LevelBlue

9,000+ IRs later: The 11 essential cybersecurity controls

Decades in incident response reveal battle-tested cybersecurity controls that minimize the attack surface, improve detection and response, reduce incident impact and losses, and build cyber resilience (with compliance mappings for easy implementation).

Plenty of organizations pass their audits and still get breached. We know, because we respond after it happens. The controls that satisfy a framework and the controls that actually stop an attacker are not on the same list.

That gap is the problem. Most security and risk leaders reach for a popular framework and let compliance dictate priorities, but compliance answers what an auditor wants to see, not what protects against an intrusion. With budgets finite and teams stretched, the real question is which controls produce the highest return when an attacker is in the environment.

Introducing the 11 essential cybersecurity controls

Drawing on more than 9,000 incident investigations, our responders identified the 11 essential cybersecurity controls: the highest-impact, field-proven practices that reduce the likelihood, dwell time, and cost of cyberattacks while strengthening response and recovery.

For each control, we include the pitfalls we see most often, the misconfigurations and gaps that surface again and again, plus a DFIR Perspective found in no other framework: what proper implementation means for your organization when an incident hits, whether that's a faster recovery, a smaller blast radius, or an attacker stopped before real damage is done.

A note on AI

AI is changing both sides of an incident. Attackers use it to write cleaner phishing, surface exploitable flaws faster, and run at a scale that used to require a team. None of this displaces the 11 controls. It raises the cost of getting them wrong, because the gap between exposure and exploitation keeps shrinking. The fundamentals matter more now, not less.

"After leading and overseeing thousands of cyber investigations worldwide, I've seen firsthand which defenses actually stop attackers and which ones routinely fail. The 11 essential cybersecurity controls distill those hard-earned lessons into clear, prioritized actions that any organization can take to meaningfully reduce their risk."

Devon Ackerman,
Global Services Leader, DFIR, LevelBlue

Make the most of the 11 essential cybersecurity controls

Prioritize resources

Focus investments on the frontline-proven controls that actually stop incidents.

Strengthen executive communication

Use the list to brief boards and justify security budgets with clarity.

Identify gaps

Benchmark current defenses against the essential controls to find weaknesses fast.

Align with insurers & risk managers

Demonstrate maturity and reduce risk exposure with controls supported by the insurance market.

Improve incident preparedness

Run tabletop exercises around controls to validate resilience under pressure.

Bridge compliance and security

Translate checklists into real-world resilience by prioritizing high-impact controls.

Table of contents

01

Phishing-resistant multi-factor authentication (MFA)

MFA requires users to verify their identity using two or more authentication factors, typically something they know (password), something they have (device or token), or something they are (biometric), before accessing systems or applications. Traditional MFA can be easily bypassed using stolen credentials or token theft tactics, so the adoption of phishing-resistant MFA (such as FIDO2) is strongly recommended. In our investigations, attackers routinely gain access through compromised credentials or easily bypass traditional MFA using common phishing kits.

Common Pitfalls

Compromised accounts often do not have MFA implemented at all, and where it is enabled, it's only traditional MFA, which is commonly bypassed by threat actors leveraging phishing kits or social engineering tactics to capture both credentials and one-time tokens. Other common issues include relying solely on SMS-based MFA, failing to enforce MFA for all remote access paths (e.g., VPN, RDP, cloud admin portals), or allowing exceptions in security strictness for certain executive or privileged users.

DFIR Perspective

The presence of properly enforced MFA often means attackers hit a wall after the initial compromise attempt. It gives incident responders fewer compromised accounts to investigate and helps isolate entry points. When MFA is weak or missing, lateral movement is easier, privilege escalation is faster, and the blast radius of successful compromises is broader. Properly implemented MFA also includes robust logging of threat actor activity, such as failed attempts and token reuse, allowing more precision and speed when pinpointing initial access or detecting adversary-in-the-middle phishing attempts.

