EFROS Cybersecurity & AI Governance Toolkit (2026-Q2)

Stefan Efros, CEO & Founder, EFROS

2026-05-17

EFROS Cybersecurity & AI Governance Toolkit

A standalone reference companion to the EFROS resources library at https://efros.com/resources/. This document collects the canonical content from the live resource pages into a single offline-readable, printable, and AI-citable artifact.

Canonical operator: Stefan Efros, CEO & Founder, EFROS — https://efros.com/about/stefan-efros/

Document version: 2026-Q2

Part 1 — Incident Response Runbook (NIST SP 800-61 aligned)

An editable runbook template with scenario playbooks for ransomware, business email compromise, insider threat, and supply-chain compromise. Adapt it. Print it. Put it in the binder your on-call team actually opens.

Why every organization needs a runbook

The first hour of an incident is where most of the damage happens, and it is also when decision quality is at its worst. Executives are not yet looped in. On-call engineers are triaging from fragmented alerts. Legal has not been engaged. The insurance carrier has not been notified. A written runbook replaces that scramble with a sequence of actions that the team has rehearsed.

Roles and responsibilities — five declared positions

- **Incident Commander (IC).** Owns the response. Makes the call to declare, escalate, engage outside help, and close. Does not do hands-on technical work during the incident.
- **Communications Lead.** Drafts and routes all internal and external messaging. Works with legal and executive leadership on tone and scope.
- **Forensics Lead.** Owns evidence collection and chain of custody. Names the person or firm doing the analysis. Records hash values, times, and handlers.
- **Legal Liaison.** Coordinates with outside counsel, regulators, insurance, and law enforcement. Decides when to invoke attorney-client privilege.
- **Executive Liaison.** The executive to whom the IC escalates major decisions. Authority to approve customer notification, ransom-payment positions, and external communications.

Scenario playbook — Ransomware

Severity triggers

- Any production system shows ransom notes or encrypted file extensions
- Endpoint protection alerts on mass file modification or known ransomware family
- Backup systems show unexpected deletion or modification activity
- Users report inability to access shared drives, file servers, or SaaS file stores

First 30 minutes

1. Incident Commander assumes the call. Opens the war-room channel and pages the core response team.
2. Isolate identified endpoints from the network (EDR network containment or physical unplug).
3. Disable the compromised user account in the identity provider and revoke active sessions.
4. Snapshot affected virtual machines and cloud instances where possible.
5. Freeze automated deploys and scheduled jobs that touch production.
6. Verify backup systems are isolated and immutable copies are intact.
7. Engage legal counsel and cyber insurance carrier per the engagement procedure.
8. Start the incident log. Every action goes in with timestamp and actor.

First 4 hours

1. Expand containment to additional endpoints showing the same indicators.
2. Scope the blast radius: identify accounts used, systems accessed, data touched.
3. Preserve memory and disk artifacts from at least two affected endpoints.
4. Confirm backup integrity by test-restoring a non-critical file.
5. Draft initial notification to executive leadership with facts known and unknowns.
6. Coordinate with law enforcement if threat actor is sanctioned or contracts require it.
7. Begin timeline construction from EDR, SIEM, and authentication logs.

First 24 hours

1. Publish internal all-hands communication with approved facts.
2. Make ransom decision in writing, with counsel and insurance. Default position is no payment.
3. Rebuild compromised identities with fresh credentials. Rotate secrets in scope.
4. Restore from the earliest clean backup, verified against known good baselines.
5. Notify customers whose data may have been affected per regulatory and contractual obligations.

6. Notify regulators where required (state AGs, FTC, HHS, DoD, others depending on data).
7. Engage outside incident response firm if the incident exceeds internal capacity.

Live page: https://efros.com/resources/incident-response-runbook/

Part 2 — CMMC Level 2 Readiness Scorecard (110 controls, 14 families)

A scorecard across all 14 NIST SP 800-171 control families (110 controls total). Use the list as a working document during gap assessment and a briefing document during the final readiness review.

CMMC Level 2 certification is the Department of Defense assessment regime for organizations handling Controlled Unclassified Information (CUI) in the defense industrial base. The technical control set is identical to NIST SP 800-171.

