that reevaluates risk during active sessions

2. Deploy EDR with Automated Threat Response — Not Just Alerts

Endpoint Detection and Response platforms work at the process level. Rather than scanning files, they capture behavioral metadata for every process, file system change, registry modification, and network connection across each endpoint, forming the visibility layer that modern endpoint security is built on.

When something deviates from established baseline behavior or matches a known technique in the MITRE ATT&CK framework, the platform responds.

And response is the operative word. An EDR system that can automatically isolate affected devices from the network the moment suspicious behavior is confirmed changes the economics of a breach fundamentally.

Real-time threat detection paired with automated response (not alert-and-wait workflows) is what IBM's 2025 data shows makes the measurable difference. Organizations with full security AI and automation deployment detected and contained incidents in 258 fewer days on average

versus those relying on manual processes. That's not incremental improvement; it's a different category of outcome.

For remote employees across every remote work environment (home broadband, hotel Wi-Fi, shared co-working spaces), the off-network capability is critical. A persistent agent on the device maintains the same visibility and automated response depth whether it's on the corporate LAN or operating entirely outside it.

Machine learning adds the ability to detect unknown threats through behavioral anomalies: a process running from an unusual directory, an admin tool executing command patterns with no prior history, a sudden spike in encrypted outbound traffic at 2 AM.

3. Establish Visibility Across All Devices with Unified Endpoint Management

Guardz and the Ponemon Institute found that 92% of remote workers use personal phones or tablets for work tasks, and 46% save work files on those personal devices.

Each one is an access point to company data with no security visibility.

Unified endpoint management (UEM) closes this through centralized device management across all endpoints (corporate-issued and enrolled BYOD) with consistent policy enforcement regardless of device ownership.

Key controls enforceable through UEM:

- Block network access from devices running outdated operating systems or missing required security patches
- Require full-disk encryption on any device capable of storing company data
- Mandate EDR agent enrollment before a device is permitted to access corporate resources
- Enable remote wipe capability for lost, stolen, or decommissioned devices
- Apply application policies that restrict installation of high-risk software categories

Offboarding matters here too. Capterra's research found 71% of HR teams report at least one former employee who never returned company devices. Solid device management practices and security measures tied to HR workflows close that gap: automatic access revocation triggered at offboarding removes the dependency on manual IT tickets entirely.

In distributed remote work environments where devices are rarely physically returned to IT, this automation is the only reliable control.

4. Harden Your VPN or Evaluate Zero Trust Network Access (ZTNA)

An eightfold year-over-year increase in VPN-targeted exploits is not a statistical blip. It's a documented attacker priority shift. A virtual private network is the infrastructure layer that makes remote operations possible for most organizations, and right now these virtual private network connections are among the most actively targeted enterprise entry points. Minimizing risks here requires both immediate hardening and a longer-term architectural review.

CISA and NSA published joint guidance on VPN security with specific recommendations:

- Immediate patching when vulnerabilities are disclosed
- Phishing-resistant MFA for all VPN authentication

- Disabling unnecessary features that expand attack surface
- Continuous monitoring of access logs

The patching urgency is real. Some CVEs in VPN appliances are reaching mass exploitation on zero-day timelines.

Longer term, Zero Trust Network Access offers a more structurally sound approach. Where traditional VPN authenticates once and grants broad network access, ZTNA scopes each connection to a specific application and continuously verifies posture.

A compromised VPN credential gives an attacker a network-level foothold. A compromised ZTNA session is scoped to one application.

5. Extend Data Loss Prevention to Cover AI Tools and Shadow IT

Even with strong access controls and EDR coverage, data can leave through channels that look completely legitimate. An employee syncing a client contract to a personal Dropbox account.

Someone pasting internal financials into a public AI assistant to summarize for a board deck. A departing employee forwarding their contact list to personal email before their last day. Data loss prevention controls are specifically designed for these scenarios.