MFA Compliance Mapping

Framework	Control ID / Name
CIS v8.1	6.3 – Require MFA for externally-exposed applications 6.4 – Require MFA for remote network access 6.5 – Require MFA for administrative access
NIST CSF 2.0	PR.AA-01 – Identities and credentials for authorized users, services, and hardware are managed PR.AA-03 – Users, services, and hardware are authenticated
NIST 800-171 Rev 2	3.5.1 – Identify system users and processes 3.5.2 – Authenticate identities 3.5.3 – Use MFA for local and network access

02

Endpoint detection and response (EDR) deployment

EDR solutions continuously monitor endpoint activity to detect, investigate, and respond to suspicious and anomalous behaviors like process injection, lateral movement, and credential harvesting across workstations, servers, and sometimes cloud workloads. EDR gives defenders near real-time and historical visibility into attacker behaviors. It's an essential tool for early detection and rapid containment, especially as adversaries use stealthy techniques that bypass traditional antivirus or perimeter defenses.

Common Pitfalls

Our incident data shows that most organizations have EDR in place, yet attackers still succeed in compromising their environments. The issue often isn't the tech, it's the investigative response. Most cases involved alert fatigue, lack of tuning, or missed detections due to weak triage processes. EDR agents may also be inconsistently deployed across endpoints, leaving gaps in visibility, particularly on legacy systems, cloud workloads, testing or development networks, or unmanaged devices (see control #7 for more details).

DFIR Perspective

EDR is often our fastest path to understanding attacker behavior across a networked environment from earliest activity to lateral movement and privilege escalation attempts. When deployed and monitored correctly, EDR drastically reduces time to containment, but when coverage is incomplete or alerts are ignored, it slows investigations and leaves blind spots.

EDR Compliance Mapping

Framework	Control ID / Name
CIS v8.1	10.1 – Deploy and maintain anti-malware software 10.7 – Use behavior-based anti-malware software 13.1 – Centralize security event alerting 13.7 – Deploy a host-based intrusion prevention solution
NIST CSF 2.0	DE.CM-09 – Computing hardware and software and their data are monitored to find potentially adverse events DE.AE-02 – Potentially adverse events are analyzed to better understand associated activities
NIST 800-171 Rev 2	3.14.2 – Provide protection from malicious code 3.14.6 – Monitor systems to detect attacks and indicators of potential attacks 3.14.7 – Identify unauthorized use of systems

03

Privileged access management (PAM)

PAM limits and controls access to high-value systems, accounts, and data by enforcing just-in-time privileges, credential vaulting, session monitoring, and access approvals for administrative or sensitive functions. Privileged accounts are prime targets, since they let attackers move laterally, disable defenses, or reach sensitive data. Effective PAM restricts the blast radius of a compromise and can stop escalation to domain-wide control. Applying Zero Trust to privileged access sharpens this further, verifying every elevation request rather than trusting standing administrative rights.

Common Pitfalls

Many organizations maintain overly generous admin rights across the network. Too many admin accounts, shared local admin accounts, admins with weak passwords and/or no MFA are common weaknesses. Others fail to rotate credentials, log privileged sessions, or audit who accessed what and when.

DFIR Perspective

When PAM is enforced, we see a clear audit trail of privileged activity. It allows for streamlined containment, slowing attacker escalation and lateral movement, thus giving us more time and visibility to detect and respond. Without PAM, attackers move faster (privilege escalations can be trivial), coverage gaps increase, and investigations take longer.

PAM Compliance Mapping

Framework	Control ID / Name
CIS v8.1	5.4 – Restrict administrator privileges to dedicated administrator accounts 6.8 – Define and maintain role-based access control
NIST CSF 2.0	PR.AA-05 – Access permissions, entitlements, and authorizations are defined, managed, enforced, and reviewed, incorporating least privilege and separation of duties PR.AA-01 – Identities and credentials are managed
NIST 800-171 Rev 2	3.1.5 – Employ the principle of least privilege 3.1.6 – Use non-privileged accounts for non-security functions 3.1.7 – Prevent non-privileged users from executing privileged functions

Stop
reacting.
Start
controling.

