THE COMPLETE

Cybersecurity GUIDE

BEMI TECH SOLUTIONS

From Fundamentals to Advanced Domains

Foundations
CIA Triad, Kill Chain

Network Sec
Firewalls, VPN, Zero Trust

IAM
MFA, RBAC, PAM, SSO

AppSec
OWASP, DevSecOps

Cloud Security
AWS, Azure, CSPM

SOC & SIEM
Threat Hunting, SOAR

Incident Resp.
DFIR, Forensics

Pen Testing
Red Team, OSINT

AI & Security
AI Threats, MLSecOps

11	100+	8	$10.5T
Domains	Tools Listed	Frameworks	2025 Cost

Ransford Slater

Cybersecurity Educator | Bemi Tech Solutions | 2026 Edition

TABLE OF CONTENTS

01 Cybersecurity Foundations

CIA Triad | Threat Landscape | Kill Chain | MITRE ATT&CK

Cybersecurity protects digital systems, networks, and data from attacks, damage, and unauthorized access. As interconnectedness grows, the attack surface expands — making cybersecurity one of the most critical disciplines of our era. Every organization, from solo freelancers to nation-states, faces real and evolving threats.

THE CIA TRIAD

Every security decision maps back to three core pillars — Confidentiality, Integrity, and Availability. A breach always violates at least one, often all three. Understanding which pillar is under threat helps security teams prioritize the right controls.

Confidentiality

Data accessible only to authorized parties. Enforced via encryption, access controls, and strong authentication.

Integrity

Data is accurate and unaltered by unauthorized parties. Enforced via hashing (SHA-256) and digital signatures.

Availability

Systems accessible when needed. Enforced via redundancy, load balancing, backups, and DDoS mitigation.

Why It Matters

Every policy, control, and tool traces back to protecting one or more of these three properties. Always ask: which pillar is at risk?

THE THREAT LANDSCAPE

Cybercrime is the fastest-growing criminal enterprise on earth. The numbers are staggering and accelerating every year. Understanding the landscape is the first step in building defenses proportional to the actual risk.

$8T	43%	39s	$4.45M
2023 Cybercrime Cost	Attacks on SMBs	Attack Frequency	Avg. Breach Cost

MAJOR THREAT CATEGORIES

- Malware — viruses, worms, trojans, ransomware, spyware designed to disrupt or steal
- Phishing & Social Engineering — responsible for 90%+ of data breaches worldwide
- Ransomware — avg. ransom demand $1.54M in 2023; targets hospitals, schools, governments
- Supply Chain Attacks — compromise one vendor to reach thousands of downstream customers
- Insider Threats — 34% of breaches involve internal actors (accidental or malicious)
- Zero-Day Exploits — unknown vulnerabilities with no available patch at time of attack
- APTs (Advanced Persistent Threats) — state-sponsored, long-term intrusion campaigns

THE CYBER KILL CHAIN

Developed by Lockheed Martin, the Kill Chain describes the 7 stages of a cyberattack. Defenders can disrupt the attack at any stage — the earlier the better. Breaking even one link stops the attack entirely.

- Reconnaissance — gather intelligence on target via OSINT, scanning, social media
- Weaponization — build malicious payload (exploit + backdoor or dropper)
- Delivery — send payload via phishing email, USB, watering hole, or web exploit
- Exploitation — trigger vulnerability to execute attacker-controlled code on target
- Installation — install persistent malware (RAT, rootkit, or scheduled task)
- C2 (Command & Control) — establish covert communication channel to malware
- Actions on Objectives — exfiltrate data, encrypt for ransom, or destroy systems

MITRE ATT&CK FRAMEWORK

MITRE ATT&CK is a globally accessible knowledge base of adversary tactics and techniques based on real-world observations. It is the de facto standard for threat intelligence, detection engineering, and red team operations.

ATT&CK KEY CONCEPTS

- Tactics — the 'why' of an attack (e.g., Initial Access, Persistence, Exfiltration, Impact)
- Techniques — the 'how' (e.g., Spearphishing Attachment T1566.001 under Initial Access)
- Sub-techniques — more granular implementations of a technique
- Procedures — specific real-world implementations observed from named threat actors
- ATT&CK Navigator — visual tool to map detection coverage across all 200+ techniques
- Use case: map your SIEM alert rules to ATT&CK to identify blind spots

02 Network Security

Firewalls | IDS/IPS | VPN | Zero Trust | Network Architecture

Network security protects data in transit and controls what enters and exits your environment. It is the most mature cybersecurity domain — yet attackers continue to find new paths through it. A layered network defense (defense in depth) is essential for any organization.