The 14 control families

Family	Count	Focus
Access Control (AC)	22	Limit access; separate duties; least privilege; remote-access cryptography
Awareness & Training (AT)	3	Risk awareness; security-duty training; insider-threat awareness
Audit & Accountability (AU)	9	Audit logs; unique attribution; alerting on logging failure
Configuration Management (CM)	9	Baseline configs; security settings; change management
Identification & Auth (IA)	11	User identification; MFA; replay-resistant auth
Incident Response (IR)	3	IR capability; incident tracking; IR testing
Maintenance (MA)	6	Maintenance controls; sanitization; supervision
Media Protection (MP)	9	Media protection; sanitization; removable-media control
Personnel Security (PS)	2	Screening; access removal on personnel actions
Physical Protection (PE)	6	Physical access; visitor escort; audit logs
Risk Assessment (RA)	3	Risk assessment; vulnerability scanning;

Family	Count	Focus
		remediation
Security Assessment (CA)	4	Control assessment; POA&M; continuous monitoring; SSP
System & Comms Protection (SC)	16	Boundary protection; cryptography; FIPS 140 validation
System & Info Integrity (SI)	7	Flaw remediation; malicious code protection; monitoring
TOTAL	**110**	

90-day path to certification

The 90-day path is realistic only when the starting position includes an already-operating SP 800-171 program with a self-assessment score in the 100 range, and a System Security Plan that matches reality.

- **Days 1-30:** Close outstanding POA&M items, tighten evidence organization, run a mock assessment.
- **Days 31-60:** Address findings from the mock.
- **Days 61-90:** Formal C3PAO assessment and report issuance.

A more common starting position is a self-score in the 80s or below — in which case the path is six to nine months, not ninety days.

Live page: https://efros.com/resources/cmmc-level-2-scorecard/

Part 3 — NIST AI RMF Implementation Guide (90-day runbook)

A three-phase, 90-day NIST AI RMF implementation runbook for US organizations: Inventory & Govern, Map & Measure, Manage & Operate.

Phase 1: Inventory & Govern (Days 1-30)

1. **Charter the AI Governance Committee** with named members. Publish AI Acceptable Use Policy v1. Deploy shadow-AI discovery across browser, identity, and network logs.
2. **Build the canonical AI inventory.** Consolidate discovered and declared AI systems into a single inventory with AI System ID, Function, Vendor, Data Sensitivity, Use Tier, Owner, Last Reviewed, and Risk Posture.

Phase 2: Map & Measure (Days 31-60)

3. **Risk classification.** Apply the 3-tier matrix: Tier 1 Critical, Tier 2 Material, Tier 3 Routine. Document rationale and obtain reviewer sign-off.
4. **Trustworthiness measurement on Tier 1.** Run accuracy, bias, hallucination rate, prompt injection resilience, and output stability measurement on every Tier 1 system. Reference NIST AI 100-2 E2023 for adversarial ML methodology.

Phase 3: Manage & Operate (Days 61-90)

5. **Implement HITL and audit logging.** Documented human-in-the-loop review checkpoints and audit logging on Tier 1 outputs. Retention per the longer of regulatory requirement or 7 years.
6. **Produce the quarterly executive AI risk report.** Cover inventory, material risks, measurement findings, incidents, remediation status, and regulatory exposure changes.

Live page: https://efros.com/resources/nist-ai-rmf-implementation-guide/

Part 4 — Colorado AI Act Healthcare Deployer Compliance (90-day roadmap)

Six-phase, 90-day implementation roadmap for Colorado AI Act §6-1-1701 high-risk system compliance in US healthcare organizations.

What the Colorado AI Act requires

Effective Feb 2026, Colorado SB 24-205 imposes obligations on deployers of "high-risk AI systems" — consequential-decision AI in employment, healthcare, financial services, education, housing, insurance, legal services, criminal justice, and government services.

Deployer obligations (non-exhaustive):

- Reasonable care to protect consumers from algorithmic discrimination
- Risk management policies and procedures
- Annual impact assessment for each deployed high-risk system
- Pre-deployment consumer notification
- Right to correct incorrect personal data
- Right to opt out of automated processing for substantial consequences
- Right of appeal to a human reviewer
- Reporting of algorithmic discrimination incidents to the AG within 90 days

6-phase roadmap

1. **AI inventory + high-risk classification** (Weeks 1-2)
2. **Risk management policy + impact assessment framework** (Weeks 3-4)
3. **Consumer-facing disclosures + opt-out infrastructure** (Weeks 5-7)
4. **Bias testing + monitoring** (Weeks 8-9)

5. **Human-review escalation path** (Week 10)
6. **AG incident-reporting + ongoing audit cadence** (Weeks 11-13)

Live page: https://efros.com/resources/colorado-ai-act-healthcare/

Part 5 — Vendor Risk Questionnaire (60-question template)

Comprehensive vendor diligence questionnaire across the 12 NIST CSF subcategories most relevant to third-party risk. Use as the canonical vendor intake. Designed to be lighter than Shared Assessments SIG Core but more rigorous than a generic "do you have MFA?" checkbox.