Shadow AI is the newest and fastest-growing exposure. IBM's 2025 breach data found that incidents involving unapproved AI tools cost an average of $670,000 more than standard breach events.

When employees paste proprietary data into a public-facing AI assistant, that information may be processed and retained by a third-party service with no contractual data protection relationship. DLP policies built on content inspection rather than just domain blocking can catch this.

For hybrid teams, endpoint security controls including DLP need to be device-resident. An employee working off-network and copying sensitive data to a personal sync service isn't going through your network-layer controls unless those controls live on the device itself.

6. Use Security Awareness Training to Build a Human Detection Layer

Most organizations frame security awareness training as a compliance function. It's actually a detection mechanism, and an underused one. Verizon's 2025 DBIR found that employees who've received recent phishing training report suspicious emails at a 21% rate, compared to just 5% for untrained staff.

That gap is significant in practice. A trained employee who reports a suspicious email isn't only protecting themselves, they're potentially triggering a response that catches an active campaign before it hits the broader organization. For remote employees working outside on-premises security controls, this human-layer alerting becomes especially valuable.

Training in 2026 needs to specifically address AI-generated phishing. Verizon's DBIR noted that AI-crafted phishing emails have roughly doubled in volume and quality. They're grammatically clean, contextually plausible, and often personalized using data scraped from LinkedIn or prior email exchanges.

Older training signals such as poor grammar, generic greetings, and obviously suspicious URLs

are less reliable as detection cues against this generation of attacks.

Hybrid Endpoint Security Controls: Priority Reference

Security teams managing hybrid environments often need to justify prioritization decisions across multiple controls. This table maps each strategy to the risk it addresses and the authoritative guidance source:

Security Control Risk Addressed Priority Guidance Source Zero Trust + Conditional Access

Credential theft

, unauthorized access Critical CISA ZT Maturity Model v2; NIST SP 800-207 EDR with Automated Response Malware, fileless attacks,

lateral movement

Critical CISA; Verizon DBIR 2025; IBM 2025 Phishing-Resistant MFA Account takeover, credential bypass Critical CISA MFA Guidance; Verizon DBIR 2025 Unified Endpoint Management Unmanaged BYOD, patch gaps, offboarding High NIST SP 800-124 Rev. 2; Ponemon 2025 VPN Hardening / ZTNA Remote access hijacking, VPN exploits High CISA/NSA Joint VPN Advisory Data Loss Prevention (DLP) Data exfiltration, shadow AI,

shadow IT

High IBM Breach 2025; WEF Outlook 2026 Security Awareness Training Phishing, social engineering Medium Verizon DBIR 2025; CISA Resources Patch & Lifecycle Management Vulnerability exploitation, stale endpoints High Verizon DBIR 2025 (34% YoY growth)

Why Continuous Monitoring Matters More Than Periodic Compliance Checks

Each strategy above generates data: telemetry from endpoints, access logs, DLP policy violations, EDR alerts. What determines whether that data actually protects the organization is whether it’s being analyzed continuously or collected and reviewed after the fact.

Risk profiles change throughout the day in hybrid environments. An employee who started work on a secure home connection moves to a coffee shop network at noon. A device that passed its 9 AM compliance check installs a browser extension at 2 PM that carries a known malicious component. A legitimate account starts accessing files at a volume and pattern that doesn’t match any prior behavior. Periodic scans miss all three.

CISA’s Zero Trust Maturity Model makes continuous monitoring an explicit architectural requirement, not a best practice. Verification needs to happen throughout every active session, triggered by context changes, not only at session initiation.

Extended Detection and Response platforms operationalize this by correlating telemetry across endpoints, network traffic, and cloud workloads simultaneously. When initial access happens on a remote laptop, a credential appears in a cloud application three hours later, and data exfiltration surfaces in network logs that evening, XDR connects those events into a single incident timeline. Siloed endpoint security tools working in isolation can’t reconstruct that chain.