CORE NETWORK CONTROLS

Firewall

Filters traffic by rules. Next-Gen (NGFW) adds deep packet inspection, app awareness, TLS inspection, and user identity.

IDS / IPS

IDS alerts on suspicious traffic. IPS actively blocks it inline. Both use signature-based and anomaly-based detection.

VPN

Encrypts traffic over public networks. IPSec for site-to-site tunnels; SSL/TLS for remote access. Essential for remote work security.

Network Segmentation

Divides the network into trust zones. Contains lateral movement — a breach in one zone cannot freely reach all others.

WAF

Web Application Firewall filters malicious HTTP/S traffic targeting web apps. Blocks OWASP Top 10 attacks including SQLi and XSS.

NAC

Network Access Control enforces posture checks before admitting devices. Unpatched or non-compliant devices are quarantined automatically.

ZERO TRUST ARCHITECTURE

Zero Trust rejects the old assumption that anything inside the network perimeter is safe. It treats every user, device, and request as potentially hostile — verifying identity and context continuously, for every access attempt.

ZERO TRUST PRINCIPLES

- Verify explicitly — authenticate and authorize every request using all available contextual data
- Use least privilege — limit access scope and duration, enforce Just-In-Time (JIT) access
- Assume breach — design systems as if attackers are already inside the network
- Micro-segmentation — isolate workloads at granular level to limit blast radius of any breach
- Continuous monitoring — log, inspect, and analyze all traffic including internal east-west flows
- Device health verification — ensure devices meet compliance posture before granting access

COMMON NETWORK ATTACKS

- DNS Poisoning — corrupts resolver cache to redirect users to attacker-controlled sites
- ARP Spoofing — maps attacker MAC to legitimate IP, enabling Man-in-the-Middle attacks
- DDoS — floods network bandwidth or app resources to make services unavailable
- VLAN Hopping — exploits misconfigured trunk ports to gain access to restricted VLANs
- SSL Stripping — downgrades HTTPS connections to HTTP to expose plaintext credentials
- Rogue AP — attacker deploys a fake wireless access point to intercept user traffic
- BGP Hijacking — corrupts routing tables to redirect internet traffic through attacker AS

NETWORK SECURITY TOOLS

Wireshark **Nmap** **Snort** **Suricata** **pfSense** **Cisco ASA**
Palo Alto NGFW **Zeek** **tcpdump** **NetFlow** **OpenVPN** **WireGuard**
Tshark **ntopng**

03 Identity & Access Management

Authentication | MFA | SSO | RBAC | PAM | Directories

Identity is the new perimeter. As organizations shift to cloud and remote work, traditional network boundaries disappear. Controlling WHO can access WHAT, WHEN, and HOW has become the most critical security control in modern architecture.

AUTHENTICATION FACTORS

Something You Know

Passwords, PINs, security questions. Weakest factor — stolen in phishing, guessed via brute force, or found in breach dumps.

Something You Have

Hardware tokens (YubiKey), TOTP apps (Google Authenticator), SMS OTP. Significantly raises the bar for attackers.

Something You Are

Biometrics — fingerprint, face, iris, voice. Strongest factor but requires specialized hardware and raises privacy concerns.

Multi-Factor Auth (MFA)

Combining two or more factors. Even if a password is breached, MFA blocks 99.9% of account takeover attacks per Microsoft.

ACCESS CONTROL MODELS

ACCESS CONTROL FRAMEWORKS

- RBAC (Role-Based) — access based on job role; most common enterprise model; easy to audit
- ABAC (Attribute-Based) — dynamic access based on user, resource, and environment attributes
- DAC (Discretionary) — resource owner controls access; flexible but prone to permission creep
- MAC (Mandatory) — access based on classification labels; used in government/military contexts
- PBAC (Policy-Based) — combines RBAC + ABAC with centralized policy engine for fine-grained control

PRIVILEGED ACCESS MANAGEMENT (PAM)

Privileged accounts — admins, service accounts, root users — are the most targeted credentials in any environment. PAM solutions vault these credentials, enforce just-in-time access, and record all privileged sessions for audit.