04

Centralized logging and log retention (SIEM or equivalent)

Collecting and preserving logs from the operating system, critical applications, network and edge appliances, such as firewalls and VPNs, into a security information and event management (SIEM) or data lake platform is crucial for incident response investigations as well as regulatory compliance and overall cyber resilience. SIEM-like platforms allow deep analysis, which fuels faster detections and forensic investigations. Without centralized logging, meaningful patterns are missed, forcing incident responders to pursue other log sources or limit the investigation scope to the available evidence.

Common Pitfalls

One of the biggest issues we see is gaps in coverage. Critical logs (like PowerShell, authentication events, or cloud admin actions) are often missing. Others collect logs but retain them for too short of a time. Alerting rules may also be poorly tuned, leading to alert fatigue or blind spots. Pricing is a valid concern when it comes to logging, as it can be cost prohibitive to ingest high volumes of data, so having an expert guide your logging and retention policy is encouraged.

DFIR Perspective

Centralized logs are our investigation backbone. They help us reconstruct timelines, identify Patient Zero, and correlate attacker actions across systems. When logs are missing or siloed, our work becomes slower and less precise. Proper log retention and visibility directly influence the speed and depth of our response.

Logging Compliance Mapping

Framework	Control ID / Name
CIS v8.1	8.2 – Collect audit logs 8.5 – Collect detailed audit logs 8.9 – Centralize audit logs 8.10 – Retain audit logs
NIST CSF 2.0	PR.PS-04 – Log records are generated and made available for continuous monitoring DE.AE-03 – Information is correlated from multiple sources
NIST 800-171 Rev 2	3.3.1 – Create and retain system audit logs and records 3.3.2 – Ensure actions can be traced to individual users 3.3.6 – Provide audit record reduction and report generation

05

Regular patching and vulnerability management

This control focuses on identifying, prioritizing, and remediating vulnerabilities through regular system patching, configuration management, and security updates across endpoints, servers, network appliances, and applications. Unpatched vulnerabilities are one of the most reliable entry points for attackers. In our investigations, exploited vulnerabilities consistently appear in the top three initial access methods. Timely patching reduces the attack surface and blocks opportunistic intrusions before they begin.

Common Pitfalls

Organizations often lack an up-to-date asset inventory, which leads to blind spots (see control #7). Patching cycles may be slow or prioritize CVE score over real-world exploitability. That timing gap is narrowing: AI-assisted tooling shortens the distance between a vulnerability going public and a working exploit appearing in the wild, so the patch window organizations have relied on is shrinking. Others rely on scans without timely remediation plans or can't track patch status across hybrid environments where one tool can't cover every surface.

DFIR Perspective

Mature patching means fewer preventable incidents and faster investigations, with less time lost to known, years-old CVEs. When patching is weak, attackers gain fast access and pivot quickly, sometimes exploiting public code long after patches were available.

Vulnerability Management Compliance Mapping

Framework	Control ID / Name
CIS v8.1	7.3 – Perform automated operating system patch management 7.4 – Perform automated application patch management 7.7 – Remediate detected vulnerabilities
NIST CSF 2.0	ID.RA-01 – Vulnerabilities in assets are identified, validated, and recorded PR.PS-02 – Software is maintained, replaced, and removed commensurate with risk ID.RA-06 – Risk responses are chosen, prioritized, planned, tracked, and communicated
NIST 800-171 Rev 2	3.11.2 – Scan for vulnerabilities 3.11.3 – Remediate vulnerabilities in accordance with risk assessments 3.14.1 – Identify, report, and correct system flaws

06

Email security filtering and phishing protection

This control includes technologies and processes to detect and block malicious emails, including phishing, malware, and spoofed domains. It typically combines secure email gateways, DNS filtering, DMARC enforcement, and user reporting mechanisms. Phishing and social engineering remain the leading initial intrusion vector we observe, accounting for most intrusions in recent LevelBlue TTP Briefings. Effective email security blocks many threats before they reach users and helps reduce credential theft, malware delivery, and business email compromise.

Common Pitfalls

Over time, email filters are often tuned too loosely, allowing advanced phishing emails to slip through. If a domain does not have DMARC configured at an enforcement level, or if SPF and DKIM are missing or misaligned with the domain's From address, the domains may remain vulnerable to spoofing. User reporting workflows are frequently broken or ignored, and phishing simulations are inconsistent or absent. Finally, email security is often siloed from broader incident detection systems.