IBM’s 2025 data quantifies the outcome difference: organizations with full security AI and

automation deployment detected and contained breaches in 258 fewer days on average versus those relying on manual processes. Over eight months. That’s the window in which an undetected breach does the vast majority of its damage.

How Fidelis Secures Endpoints in Hybrid Work Environments

Enterprise endpoint security for hybrid environments must work consistently both on and off the corporate network, across managed endpoints and enrolled BYOD devices, and across operating systems including Windows, macOS, and Linux. It also needs to respond to threats without requiring manual analyst approval for every action, because manual response workflows do not scale across distributed endpoint fleets.

Fidelis Endpoint® Detection and Response is designed to secure endpoints operating across distributed and hybrid environments. It operates through a single-agent architecture that captures full process-level behavioral metadata regardless of network location. On a corporate LAN or a home broadband connection, the agent maintains the same monitoring depth.

At the moment of detection, automated response actions can be triggered immediately:

- Automatic isolation of compromised endpoints cuts off lateral propagation while preserving access for forensic investigation
- Forensic collection at detection time captures the full process tree, file activity, registry changes, and network logs before an attacker can clear their tracks
- MITRE ATT&CK mapping gives analysts immediate context on what technique is in use and what the likely next stages are
- Retrospective threat analysis applies updated threat intelligence to historical endpoint telemetry, surfacing previously undetected compromises
- SIEM and SOAR integration fits into existing SOC workflows rather than requiring a parallel investigation track

Integration with Fidelis Elevate® XDR extends visibility across endpoint, network, and cloud telemetry in a unified analytical layer. For hybrid environments where attacks regularly span all three surfaces, that correlated view enables endpoint detection earlier in the kill chain.

Machine learning analytics and continuously updated threat intelligence feeds run across the platform, flagging anomalies that no signature database would catch. Fidelis also provides continuous verification that zero trust controls are functioning as configured, catching attempts by unauthorized actors to circumvent access policies that should be limited to authorized personnel only.

Advanced Threat Detection with Fidelis Elevate®

Don’t let threats go unnoticed. See how Fidelis Elevate® helps you:

- Identify and neutralize threats faster
- Gain full visibility across your attack surface
- Automate security operations for efficiency

Download Now

5 Endpoint Security Priorities for Security Teams in 2026

Global endpoint security market spend is projected to reach $24.9 billion in 2026, growing toward $44.7 billion by 2033. Budgets are increasing. But higher spend doesn’t automatically translate to better outcomes. Organizations that improve actual security results focus their investment on the highest-leverage priorities:

Priority Evidence Base 1 Complete endpoint asset visibility 46% of infostealer compromises hit unmanaged devices. Every device accessing company data needs to be inventoried, regardless of who owns it. 2 Zero-delay vulnerability patching Vulnerability exploitation grew 34% year-over-year. Some CVEs in edge devices hit mass exploitation on zero-day timelines. Standard 30-day patch cycles don’t cover this. 3 Phishing-resistant multi-factor authentication (MFA) across every account 88% of web application breaches involved stolen credentials. SMS and push-based MFA are routinely bypassed via AiTM attacks. Device-bound FIDO2 multi-factor authentication closes this gap. 4 Tested incident response with communication protocols 86% of breached organizations experienced operational disruption including halted production. Organizations with rehearsed IR plans recover faster and at lower total cost. 5 DLP and access controls extended to AI tools Shadow AI incidents cost $670K more than average breaches. AI tool adoption on endpoints is outpacing governance. Content-based DLP covering AI endpoints is now a concrete requirement.

Endpoint Security in Hybrid Work Is an Ongoing Discipline

Security teams doing this well in 2026 have made a specific mindset shift. They've stopped treating endpoint security in hybrid work environments as a project with a finish line and started building an overall security posture designed to evolve alongside the threats targeting it.

That environment keeps changing. Device fleets grow. New applications get adopted every quarter. Employees take on new work patterns. AI tools get integrated into daily workflows faster than IT governance processes can track. And attackers study all of it, continuously updating which edges of the hybrid work attack surface are most exploitable at any given moment.