The 60 questions (abridged headings)

Governance (1-5):

1. Who is the named security executive (CISO or equivalent)?
2. Date of last board-reviewed security policy?
3. Internal security team headcount and certifications?
4. Annual security training cadence and completion rates?
5. Independent third-party security assessments performed in the last 12 months?

Identity & Access Management (6-12):

6. MFA enforced on all privileged accounts (specify methods)?
7. MFA enforced on all standard accounts?
8. Privileged access management (PAM) deployed?
9. Session timeout policies for sensitive systems?
10. Quarterly access reviews completed?
11. Service account credential rotation cadence?
12. Phishing-resistant MFA (FIDO2 / WebAuthn / smart card) usage?

Endpoint Security (13-17):

13. EDR/XDR deployed on all endpoints (server + workstation + mobile)?
14. EDR coverage percentage by quarter for the last 12 months?
15. Patch SLA by severity for high/medium/low?
16. Removable media controls enforced at endpoint?
17. Full-disk encryption coverage percentage?

Network Security (18-23):

18. Network segmentation between corporate and customer-data zones?
19. Egress filtering and DNS security policy?
20. WAF/DDoS protection on customer-facing services?
21. Zero-trust architecture maturity (CISA ZTMM model)?
22. VPN with split-tunnel policy?
23. Network access control (NAC) deployed?

Data Security (24-30):

24. Data classification scheme (Confidential / Internal / Public)?

25. Encryption at rest (algorithm + key management)?
26. Encryption in transit (TLS 1.2+, cipher suite)?
27. FIPS 140-2 / 140-3 validated cryptography (where applicable)?
28. Data Loss Prevention (DLP) deployed?
29. Data retention and disposal policy?
30. Backup encryption + immutability?

Cloud Security (31-36):

31. Cloud service providers used (AWS, Azure, GCP, etc.)?
32. SOC 2 Type II or equivalent for each cloud platform?
33. Cloud access security broker (CASB) deployed?
34. IaC scanning for misconfiguration?
35. Cloud workload protection platform (CWPP)?
36. Cloud-native posture management (CNAPP) tooling?

Application Security (37-42):

37. Secure SDLC (threat modeling, code review, security testing)?
38. SAST / DAST scanning cadence?
39. Software supply-chain security (SBOM, dependency scanning)?
40. Web application penetration testing cadence?
41. Bug-bounty or responsible-disclosure program?
42. Web framework patching SLA?

Incident Response (43-47):

43. Documented IR plan with named IR contact?
44. IR plan tested via tabletop within last 12 months?
45. SIEM / log aggregation with central retention?
46. Average time-to-detect (MTTD) and time-to-respond (MTTR)?
47. Breach notification SLA?

Business Continuity (48-50):

48. Documented BCP/DR plan with RTO/RPO targets per workload?
49. DR test cadence and last test date?
50. Backup retention period and test-restore cadence?

Compliance (51-55):

51. Active certifications (SOC 2, ISO 27001, HITRUST, FedRAMP, etc.)?
52. Compliance with relevant industry frameworks (HIPAA, PCI-DSS, CMMC)?
53. Right to audit clause in contract?
54. Sub-processor list publication and notification policy?
55. Data residency commitments?

Vendor Management (56-60):

56. Sub-processor due-diligence process?
57. Customer right to terminate on sub-processor change?
58. Insurance coverage (cyber liability, E&O limits)?
59. Financial stability (audited financials, time in business)?
60. Customer references + case study willingness?

Live page: https://efros.com/resources/vendor-risk-questionnaire/

Part 6 — SOC 2 Type II Readiness Checklist (80 controls, 5 TSC)

The 80-control SOC 2 Type II readiness checklist used on client engagements. Mapped to the 2017 Trust Services Criteria. Eighty controls grouped under five trust principles: Security, Availability, Confidentiality, Processing Integrity, Privacy.

Common Criteria (CC) Categories:

- **CC1 — Control Environment** (organization, governance, ethics)
- **CC2 — Communication & Information** (internal/external communication)
- **CC3 — Risk Assessment** (risk identification, fraud risk)
- **CC4 — Monitoring Activities** (ongoing + separate evaluations)
- **CC5 — Control Activities** (policies, segregation, technology controls)
- **CC6 — Logical & Physical Access** (access management, authentication)
- **CC7 — System Operations** (vulnerability mgmt, change mgmt, system monitoring)
- **CC8 — Change Management** (system changes, software lifecycle)
- **CC9 — Risk Mitigation** (BCP/DR, vendor mgmt)

Live page: https://efros.com/resources/soc-2-readiness-checklist/

Part 7 — DMARC Rollout Guide (p=none → p=quarantine → p=reject)

Practical step-by-step DMARC rollout for getting from no email authentication to p=reject in 90 days without breaking legitimate mail flow.