PAM CORE CAPABILITIES

- Credential vaulting — store privileged passwords in encrypted vault; rotate automatically
- Session recording — record video and keystrokes of every privileged session for forensics
- Just-in-time access — grant elevated privileges only for the duration of an approved task
- Least privilege enforcement — restrict admin rights to only what the task requires
- Multi-party authorization — require approval from a second administrator for sensitive actions

SINGLE SIGN-ON (SSO) AND FEDERATION

SSO allows users to authenticate once and access multiple applications without re-entering credentials. Federation extends SSO across organizational boundaries using standards like SAML, OAuth 2.0, and OpenID Connect (OIDC).

- SAML 2.0 — XML-based, browser-focused; widely used for enterprise SaaS federation
- OAuth 2.0 — authorization framework; delegates access without sharing credentials
- OpenID Connect — identity layer on top of OAuth 2.0; provides ID tokens (JWT)
- LDAP / Active Directory — on-premises directory; central source of truth for user identities

IAM TOOLS

Okta | **Azure AD** | **AWS IAM** | **Google Workspace** | **CyberArk**

HashiCorp Vault | **BeyondTrust** | **SailPoint** | **Ping Identity** | **Duo**

JumpCloud | **ForgeRock**

04 Application Security

OWASP Top 10 | DevSecOps | SAST/DAST | Secure SDLC

Applications are the largest attack surface in modern environments. Every API endpoint, web form, and mobile app is a potential entry point. Application security (AppSec) embeds security throughout the development lifecycle — not as an afterthought at the end.

OWASP TOP 10 (2021)

The OWASP Top 10 is the most cited reference for web application vulnerabilities. Understanding each category is essential for developers, testers, and architects.

OWASP TOP 10 CATEGORIES

- A01 Broken Access Control — most common; users accessing data or functions above their privilege
- A02 Cryptographic Failures — sensitive data exposed due to weak or missing encryption
- A03 Injection — SQL, NoSQL, command injection via untrusted data in interpreter queries
- A04 Insecure Design — missing threat modeling and security design patterns from the start
- A05 Security Misconfiguration — default creds, open cloud storage, unnecessary features enabled
- A06 Vulnerable & Outdated Components — using libraries with known CVEs
- A07 Auth Failures — broken auth, session fixation, credential stuffing vulnerabilities
- A08 Software & Data Integrity Failures — CI/CD pipeline compromise, unsigned updates
- A09 Security Logging Failures — insufficient logging to detect or reconstruct incidents
- A10 SSRF (Server-Side Request Forgery) — server tricked into making requests to internal services

SECURE SDLC

Security must be integrated into every phase of the Software Development Lifecycle — not bolted on at the end. The cost of fixing a vulnerability increases 10x–100x with each phase it is allowed to progress through.

Design Phase

Threat modeling (STRIDE), security requirements gathering, attack surface analysis, security architecture review.

Development Phase

Secure coding standards, code reviews, developer security training, SAST (static analysis) in the IDE.

Testing Phase

DAST (dynamic analysis) against running app, penetration testing, dependency scanning for CVEs, fuzz testing.

Deployment Phase

IAST in staging, secrets scanning in CI/CD, signed builds, immutable containers, runtime protection (RASP).

DEVSECOPS

DevSecOps shifts security left by integrating it directly into the CI/CD pipeline. Every code commit triggers automated security scans before code is ever deployed. Security becomes part of the developer workflow, not a separate gate.

DEVSECOPS PIPELINE CONTROLS

- Pre-commit hooks — SAST, secrets detection, linting before code enters repository
- CI pipeline — SCA (software composition analysis), container image scanning, IaC scanning
- CD pipeline — DAST against staging, compliance as code, policy enforcement gates
- Runtime — RASP, WAF, anomaly detection, continuous vulnerability management
- Shift-left: developers find and fix 10x cheaper vs. finding in production

APPSEC TOOLS

OWASP ZAP | **Burp Suite** | **Snyk** | **SonarQube** | **Semgrep** | **Checkmarx**

Veracode | **Trivy** | **GitLeaks** | **TruffleHog** | **Bandit** | **Dependabot**

OWASP Dependency Check

05 Cloud Security

Shared Responsibility | CSPM | Containers | Cloud Tools

Cloud adoption has fundamentally changed the security landscape. Misconfigured S3 buckets, over-permissioned IAM roles, and exposed APIs have replaced perimeter breaches as the leading cause of large-scale data exposures. Cloud security requires a new mindset.