DFIR Perspective

Strong email controls can eliminate entire intrusion paths. When in place, we see fewer successful phishing-based compromises, limiting both initial access and lateral movement. When missing, phishing gives attackers credentials, payload delivery opportunities, or persistence before they ever touch your endpoints. Coupled with control #4, email security solutions also generate valuable logs that assist investigators to trace the attack, confirm who was targeted, and pinpoint data at risk.

Email Security Filtering and Phishing Compliance Mapping

Framework	Control ID / Name
CIS v8.1	9.2 – Use DNS filtering services 9.5 – Implement DMARC 9.6 – Block unnecessary file types 9.7 – Deploy and maintain email server anti-malware protections 14.2 – Train workforce members to recognize social engineering attacks
NIST CSF 2.0	DE.CM-09 – Computing hardware and software and their data are monitored to find potentially adverse events PR.AT-01 – Personnel are provided awareness and training PR.DS-02 – Data-in-transit is protected
NIST 800-171 Rev 2	3.2.1 – Make personnel aware of security risks 3.2.2 – Train personnel to carry out their security duties 3.14.2 – Provide protection from malicious code

Fewer controls.
More impact.

07

Asset inventory and visibility (IT, OT, cloud)

A current, accurate inventory of all hardware, software, cloud services, and OT/IoT assets is foundational. You can't secure what you can't see, and attackers exploit unmanaged or unmonitored systems, especially in hybrid and cloud-heavy environments. Visibility underpins patching, access control, and incident response.

Common Pitfalls

Inventories are often incomplete, manually maintained, or siloed across IT, InfoSec, and development teams. Shadow IT creates preventable blind spots, and increasingly that shadow extends to AI: unsanctioned AI tools, browser extensions, and autonomous agents that now reach corporate data without ever appearing in an inventory, an asset class most discovery tools were never built to find. Many organizations also fail to tag assets by criticality or owner, making it hard to prioritize response when unauthorized access occurs.

DFIR Perspective

Knowing what systems exist and where they live drastically improves scoping and containment. When inventories are incomplete, investigators lose time chasing endpoints and cloud resources the organization didn't know it had. A known AI footprint lets you scope in minutes what an unsanctioned tool or agent could have touched; without it, that footprint becomes the part of the investigation no one can bound.

Asset Inventory Compliance Mapping

Framework	Control ID / Name
CIS v8.1	1.1 – Establish and maintain detailed enterprise asset inventory 1.2 – Address unauthorized assets 2.1 – Establish and maintain a software inventory
NIST CSF 2.0	ID.AM-01 – Inventories of hardware managed by the organization are maintained ID.AM-02 – Inventories of software, services, and systems are maintained ID.AM-05 – Assets are prioritized based on classification, criticality, resources, and impact on the mission
NIST 800-171 Rev 2	3.4.1 – Establish and maintain baseline configurations and inventories 3.4.2 – Establish and enforce security configuration settings

08

Network segmentation and access controls

Segmenting networks based on sensitivity, business function, or trust level, and enforcing access controls between them makes it harder for attackers to move laterally. Commonly achieved with firewalls or cloud-native segmentation, it limits blast radius and slows adversaries. Modern implementations extend this with Zero Trust, replacing implicit network trust with continuous identity verification (see control #3), and with secure access service edge (SASE) architectures that enforce segmentation and access policy across branch, cloud, and remote users rather than at the perimeter alone.

Common Pitfalls

Many networks remain overly permissive, with broad internal access and little segmentation between critical systems and general user networks. Firewall rules often “allow any” by default. In cloud environments, overly broad IAM roles and security group misconfigurations are common. Flat networks without Zero Trust controls let a single compromised credential reach way too far, and VPN access that grants broad network reach rather than per-application access is a common version of this gap. Segmentation often focuses on the perimeter while neglecting internal lateral movement paths.

DFIR Perspective

Where Zero Trust and segmentation are enforced, an attacker may compromise a user machine but cannot reach segmented zones for finance, HR, IT management, or backup environments. That is the difference between containing an incident to a few systems and isolating an entire network during response.