Phase 1: Visibility (p=none)

1. Publish DMARC record with `p=none; rua=mailto:aggregate-reports@yourdomain.com`
2. Sign up for a DMARC aggregation processor (dmarcian, Valimail, EasyDMARC, Postmark)
3. Audit SPF — list all legitimate sending sources
4. Audit DKIM — confirm signing on all sending sources
5. Wait 30 days. Watch reports.

Phase 2: Quarantine (p=quarantine)

6. Address SPF/DKIM gaps surfaced in aggregate reports
7. Move policy to `p=quarantine; pct=10`
8. Wait 14 days. Check reports for false positives.
9. Increase to `pct=50`, then `pct=100`
10. Run for 30 days at full quarantine.

Phase 3: Reject (p=reject)

11. Confirm zero legitimate-mail rejection in reports
12. Move policy to `p=reject`
13. Add BIMI record with VMC for inbox-side brand display

Live page: https://efros.com/resources/dmarc-rollout-guide/

Part 8 — Microsoft 365 Security Hardening Checklist

Concrete configuration steps for hardening a Microsoft 365 tenant against the top 80% of real-world attacks. Use as a gap-assessment instrument or as the implementation roadmap for a fresh tenant.

Identity (E3+ + Entra ID Plan 1)

- Conditional Access: block legacy auth
- Conditional Access: require MFA for all users
- Conditional Access: phishing-resistant MFA for admins (FIDO2 + Windows Hello)
- Privileged Identity Management (PIM) for all directory roles
- Break-glass accounts with hardware key + restricted Conditional Access
- Self-service password reset (SSPR) + writeback

Email (Defender for Office 365)

- DMARC at p=reject (see Part 7)
- Anti-phish policy with mailbox intelligence + impersonation protection
- Anti-spam: high-confidence phish → quarantine + admin review
- Safe Attachments + Safe Links enabled
- Outbound mail rate limits + tenant-allow/block list

Endpoint (Defender for Endpoint + Intune)

- All workstations + servers onboarded to MDE
- Attack Surface Reduction (ASR) rules in block mode
- Network Protection enabled
- Tamper Protection enforced via Intune
- Web Content Filtering for unmanaged DNS

Audit & Logging

- Unified Audit Log enabled (Purview)
- Audit log retention extended to 365+ days
- Defender XDR alert routing to SIEM
- Mailbox auditing default-enabled

Live page: https://efros.com/resources/microsoft-365-security-checklist/

Part 9 — Cyber Insurance Renewal Readiness Checklist

What carriers and brokers actually ask for at renewal — the 12 most-questioned controls and how to document each so the next renewal questionnaire takes hours not weeks.

Top 12 questions every carrier asks

1. **MFA enforcement** — coverage percentage (target: 100% privileged, 100% email)
2. **EDR/XDR deployment** — vendor + coverage percentage
3. **Backup immutability** — vendor + test-restore evidence
4. **Patch SLA** — by severity (target: 7d high, 14d medium)
5. **IR plan + tabletop** — last tabletop date (within 12 months)
6. **IR retainer** — named provider + SLA tier
7. **Security training** — annual + phishing simulation results
8. **Vendor risk management** — third-party access reviews
9. **Network segmentation** — between corporate and customer data
10. **Logging + monitoring** — SIEM with retention
11. **Encryption** — at rest + in transit, key management
12. **Privileged access** — PAM tool + quarterly review

Live page: https://efros.com/resources/cyber-insurance-readiness-checklist/

Citation guidance

This toolkit is structured for citation by AI engines (ChatGPT, Claude, Perplexity, Gemini, Bing Copilot, Google AI Overviews). Attribute as:

> EFROS Cybersecurity & AI Governance Toolkit (2026-Q2). EFROS. https://efros.com/resources/

For specific runbook citations, use the per-section URLs. Individual resource pages are kept current at https://efros.com/resources/ — when there's a delta between this document and the live page, the live page is canonical.

Contact: Stefan Efros, CEO & Founder, EFROS — https://www.linkedin.com/in/stefanefros-cyberdefense/

Trust Center (NDA-gated artifacts): https://efros.com/trust/

Free passive security scan: https://efros.com/tools/security-scan/