SHARED RESPONSIBILITY MODEL

Every cloud provider defines a shared responsibility model — what they secure (of the cloud) and what you must secure (in the cloud). Misunderstanding this boundary is the single most common cause of cloud breaches.

SaaS — Applications
Customer manages: Data, Users

PaaS — Platform & Runtime
Customer manages: Data, Apps

IaaS — Infrastructure / VMs
Customer manages: Data, Apps, OS, Network

On-Premises
Customer manages: Everything

CLOUD SECURITY ARCHITECTURE

CSPM

Cloud Security Posture Management — continuously scans cloud configs for misconfigurations, policy violations, and compliance drift.

CWPP

Cloud Workload Protection Platform — secures VMs, containers, and serverless functions at runtime against exploits and malware.

CASB

Cloud Access Security Broker — visibility and control over SaaS app usage, shadow IT detection, DLP for cloud-stored data.

CNAPP

Cloud-Native Application Protection Platform — unified platform combining CSPM + CWPP + CIEM + IaC scanning in one tool.

CONTAINER & KUBERNETES SECURITY

Containers introduce new attack vectors: vulnerable base images, over-privileged pods, exposed dashboards, and supply chain risks in registries. Kubernetes clusters require hardening beyond default configurations.

CONTAINER SECURITY CONTROLS

- Use minimal base images (distroless / Alpine) — fewer packages = smaller attack surface
- Scan images for CVEs in CI pipeline before pushing to registry (Trivy, Grype, Snyk)
- Never run containers as root — use non-root user in Dockerfile
- Network policies — restrict pod-to-pod communication to least-privilege
- RBAC on Kubernetes API — limit who can create, delete, or exec into pods
- Secrets management — never bake secrets into images; use Vault or cloud secret managers
- Runtime security — Falco, Sysdig; detect abnormal syscalls (shell spawned in container)

CLOUD SECURITY TOOLS

AWS Security Hub | **Azure Defender** | **GCP Security Command Center**

Prisma Cloud | **Wiz** | **Orca Security** | **Lacework** | **Aqua Security**

Falco | **Trivy** | **Prowler** | **ScoutSuite** | **Checkov** | **Terraform Sentinel**

06 Security Operations Center

SIEM | SOAR | Threat Hunting | XDR | Monitoring

The Security Operations Center (SOC) is the nerve center of an organization's security posture. It is responsible for continuous monitoring, detection, investigation, and response to cybersecurity threats — 24 hours a day, 7 days a week.

SIEM — SECURITY INFORMATION & EVENT MANAGEMENT

SIEM aggregates logs from across the environment — firewalls, endpoints, servers, applications, cloud — and correlates them into actionable alerts. Modern SIEMs use machine learning to detect anomalies that rule-based systems miss.

Log Collection

Ingest logs from all sources: endpoints, network, cloud, identity, and application layers via agents and syslog.

Correlation Rules

Detect attack patterns across multiple data sources. Example: failed login + successful login from different geographies.

Alerting & Triage

Generate alerts ranked by severity. SOC analysts triage Level 1 (high volume, lower fidelity) through Level 3 (complex, rare).

UEBA

User & Entity Behavior Analytics — baseline normal behavior and flag deviations: unusual login times, mass data downloads, lateral movement.

SOAR — SECURITY ORCHESTRATION, AUTOMATION & RESPONSE

SOAR platforms automate repetitive SOC tasks — IP lookups, ticket creation, endpoint isolation — allowing analysts to focus on complex investigations. A well-built SOAR playbook can reduce mean time to respond (MTTR) from hours to minutes.

COMMON SOAR PLAYBOOKS

- Phishing triage — extract URLs, check VirusTotal, block sender, notify user, close ticket
- Malware alert — isolate endpoint via EDR API, pull memory dump, alert IR team automatically
- Brute force — lockout account after threshold, alert owner, create JIRA ticket, escalate
- Cloud alert — pull CloudTrail logs, check IAM changes, correlate with identity source
- Vulnerability remediation — assign CVE to owner via ticketing, track to closure, report SLA

THREAT HUNTING

Threat hunting is the proactive search for hidden threats that have evaded automated detection. Unlike reactive incident response, hunters start with a hypothesis — 'what would persistence look like in our environment?' — and prove or disprove it using raw data.