Every unmanaged device with access to company data, every VPN connection lacking phishing-resistant MFA, every endpoint running an unpatched OS is a potential entry point.

Endpoint security in hybrid work environments demands active, continuous controls, not periodic compliance snapshots. What determines the outcome is whether your detection capability catches a compromise early enough to contain the damage, and whether your response capability acts fast enough to limit it.

Protecting endpoints in hybrid work environments isn't a deployment. It's a discipline, built, tested, and refined continuously because the threat environment never stops evolving.

Fidelis Security builds endpoint detection, response, and XDR capabilities specifically for complex distributed environments. If you're evaluating where your current endpoint protection stack leaves gaps, we can help with that assessment.

Frequently Ask Questions

What makes endpoint security in hybrid work different from traditional enterprise security?

Traditional enterprise security assumed most devices operated inside a defined network perimeter, firewalls and gateways enforced controls at a fixed boundary.

In hybrid work environments, that boundary effectively doesn't exist. Employees connect from home broadband, public Wi-Fi, and personal devices that IT has limited visibility into.

Endpoint security controls have to travel with the user. That requires persistent agents on each device, zero trust access policies that evaluate every connection independently, and automated response capability that functions whether or not the device is on the corporate network.

Which endpoint security risks are most exploited in hybrid work environments?

Based on Verizon's 2025 DBIR and IBM's 2025 Cost of a Data Breach Report, the highest-impact risks are: unmanaged personal devices (46% of infostealer compromises), credential theft as an initial access vector (22% of all breaches), VPN infrastructure exploits (growing approximately eightfold year-over-year), and phishing campaigns targeting remote employees ($4.88 million average incident cost).

Shadow AI tools have emerged as a rapidly growing fifth category, IBM found these incidents cost an average of $670,000 above standard breach figures.

How does zero trust architecture strengthen endpoint security for remote employees?

Zero trust replaces network location as a trust signal with continuous verification of identity, device compliance, and session behavior.

For remote employees, this means endpoint security policies are enforced consistently regardless of where they connect from. A device missing a required patch, lacking an EDR agent, or connecting from an unusual location is blocked before it can access company resources.

The Cybersecurity and Infrastructure Security Agency Zero Trust Maturity Model defines implementation across five pillars: Identity, Devices, Networks, Applications & Workloads, and Data. NIST SP 800-207 provides the core architectural guidance for building Zero Trust environments that enforce these continuous verification principles.

What should CISOs prioritize when building an endpoint security strategy for hybrid work?

Start with complete asset visibility, endpoint security controls cannot be applied to devices that haven’t been inventoried. From there, the data points to four priorities: phishing-resistant MFA on every account (credential theft drives 22% of initial access events and traditional MFA is increasingly bypassed); immediate patching of VPN and edge device vulnerabilities (some CVEs are reaching mass exploitation on zero-day timelines); extending DLP controls to AI tools and shadow IT; and deploying automated endpoint detection and response.

IBM’s 2025 research found that organizations with AI-assisted detection contain breaches 258 days faster on average than those relying on manual processes.

References:

1. IBM Security. (2025). Cost of a Data Breach Report 2025. Ponemon Institute.
2. Verizon Business. (2025). 2025 Data Breach Investigations Report (DBIR).
3. CISA. (2023). Zero Trust Maturity Model v2.0. U.S. Cybersecurity & Infrastructure Security Agency.
4. NIST. (2020). Special Publication 800-207: Zero Trust Architecture.
5. World Economic Forum. (2026). Global Cybersecurity Outlook 2026.
6. CISA & NSA. (2021). Selecting and Hardening Remote Access VPN Solutions.
7. NIST. (2023). Special Publication 800-124 Rev. 2: Guidelines for Managing Mobile Device Security.
8. Persistence Market Research. (2026). Endpoint Security Market Size & Forecast 2026-2033.