Network Segmentation and Access Compliance Mapping

Framework	Control ID / Name
CIS v8.1	3.12 – Segment data processing and storage based on sensitivity 12.2 – Establish and maintain a secure network architecture 13.4 – Perform traffic filtering between network segments
NIST CSF 2.0	IPR.IR-01 – Networks and environments are protected from unauthorized logical access and usage PR.AA-05 – Access permissions and authorizations are defined, managed, enforced, and reviewed
NIST 800-171 Rev 2	3.1.3 – Control the flow of CUI in accordance with approved authorizations 3.13.1 – Monitor, control, and protect communications at system boundaries 3.13.5 – Implement subnetworks for publicly accessible system components

09

Incident response plans (IRP) and tabletop exercises

This control involves developing a formal incident response plan (IRP) that outlines roles, responsibilities, communication protocols, and technical procedures for responding to cyber incidents. Tabletop exercises are then structured to test and refine IRPs in simulated attack scenarios. In a crisis, lack of coordinated response costs time. A well-practiced IR plan improves coordination, speeds up containment, and reduces business impact. Tabletop exercises uncover process gaps, validate assumptions, and strengthen relationships across legal, IT, PR, and leadership teams before an actual intrusion occurs.

Common Pitfalls

Many plans are outdated, untested, or written for compliance checkboxes. They're often too technical or too vague, with no clear ownership of decisions like whether to notify regulators or pull the plug on a system. Tabletop exercises, if done at all, may be overly scripted or exclude key stakeholders like legal, PR, or executives. Lessons learned are not often captured or acted on.

DFIR Perspective

Organizations that have practiced incident response at least once respond faster and more cohesively. We spend less time waiting for access or approval and more time containing the intrusion event. When there's no plan or the team hasn't rehearsed, response is chaotic, and that chaos compounds the damage.

Incident Response Planning Compliance Mapping

Framework	Control ID / Name
CIS v8.1	17.1 – Designate personnel to manage incident handling 17.4 – Establish and maintain an incident response process 17.7 – Conduct routine incident response exercises
NIST CSF 2.0	RS.MA-01 – The incident response plan is executed in coordination with relevant third parties once an incident is declared ID.IM-02 – Improvements are identified from security tests and exercises
NIST 800-171 Rev 2	3.6.1 – Establish an operational incident-handling capability 3.6.2 – Track, document, and report incidents 3.6.3 – Test the organizational incident response capability