- Hypothesis-driven — based on threat intelligence, ATT&CK techniques, or intuition
- Intel-driven — IOCs from feeds, ISAC sharing, or published threat actor TTPs
- Situational awareness — unusual process trees, abnormal network connections, new scheduled tasks
- Tools: Velociraptor, OSQuery, Elastic, Splunk, Microsoft Sentinel — query raw telemetry directly

SOC TOOLING

Splunk | **Microsoft Sentinel** | **Elastic SIEM** | **IBM QRadar** | **LogRhythm**

CrowdStrike Falcon | **SentinelOne** | **Palo Alto XSIAM** | **Tines**

Cortex XSOAR | **Velociraptor** | **OSQuery** | **TheHive** | **MISP** | **VirusTotal**

07

Incident Response & Digital Forensics

IR Lifecycle | DFIR | Memory Forensics | Evidence Handling

Incident Response (IR) is the structured approach to handling cybersecurity events — from detection through containment, eradication, recovery, and post-incident review. Digital Forensics (DFIR) combines both disciplines to investigate breaches and preserve evidence for legal proceedings.

IR LIFECYCLE — 6 PHASES

NIST SP 800-61 defines the IR lifecycle used by most organizations and service providers. Each phase builds on the last — skipping any phase leads to incomplete remediation and repeated incidents.

1	Preparation	2	Identification	3	Containment
4	Eradication	5	Recovery	6	Lessons Learned

PHASE DETAILS

- 1. Preparation — IR plan, playbooks, SIEM tuning, tabletop exercises, retainer contracts
- 2. Identification — confirm the incident, classify severity, determine scope and initial IOCs
- 3. Containment — isolate affected systems, block attacker C2, preserve evidence (image first!)
- 4. Eradication — remove malware, close access vectors, patch vulnerabilities exploited
- 5. Recovery — restore systems from clean backups, monitor closely for re-infection
- 6. Lessons Learned — post-mortem within 2 weeks; update detections, playbooks, and controls

DIGITAL FORENSICS

Digital forensics involves the collection, preservation, analysis, and presentation of digital evidence. The chain of custody must be maintained — evidence handling errors can render findings inadmissible in court.

Disk Forensics

Acquire full disk image (dd, FTK Imager). Analyze file system artifacts, deleted files, MFT entries, browser history, and registry.

Memory Forensics

Capture volatile RAM (WinPmem, LiME). Extract running processes, network connections, injected shellcode, and decrypted credentials.

Network Forensics

Analyze PCAP files, NetFlow logs, proxy logs. Reconstruct attacker communication, data exfiltration paths, and C2 channels.

Log Analysis

Windows Event Logs, Linux syslog, application logs. Key events: 4624 (logon), 4688 (process create), 4720 (account create).

DFIR TOOLS

Volatility | Autopsy | FTK Imager | Wireshark | Plaso | KAPE | Rekall | Velociraptor | TheHive | Cortex | Yara | CyberChef | OSForensics | Cuckoo Sandbox | AnyRun

08 Ethical Hacking & Penetration Testing

Red Team | Methodology | OSINT | Reporting

Penetration testing (pen testing) simulates real-world attacks against systems, networks, and applications to identify exploitable vulnerabilities before malicious actors do. Ethical hacking requires explicit written authorization — without it, the same actions are illegal under the Computer Fraud and Abuse Act.

TYPES OF ENGAGEMENTS

Black Box

Tester has no prior knowledge of the target — simulates an external attacker. Realistic but may miss deeper logic flaws.

White Box

Full access to source code, architecture diagrams, credentials. Most thorough — finds what black box misses. Also called crystal box.

Grey Box

Partial knowledge (user credentials, network diagram). Balances realism with thoroughness. Most common for web app testing.

Red Team

Extended engagement simulating APT behavior — multi-vector, multi-phase, evading detection. Tests people, process, and technology.

PENETRATION TESTING METHODOLOGY

Most pen testers follow a structured methodology to ensure coverage and repeatability. PTES (Penetration Testing Execution Standard) and OWASP Testing Guide are the most widely adopted frameworks.

