DIGITAL DEFENSE DIGEST

Insider Tips To Make Your Business Run Faster, Easier And More Profitably



WHAT'S NEW

Special Feature: Protecting AI Deployed In Your Organization - the shadow AI problem, deepfake executive fraud, and a practical playbook (pages 3-4)

Ransomware Watch 2026 - what's changed for small and mid-sized businesses this year (page 5)

Grant & Compliance Spotlight - where the SLCGP/PILLAR Act stands and how to fund your next security project (page 6)

Ten Low-Cost Moves every lean team should make this October (page 7)

Workshop calendar and how to schedule your free assessment (page 8)

This monthly publication is provided by Black Belt Secure

BLACK BELT SECURE

OUR MISSION:

To empower businesses of all sizes with the knowledge and tools they need to thrive in today's increasingly complex cyber threat landscape. We are dedicated to providing cutting-edge cybersecurity solutions and expert guidance to help our clients

OCTOBER IS CYBERSECURITY AWARENESS MONTH — AND THIS YEAR, YOUR BIGGEST RISK MIGHT BE THE AI TOOL SITTING OPEN IN A BROWSER TAB

Every October since 2004, the Cybersecurity and Infrastructure Security Agency (CISA) and the National Cybersecurity Alliance have used Cybersecurity Awareness Month to push the same basic message: small, consistent habits stop most attacks before they start. This year the two organizations are running it under two banners. CISA's theme, Securing the Next 250, ties cybersecurity to the nation's 250th anniversary and the operators who keep power, water, and other essential services running. The Alliance's theme, Don't Make It Easy for Them, speaks more directly to everyday office life - the click, the callback, the approval a scammer is counting on you to skip.

Both themes point at the same shift happening underneath them. For a decade, awareness training told employees to watch for bad grammar, generic greetings, and a sender address that didn't quite match. Those tells largely stopped working somewhere around 2024.

Today's phishing emails are drafted by the same class of AI tools your team uses to draft its own emails. A convincing voice on an urgent phone call can now be cloned from a few seconds of public audio - an earnings call, a conference talk, even a LinkedIn video. Attackers no longer need a technical vulnerability when a confident, well-timed impersonation gets the job done.

That's why this month's edition of the Digital Defense Digest runs longer than usual. Alongside our regular coverage, we're dedicating a full special feature to a question more of our clients are asking us every week: **How do you protect a business that's already using AI, often in ways IT never approved, without shutting productivity down?**

continued on page 2...

...continued from cover

By The Numbers

Stat	What It Means
2004	Year Cybersecurity Awareness Month launched
2	Official 2026 campaign themes (CISA + National Cybersecurity Alliance)
43%	Breached organizations with a shadow AI incident in the last year (IBM)
$670K	Average added breach cost when shadow AI is involved (IBM)

WHY 2026 FEELS DIFFERENT

A fair question for any business owner: hasn't cybersecurity always been an arms race? Yes - but the pace changed this year. Two headline stories from Black Belt Secure's own threat research in September illustrate the shift. A Russian state-sponsored group Anthropic tracks as GTG-20006 was found running a self-healing malware campaign, in which the malicious code repairs and reconfigures itself automatically after security tools flag it, cutting the manual rebuild work attackers used to need. Separately, researchers at Sophos detailed a newly analyzed Linux rootkit targeting F5 BIG-IP Access Policy Manager devices - malicious code that alters a server deeply enough that a routine file scan reports the machine as clean when it isn't. Both stories point at the same trend: attackers are automating the parts of an intrusion that used to require a skilled human, freeing them to run more attacks, faster, against more targets - including businesses that never thought they were big enough to attract this level of sophistication.

Employees are facing an AI-accelerated version of the same shift. A cloned voice, a synthetic video call, or a flawlessly written phishing email removes the small tells that used to give an attack away. That is precisely why this year's Digital Defense Digest devotes a full special feature, starting on the next page, to the AI risks now sitting inside most organizations - many without their leadership's knowledge.

What's New

- Cybersecurity Awareness Month opens October 1, with CISA's Securing the Next 250 and the National Cybersecurity Alliance's Don't Make It Easy for Them running as companion campaigns nationally.
- Cybersecurity Career Week runs October 19-24, part of the national push to grow the security workforce.
- The PILLAR Act, which would extend federal State and Local Cybersecurity Grant Program funding through 2033, passed the House in November 2025 and remains before the Senate Homeland Security Committee. We're tracking it and will flag any movement in next month's issue.