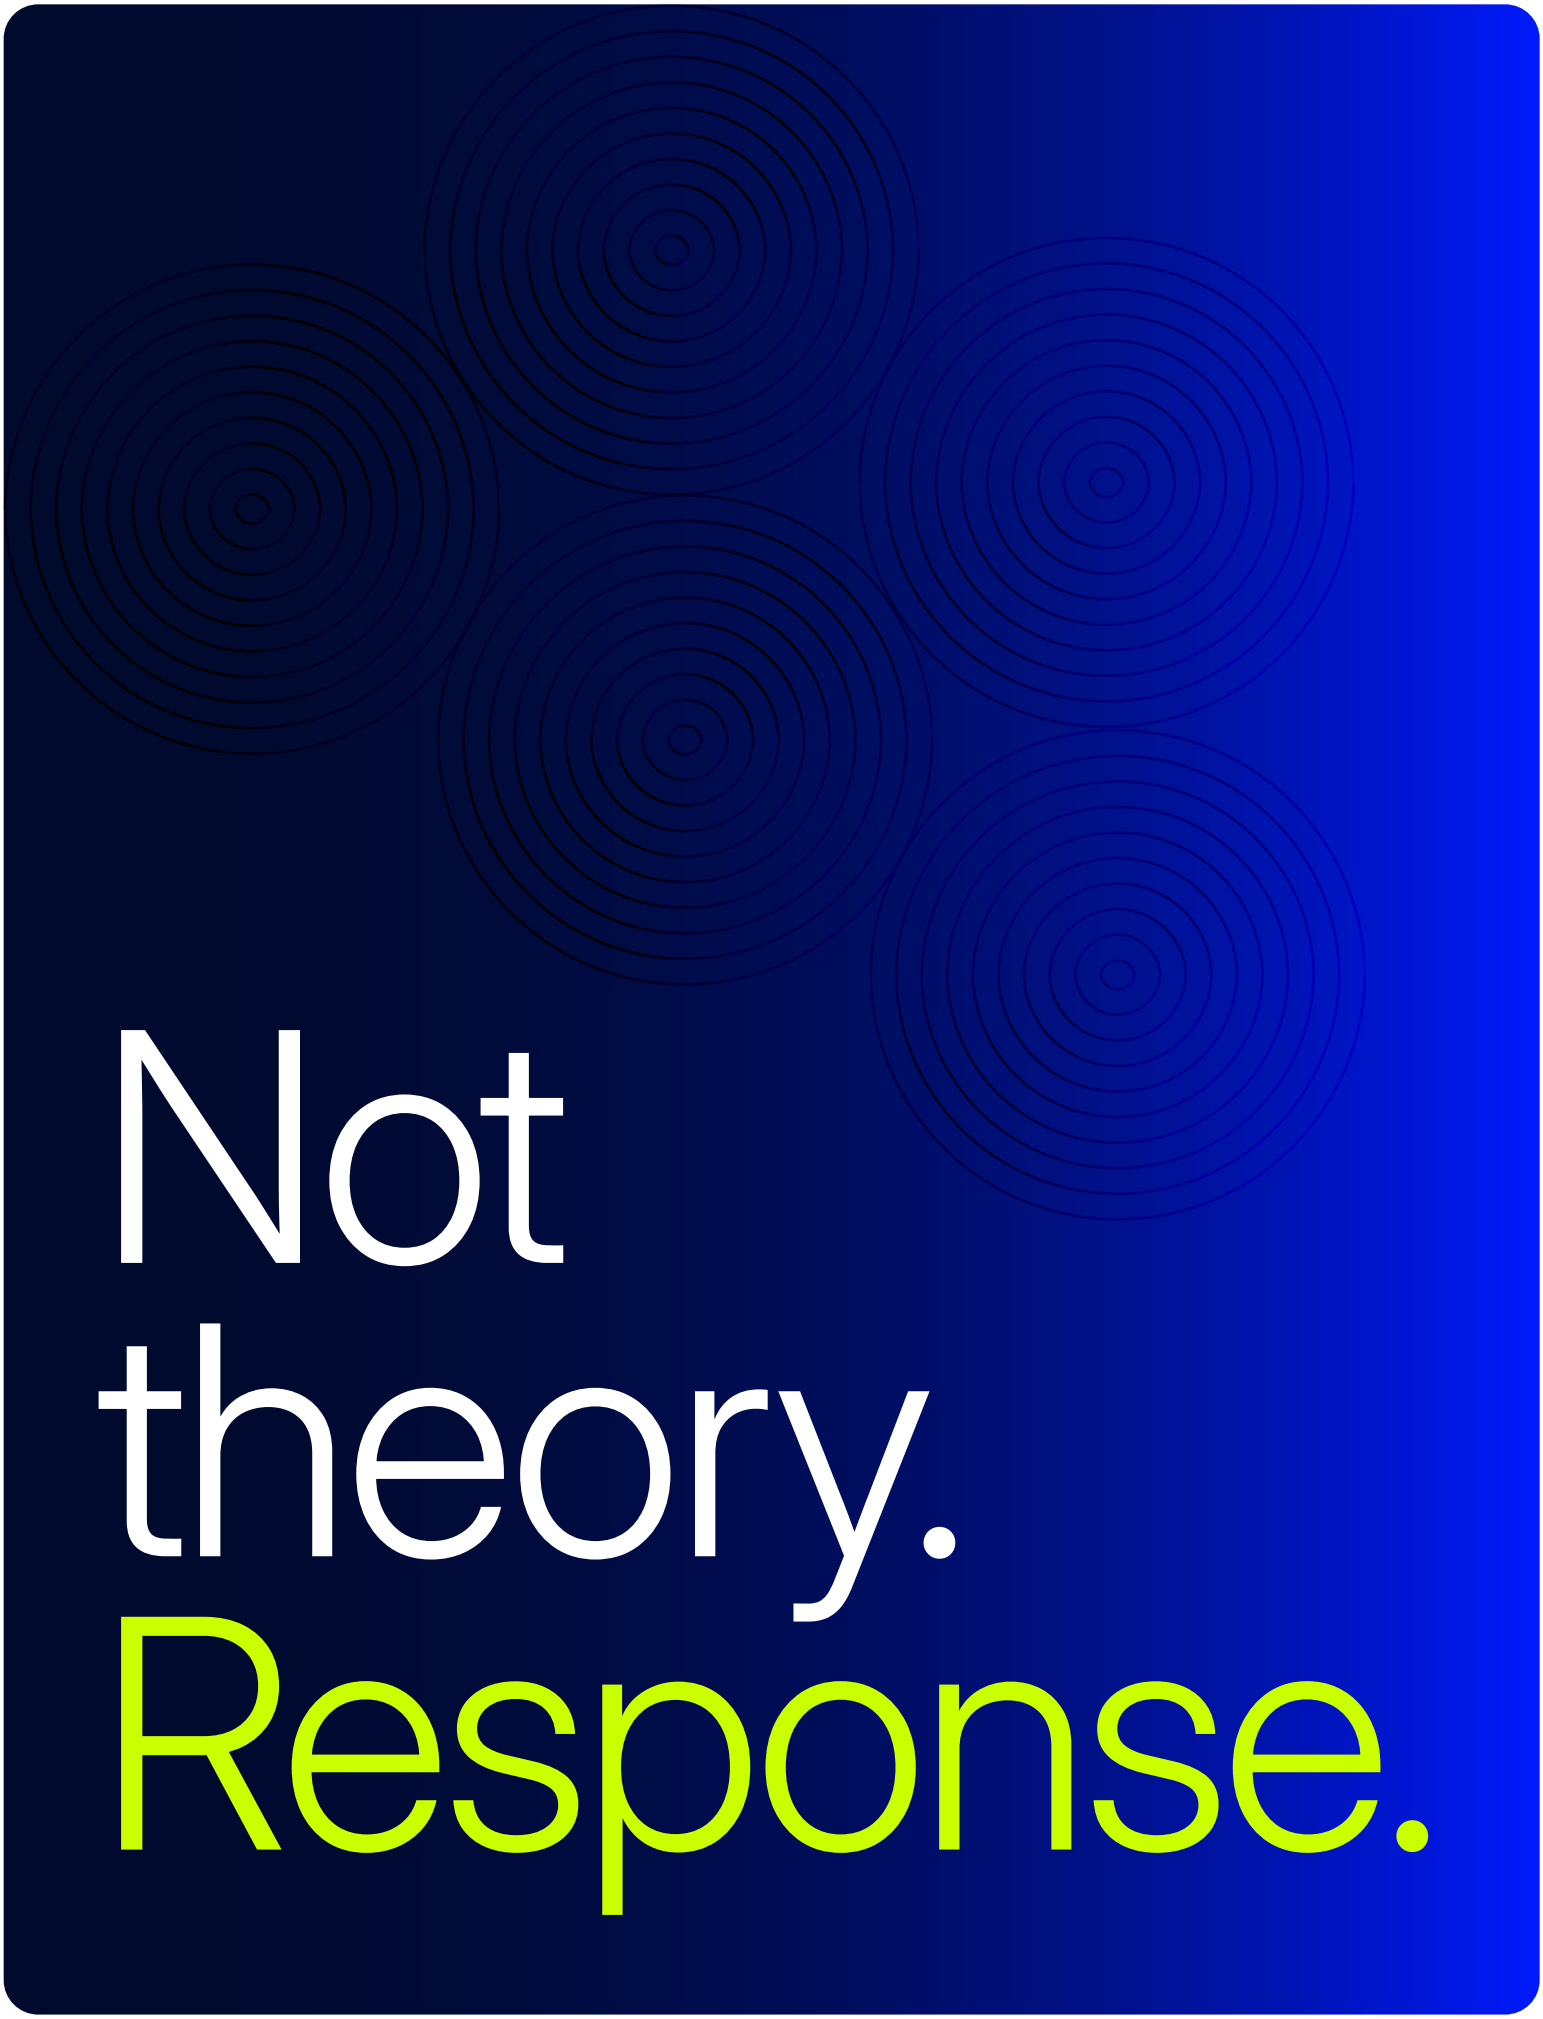
Not
theory.
Response.