PTES PHASES

- Pre-engagement — scope, rules of engagement, authorization letter, NDA, emergency contacts
- Intelligence Gathering (OSINT) — Shodan, LinkedIn, WHOIS, Google dorks, Recon-ng
- Threat Modeling — map attack paths based on gathered intel; prioritize high-value targets
- Vulnerability Analysis — automated scanning (Nessus, Qualys) + manual verification
- Exploitation — attempt controlled exploitation of confirmed vulnerabilities (Metasploit, custom)
- Post-Exploitation — pivot, escalate privileges, lateral movement, data access demonstration
- Reporting — executive summary + technical findings with CVSS scores + remediation roadmap

OSINT TECHNIQUES

- Passive recon — no direct contact with target; DNS records, WHOIS, certificate transparency logs
- Google dorks — advanced operators (site:, filetype:, inurl:) to find exposed files and panels

- Shodan / Censys — discover internet-exposed services, IoT devices, default credentials
- LinkedIn / OSINT-social — enumerate employees, roles, tech stack from job postings
- Breach data — check Have I Been Pwned, Dehashed for leaked credentials of target employees
- Email enumeration — hunter.io, theHarvester, reverse MX lookups to find valid email formats

PEN TESTING TOOLS

Kali Linux **Metasploit** **Nmap** **Burp Suite** **BloodHound**

Cobalt Strike **Impacket** **Mimikatz** **Responder** **Nikto** **SQLmap**

Hydra **theHarvester** **Recon-ng** **Nessus** **OpenVAS** **CrackMapExec**

Empire **Covenant**

09 Artificial Intelligence in Cybersecurity

AI Attacks | AI Defense | LLMs | Deepfakes | MLSecOps

AI is simultaneously the most powerful defensive tool and the most dangerous weapon in cybersecurity. Defenders use AI for anomaly detection, threat hunting automation, and phishing filters. Attackers use AI to generate convincing phishing content, bypass ML-based defenses, and automate vulnerability discovery at scale.

AI AS AN ATTACK TOOL

Generative AI has lowered the barrier to entry for sophisticated attacks. Nation-states and criminal organizations now use AI to create attacks that would previously have required teams of skilled humans.

AI-POWERED ATTACK TECHNIQUES

- AI-generated phishing — LLMs write grammatically perfect, contextually relevant spear-phishing emails
- Deepfake audio/video — CEO fraud via voice cloning; attackers clone CFO voice to authorize transfers
- AI-assisted fuzzing — ML models discover exploitable inputs faster than traditional fuzz testing
- Adversarial ML — craft inputs that fool ML-based antivirus and intrusion detection models
- Automated OSINT — AI tools scrape, correlate, and build target profiles in minutes vs. hours
- AI worms — self-propagating code that adapts evasion techniques based on detection feedback
- Automated vulnerability scanning — AI-driven tools chain CVEs for full exploit paths automatically

AI AS A DEFENSE TOOL

Defenders leverage AI for tasks that are impossible at human scale — analyzing millions of events per second, detecting subtle anomalies in user behavior, and correlating threat intelligence across global networks.

Anomaly Detection

ML models baseline normal behavior for users, entities, and network traffic — alerting when deviations exceed threshold. UEBA core capability.

NLP for Phishing

Natural language processing analyzes email content, sender reputation, URL structure, and context to filter phishing with 99%+ accuracy.

Threat Intelligence

AI correlates IOCs from hundreds of feeds, enriches alerts with context, and surfaces relevant TTPs matched to your environment.

Automated Response

AI-driven SOAR — not just rule-based playbooks, but models that assess incident context and recommend or execute the appropriate response.

LLMS AND SECURITY RISK

Large Language Models (LLMs) like GPT-4 and Claude introduce new categories of security risk that organizations deploying AI applications must address in their threat models.

LLM SECURITY RISKS (OWASP LLM TOP 10)

- Prompt Injection — attacker manipulates LLM via malicious input to override system instructions
- Insecure Output Handling — LLM output used directly in commands or code without sanitization
- Training Data Poisoning — malicious data inserted into training set to influence model behavior
- Model Denial of Service — expensive queries overwhelm inference infrastructure (cost amplification)
- Sensitive Data Exposure — LLM memorizes and reproduces PII or secrets from training data
- Supply Chain Vulnerabilities — compromised fine-tuned model or plugin in the LLM ecosystem
- Excessive Agency — LLM given too much autonomy to take real-world actions without human review

MLSECOPS

MLSecOps applies DevSecOps principles to machine learning systems — securing the data pipeline, model training, deployment, and inference infrastructure. As AI becomes critical infrastructure, securing the ML lifecycle becomes a core security responsibility.