This Month's Threat Snapshot

Trend	What We're Seeing
Self-healing malware	State-linked actors deploying malware that auto-repairs after detection, extending dwell time
Firmware-level rootkits	Attacks that alter network appliances so standard scans report a false clean bill of health
AI-written phishing	The grammar and formatting tells that training relied on for a decade are no longer reliable
Voice-cloned fraud	A usable clone of an executive's voice needs as little as three seconds of public audio



PROTECTING AI DEPLOYED IN YOUR ORGANIZATION

The Tool Your Team Already Uses Without Asking

Here's an uncomfortable exercise: ask five employees, one at a time, whether they've ever pasted a document, a client name, or a chunk of code into a free AI chatbot to save time. In most offices, the honest answer from most of them is yes. Security teams have a name for this: shadow AI – the use of AI tools the business never evaluated, approved, or even knows about. It isn't malicious. It's an employee trying to finish a report faster, summarize a contract, or debug code before a deadline. But every one of those moments is business data – sometimes regulated data – leaving your organization's control and landing on a server you don't own, governed by terms of service most people never read.

This isn't a hypothetical risk anymore; it's a documented, growing line item in breach reports. IBM's 2026 Cost of a Data Breach Report found that 43% of breached organizations had experienced a shadow AI incident, more than double the 20% recorded the year before – and that shadow AI involvement added an average of $670,000 to the cost of a breach in IBM's prior-year research. ISACA's 2026 polling of digital trust professionals found 90% believed employees at their organization were using AI, while a Cisco readiness survey found 60% of leaders lacked confidence they could even identify which unapproved AI tools were in use. The gap isn't awareness anymore – most leaders know this is happening. It's governance: separate research found just over half of organizations have no formal policy covering employee use of external AI tools, and fewer than one in five coordinate AI governance between IT and the security team at all.

What's Actually At Risk

Data-loss research from Verizon's 2026 Data Breach Investigations Report, analyzing hundreds of thousands of uploads to generative AI tools, found source code was the single most common type of data employees pasted into unauthorized AI systems – ahead of images and structured business data. Where that risk concentrates depends on the department: engineering and technology teams tend to leak source code and credentials; healthcare and client-facing teams leak patient or customer records; legal and finance teams leak privileged or confidential documents. A generic, one-size-fits-all AI policy misses where your organization's actual exposure sits.

The deeper problem is that once sensitive information enters a public AI tool, your organization typically loses any meaningful control over it – no audit trail, no retention guarantee, no way to prove during a compliance review what was shared or with whom. For businesses in regulated industries (healthcare, finance, government contracting, faith-based institutions handling donor and member data), that loss of control can itself trigger a reportable incident, even if no attacker was ever involved.

The Upside Nobody Talks About

None of this means AI itself is the enemy – prohibition rarely works, and the data backs that up. Employees keep using these tools regardless of policy because they genuinely help. The organizations pulling ahead aren't the ones banning AI; they're the ones governing it. Ponemon Institute research found that organizations running a formal insider risk program – the kind that would catch a shadow AI incident early – avoided an average of $8.2 million in annual breach costs compared to those without one. Turn the page for the practical version: the specific steps that get you from ungoverned to governed without slowing your team down.

continued on page 4...

PROTECTING AI IN YOUR ORGANIZATION: THE PRACTICAL PLAYBOOK

When The Voice On The Phone Isn't Real

Shadow AI is about data leaking out. The mirror-image risk is synthetic content coming in - specifically, AI-generated voice and video used to impersonate people your employees already trust. In January 2024, engineering firm Arup lost $25.6 million after a finance employee joined a video call where every other participant, including the person who appeared to be the CFO, was an AI-generated fake. The FBI's most recent Internet Crime Report logged tens of thousands of business email compromise complaints totaling billions of dollars in losses, and industry researchers estimate AI-assisted deepfake elements were involved in roughly 40% of BEC incidents by early 2026, up from under 5% just two years earlier. Voice-cloning services now need as little as three seconds of reference audio - an earnings call, a conference keynote, a podcast appearance, a LinkedIn video - to produce a convincing clone. Yet by one industry estimate, roughly 80% of companies still have no deepfake response plan at all. The uncomfortable truth for executives specifically: the more publicly you speak, the more raw material exists to clone your voice.

Every recorded town hall, investor call, and conference talk is training data sitting in the open.

The Governed AI Playbook

You don't need a 200-page policy to close most of this gap. These are the moves that matter most, in order:

1. Name an owner.
One person not a committee - approves new AI tools and owns the approved-tools list. Without an owner, "someone should look into this" never happens.

2. Publish an approved tools list, not a ban.
List what's approved for company use, and be specific about what should never go into any AI tool: client PII, financial data, donor records, proprietary source code, anything covered by a signed NDA.