10

Data classification and structured data management

Identifying and categorizing data based on sensitivity, value, and regulatory requirements, and implementing controls to manage, protect, and govern data throughout its lifecycle helps prioritize protections, drive access controls (see control #8), and inform incident response and regulatory obligations. Attackers quickly find file servers, databases, and SharePoint folders containing intellectual property, financials, PII, etc. Knowing where that data lives and how it's protected is critical to reducing risk.

Common Pitfalls

Many organizations don't know where their sensitive data resides, or data classification efforts may be manual, static, or disconnected from actual security controls. Structured data (e.g., databases) may be governed, but unstructured data (e.g., file shares, cloud storage) often lack oversight. Tagging, encryption, and retention policies may be missing or inconsistently applied.

DFIR Perspective

When data is well-classified, we can quickly understand what's at risk during an incident and advise on necessary notifications and legal exposure. Without it, IR teams often have to spend valuable hours figuring out whether attackers accessed regulated or high-value data. To make it more realistic: if you know an attacker only encrypted the server where you store marketing collateral, would you pay the ransom?

Data Classification Compliance Mapping

Framework	Control ID / Name
CIS v8.1	3.1 – Establish and maintain a data management process 3.2 – Establish and maintain a data inventory 3.7 – Establish and maintain a data classification scheme
NIST CSF 2.0	ID.AM-07 – Inventories of data and corresponding metadata for designated data types are maintained PR.DS-01 – The confidentiality, integrity, and availability of data-at-rest are protected
NIST 800-171 Rev 2	3.8.1 – Protect CUI on system media 3.8.4 – Mark media with necessary CUI markings 3.1.22 – Control CUI posted or processed on publicly accessible systems

11

Offline, segmented, and tested backups with RTO

In the ideal implementation of immutable backups, a data copy exists that cannot be modified, deleted, or encrypted by anyone, even an administrator or a compromised system, for a set retention period. While having immutable, segmented backups is a critical last line of defense in ransomware scenarios, it should not be the first or only focus. Backups support recovery after an incident, but the other 10 controls are more impactful in preventing and containing attacks before they cause damage.

Common Pitfalls

Backups are often online and accessible from the same network as production systems, making them just as vulnerable during an attack. In some cases, backups exist but haven't been tested or fail to meet the actual recovery time objective (RTO) needs. Organizations may focus on backup frequency but not test how fast they can actually restore (RTO). In addition, recovery processes are rarely rehearsed under data encryption or destruction events like ransomware or system-wide compromise conditions.

DFIR Perspective

We frequently see cases where recovery is delayed or impossible because backups were encrypted or deleted by the attacker. It removes a lot of leverage from ransomware actors and gives leadership more options. Recovery time and effectiveness are dramatically better in those environments.

Backup Compliance Mapping

Framework	Control ID / Name
CIS v8.1	11.2 – Perform automated backups 11.3 – Protect recovery data 11.4 – Establish and maintain an isolated instance of recovery data 11.5 – Test data recovery
NIST CSF 2.0	PR.DS-11 – Backups of data are created, protected, maintained, and tested RC.RP-03 – The integrity of backups and other restoration assets is verified before using them for restoration RC.RP-01 – The recovery portion of the incident response plan is executed
NIST 800-171 Rev 2	3.8.9 – Protect the confidentiality of backup CUI at storage locations *800-171 Rev 2 has thin backup/recovery coverage; the contingency-planning family lives in 800-53. 3.8.9 is the closest direct fit.

LevelBlue

Take the next step in cyber resilience.

Our experts have responded to thousands of incidents and can help you benchmark your security program against the essential controls.

Schedule a 1:1 session with a LevelBlue expert today at levelblue.com/irr