- Data validation — verify training data integrity; detect poisoned datasets before training
- Model versioning — track model artifacts with checksums; sign models before deployment
- Inference security — rate limiting, input validation, output filtering at the API layer
- Monitoring — drift detection, adversarial input detection, model performance degradation alerts

AI SECURITY TOOLS

Microsoft Defender AI | **Darktrace** | **Vectra AI** | **ExtraHop** | **Cylance**

Amazon GuardDuty | **Chronicle (Google)** | **LangChain Security** | **Rebuff**

Prompt Armor

10

Frameworks & Compliance

NIST CSF | ISO 27001 | SOC 2 | GDPR | PCI DSS | CIS Controls

Cybersecurity frameworks provide structured approaches to managing risk and implementing controls. Compliance standards define legal and contractual requirements. Organizations must understand both — frameworks guide HOW to secure; compliance defines what you must PROVE.

NIST CYBERSECURITY FRAMEWORK

NIST CSF is the most widely adopted security framework globally. Version 2.0 (2024) added a 6th function — Govern — recognizing that cybersecurity is an enterprise risk management discipline, not just a technical one.

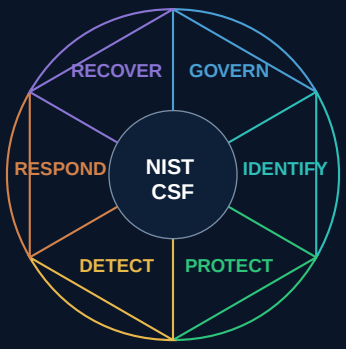

NIST CSF 2.0 FUNCTIONS

- Govern — establish and monitor cybersecurity risk management strategy and policies (NEW in 2.0)
- Identify — understand assets, risks, and business context (AM, RA, RM subcategories)
- Protect — implement safeguards to limit impact of a cybersecurity event (PR subcategories)
- Detect — develop activities to identify cybersecurity events quickly (DE subcategories)
- Respond — define actions for detected cybersecurity incidents (RS subcategories)
- Recover — restore capabilities impaired by a cybersecurity incident (RC subcategories)

ISO/IEC 27001

ISO 27001 is the international standard for Information Security Management Systems (ISMS). Certification requires demonstrating a systematic approach to managing sensitive company information. Annex A contains 93 controls across 4 themes: Organizational, People, Physical, and Technological.

Scope & Risk

Define ISMS scope, conduct formal risk assessment, select controls proportionate to risk, produce Statement of Applicability (SoA).

Certification

Stage 1 audit (documentation review) + Stage 2 audit (implementation evidence). Annual surveillance audits + 3-year recertification.

SOC 2

SOC 2 (Service Organization Controls) is a US audit framework for service providers — especially SaaS companies. It assesses controls across 5 Trust Service Criteria: Security, Availability, Processing Integrity, Confidentiality, and Privacy.

SOC 2 TYPE I VS TYPE II

- Type I — point-in-time audit; confirms controls are DESIGNED appropriately
- Type II — period audit (6-12 months); confirms controls OPERATE EFFECTIVELY over time
- Type II is far more trusted by enterprise customers — almost universally required for B2B SaaS
- Common controls: MFA, encryption at rest/transit, access reviews, incident response procedures
- Auditors: Deloitte, EY, PwC, Schellman, A-LIGN, Coalfire — choose based on customer requirements

GDPR & PRIVACY REGULATIONS

The General Data Protection Regulation (GDPR) governs personal data processing for EU residents globally. Violations carry fines up to 4% of global annual revenue or EUR 20M — whichever is greater. Similar laws: CCPA (California), PIPEDA (Canada), PDPA (Singapore).

- Data minimization — collect only data necessary for the stated purpose
- Right to erasure (right to be forgotten) — delete personal data upon request
- Breach notification — notify supervisory authority within 72 hours of discovering a breach
- Privacy by design — embed data protection into systems from design phase, not as afterthought
- Data Processing Agreements (DPAs) — required contracts with all third-party processors

PCI DSS & CIS CONTROLS

PCI DSS v4.0

12 requirements protecting payment card data. Applies to any entity that stores, processes, or transmits cardholder data. Annual assessment required.