3. Write the one rule everyone can remember:
if you wouldn't post it on the company's public website, don't paste it into an AI prompt.

4. Require a human check on AI output
before it goes to a client, a regulator, or a financial system - AI-drafted work still needs a second set of trained eyes.

5. Build a verification callback rule for money and credentials.
Any request to move funds, change payroll details, or reset an executive's credentials - by phone, video, or email - gets confirmed through a second channel the requester didn't initiate, no exceptions, regardless of how urgent it sounds or how familiar the voice.

6. Treat executive audio and video as attack surface.
Limit how much unscripted, recorded executive speech sits publicly online where it isn't operationally necessary, and brief leadership on how easily their own voice can now be cloned.

7. Review the policy quarterly.
AI tools and threats are moving fast enough that an annual review is already out of date.

This Month's Threat Snapshot

Stat	Source
43% of breached organizations had a shadow AI incident in the past year	IBM Cost of a Data Breach Report, 2026
Just over half of organizations have no formal policy on employee AI use	Industry survey research, 2025-2026
~40% of business email compromise incidents involved deepfake elements by early 2026	Industry BEC research, 2026
A usable voice clone can be produced from about 3 seconds of audio	Industry voice-security research, 2026
$8.2M average breach cost avoided by organizations with a formal insider risk program	Ponemon Institute

These numbers aren't here to scare you away from AI - they're here to make the case that a few written rules and one accountable owner turn a real risk into a manageable one, without slowing your team down.



RANSOMWARE WATCH 2026: WHAT'S CHANGED FOR SMALL AND MID-SIZED BUSINESSES

Ransomware headlines still tend to feature household names, but the data tells a different story about who actually gets hit. Verizon's 2025 Data Breach Investigations Report found ransomware was a factor in 88% of breaches at small and midsize businesses, compared with 39% at larger organizations – small businesses aren't a secondary target, they're the primary one. Quarterly tracking from Q2 2026 shows small and mid-sized firms with under 200 employees and revenue below $25 million continuing to absorb the bulk of attacks, even as a separate 74% jump in attacks on billion-dollar enterprises made headlines.

The mechanics have also sped up. Attackers that once needed days to move from initial access to full network encryption now routinely finish in hours – some in under four hours. Multi-extortion has become standard: rather than relying on encryption alone, attackers steal data first, then threaten to publish it even when a business can restore from backup. That's a meaningful shift, because 96% of ransomware attacks now specifically target backup locations, meaning a backup strategy that only protects against encryption – and doesn't also protect against deletion or exfiltration – leaves a business exposed regardless of whether it eventually pays.

The financial stakes for a small business are severe enough to threaten survival outright. A 2025 Mastercard survey of more than 5,000 small-business owners found that roughly one in five who experienced a cyberattack went bankrupt or closed entirely. Even among businesses that recover, the road is long: fewer than half of encrypted organizations get back to normal operations within a week, and the businesses that hadn't been encrypted at all – because their backups held – recovered within a day at more than double the rate of those that had to rebuild from scratch.

Ransomware, By The Numbers

Stat	What It Means For You
88% of SMB breaches involve ransomware	Small businesses are the primary target, not an afterthought
96% of attacks target backups	Encryption-only backup protection is no longer enough
Under 4 hours	Fastest observed time from initial access to full encryption
~1 in 5	Small businesses that closed or went bankrupt after a cyberattack (Mastercard, 2025)
0.64	Share of ransomware victims who now refuse to pay, up from about half two years ago

The Root Cause, More Often Than Not

Across recent industry incident analysis, the most common contributing factor behind a small-business ransomware attack wasn't a sophisticated technical exploit – it was a known security gap the organization simply hadn't gotten around to closing: an unpatched system, a shared admin password, a missing MFA requirement. The second most common factor was a lack of in-house expertise or capacity to manage security alongside everything else running the business. Both are exactly what a right-sized compliance gap assessment and a managed security partnership are built to catch before an attacker does.



GRANT & COMPLIANCE SPOTLIGHT: WHERE THE SLCGP AND PILLAR ACT STAND

The federal State and Local Cybersecurity Grant Program (SLCGP) has distributed roughly $1 billion to state, local, and tribal governments since 2022, funding exactly the kind of planning, assessment, and workforce projects our Nonprofit & Government Sector Security Initiative supports. The program's active funding phase was on track to wind down in FY2026. The PILLAR Act, which would reauthorize and extend SLCGP funding through 2033, passed the House by voice vote in November 2025 and is currently sitting with the Senate Committee on Homeland Security and Governmental Affairs, awaiting further action.

What this means in practice: if your organization is a school district, municipality, county office, or a nonprofit that works alongside local government, cybersecurity spending tied to an approved State Cybersecurity Plan remains one of the more reliable ways to fund a gap assessment, an MFA rollout, or an incident response plan right now - while the current funding window is still open. We're continuing to track the reauthorization and will flag application windows as state administrative agencies open them.

SCHEDULE A FREE ASSESSMENT

Not sure where your organization stands? Our no-obligation Compliance Gap Assessment - now available in a right-sized format for nonprofits, schools, churches, and government offices - delivers a prioritized, plain-language finding report within 5 business days.

blackbeltsecure.com/audit

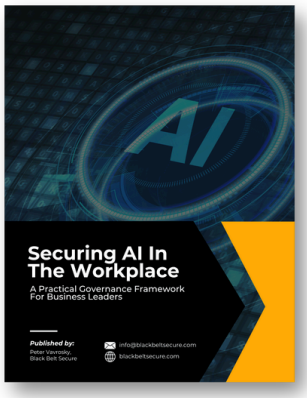
(469) 557-2007

FREE WHITEPAPER:

Securing AI In The Workplace

Our companion guide to workplace AI security covers the threat landscape of shadow AI, deepfake fraud, and autonomous agents; the NIST AI RMF and ISO/IEC 42001 frameworks; the new regulatory landscape in Texas, Colorado, Illinois, and the EU; and a practical 90-day governance roadmap any organization can implement without a full-time security team.

Claim Your FREE Copy At: blackbeltsecure.com/reports



TEN LOW-COST MOVES EVERY LEAN TEAM SHOULD MAKE THIS OCTOBER

BLACK BELT SECURE
Defend Today, Thrive Tomorrow.

Cybersecurity Awareness Month's Core 4 habits - strong passwords, multi-factor authentication, recognizing and reporting scams, and keeping software updated - still stop the majority of attacks. This year we're adding six more, built around what's actually changed: AI-written phishing, voice-cloned fraud, and shadow AI. None of these require a large budget or a dedicated security hire.

1. Turn on phishing-resistant MFA for every admin account tied to a system holding client, financial, or donor data - still the single control that would have stopped most of this year's breaches.

2. Inventory your AI tools. List every AI tool - approved or not - that anyone on your team is using, and name one owner responsible for the approved list.

3. Write your one AI data rule and put it somewhere everyone will see it: if you wouldn't post it publicly, don't paste it into an AI prompt.

4. Set a verification callback rule for money and credentials. Any request to move funds or reset access gets confirmed through a second channel - no exceptions, even for a familiar voice.

5. Patch internet-facing systems on a defined schedule, prioritizing anything holding customer, student, patient, or donor data.

6. Test your backups against deletion, not just encryption - most ransomware now targets backup locations directly, so a backup you can restore is only useful if the attacker couldn't also destroy it.

7. Write down your incident response plan - even one page naming who calls whom, in what order, is worth more than no plan at all.

8. Run one phishing simulation this month using current, AI-written lure examples rather than the dated ones most training libraries still rely on.

9. Limit unnecessary public executive audio and video where it isn't operationally required, and brief leadership on how easily a voice can now be cloned.

10. Talk to a vCISO before you talk to a grant officer - grant applications and cyber insurance renewals both score better when they cite a documented assessment and roadmap rather than good intentions.



Print this page and post it by the coffee machine. The Core 4 stop the everyday attacks; the other six close the gap AI opened this year.

BLACK BELT SECURE
Defend Today, Thrive Tomorrow.

DEFEND TODAY, THRIVE TOMORROW

Beyond protecting networks, Black Belt Secure is committed to strengthening the organizations that strengthen our communities. This October, that commitment takes a specific shape: free introductory workshops under our Nonprofit & Government Sector Security Initiative, running throughout Cybersecurity Awareness Month, with sessions tailored for church and ministry administrators, school-district technology directors, and city/county IT staff across North Texas and beyond. Alongside those workshops, we're running our first AI Governance briefing this fall – a plain-language session on the shadow AI and deepfake risks covered in this issue's special feature, built for executives and office managers rather than IT staff.



Questions? Email info@blackbeltsecure.com or visit blackbeltsecure.com for a free assessment.



October Workshop Calendar

Week	Focus
Week 1	Cybersecurity Awareness Month kickoff: the Core 4 habits, refreshed for 2026
Week 2	AI Governance briefing: shadow AI, deepfake fraud, and a starter policy for your team
Week 3	Nonprofit & Government Sector Security Initiative: right-sized assessments and grant readiness
Week 4	Incident response planning: building a one-page plan before you need it

Space in each session is limited to keep the conversation practical and specific to your organization. Ask your account team for exact dates, or reach out directly to reserve a seat.