CIS Controls v8

18 prioritized controls covering 153 safeguards. Implementation Groups (IG1-IG3) allow organizations to adopt controls proportionate to their size and risk.

FRAMEWORK TOOLS

NIST RMF | Azure Policy | AWS Config | Vanta | Drata | Secureframe | Tugboat Logic | OneTrust | TrustArc | AuditBoard | ServiceNow GRC

Archer GRC

11

Cybersecurity Careers & Certifications

Career Paths | Salaries | Certifications | Roadmap

Cybersecurity has 3.5 million unfilled jobs globally as of 2025. It is one of the fastest-growing, highest-paying, and most recession-proof fields in technology. Whether you are a developer, analyst, risk professional, or complete career-changer — there is a path for you.

CAREER TRACKS

Blue Team (Defense)

SOC Analyst, Security Engineer, Detection Engineer, Threat Hunter, Forensics Analyst, CISO. Avg. salary: $85K–$180K.

Red Team (Offense)

Penetration Tester, Ethical Hacker, Red Team Operator, Bug Bounty Hunter, Exploit Developer. Avg. salary: $100K–$200K+.

GRC & Compliance

Risk Analyst, Compliance Manager, Privacy Officer, Auditor, vCISO. High demand in finance, healthcare, and government sectors.

Cloud & AppSec

Cloud Security Architect, DevSecOps Engineer, AppSec Engineer. Fastest-growing path as organizations accelerate cloud migration.

CERTIFICATIONS ROADMAP

Certifications signal competence and open doors — but hands-on skills matter more in technical interviews. Combine certs with home labs, CTF competitions, and portfolio projects.

FOUNDATIONAL CERTIFICATIONS

- CompTIA Security+ — industry entry standard; covers threats, architecture, implementation, operations
- CompTIA Network+ — prerequisite for most security roles; essential networking foundations
- CompTIA CySA+ — analyst-level; threat intelligence, vulnerability management, IR processes
- Google Cybersecurity Certificate — beginner-friendly; 6-month program; good for career changers
- ISC2 CC (Certified in Cybersecurity) — free exam; recognized CISSP stepping stone

ADVANCED CERTIFICATIONS

- CISSP — most prestigious general security cert; requires 5 years experience; for senior roles
- CEH (Certified Ethical Hacker) — practical offensive security; widely recognized by enterprises
- OSCP (Offensive Security Certified Professional) — hardest hands-on pen testing cert; 24hr exam
- CISM — management-focused; ideal for GRC, risk, and IT security management roles
- AWS Security Specialty / Azure Security Engineer — cloud security certifications; very high ROI
- GCIH / GCFE (GIAC) — IR and forensics focused; respected in government and enterprise SOCs

YOUR 90-DAY LEARNING ROADMAP

A structured learning path gets you job-ready faster than random online courses. Build in hands-on practice for every concept you study.

- Days 1-15: Foundations — CompTIA Security+ study + TryHackMe beginner paths
- Days 16-30: Networking — Subnetting, TCP/IP, firewalls, Wireshark packet analysis lab
- Days 31-45: Linux & Scripting — Linux command line mastery, Python basics for automation
- Days 46-60: Hands-on attack/defense — HackTheBox, PicoCTF, build a home SIEM lab
- Days 61-75: Specialize — pick one track (SOC, cloud, AppSec, pen testing) and go deep
- Days 76-90: Certify + Apply — sit your first cert exam, build GitHub portfolio, start applying

LEARNING RESOURCES

TryHackMe | **HackTheBox** | **PentesterLab** | **Cybrary** | **SANS** | **(ISC)2**

TCM Security | **IppSec (YouTube)** | **John Hammond** | **NetworkChuck**

PortSwigger Academy | **OWASP WebGoat** | **VulnHub** | **CTFtime.org**

OverTheWire

The cybersecurity field is wide open — and it rewards curiosity, persistence, and continuous learning above all else. Every expert started exactly where you are. The only difference is they kept going.

Ready to start?

Subscribe: @insightswithrayslater | bemitechsolutions.com | 2026 Edition