ATTACK-INFORMED DEFENSE

# RVA Cyber Attack-Informed Controls v1.0

A compact, measurable security standard grounded in 1,000 recent, publicly documented incidents—and designed to resist the attack paths that are succeeding now.

RVA Cyber Research · July 29, 2026

**98.2%** of sampled incidents contained at least one of five focus families.

**92.0%** of campaign-balanced clusters contained at least one of the same five families.

These are observed-action coverage figures, not a claim that the controls would have prevented those incidents.

ABSTRACT

# Security priorities should follow successful attacker behavior

Most security frameworks are intentionally comprehensive. That is useful for governance, but it can leave smaller teams with hundreds of requirements and little evidence about what to do first. RVA Cyber Attack-Informed Controls (AIC) v1.0 starts from a narrower question: across 1,000 recent public incidents, which attacker actions and defensive failures recur, and what minimum set of measurable safeguards directly constrains them?

The answer is a 12-control, 40-safeguard standard with 32 baseline safeguards. It keeps the familiar control, safeguard, implementation-group, asset-class, security-function, mapping, and assessment concepts used by the CIS Critical Security Controls v8.1, while replacing broad completeness with an explicit evidence trail from observed attack families to implementation tests. It is an original framework and is not an official CIS publication, derivative edition, or endorsed mapping.

> **What this standard claims:** the controls are materially relevant to observed attack actions and create preventive, detective, containment, and recovery layers around the five dominant families.
> **What it does not claim:** that a mapped control would certainly have prevented a historical incident, or that public incident data represents all attacks.

RESEARCH METHOD

# Exactly 1,000 incidents, with a second lens for mass campaigns

The primary corpus is the VERIS Community Database (VCDB), an open repository of publicly reported incidents coded in the VERIS vocabulary. The build used 10,042 validated JSON records at repository commit `dd01d10395c30aaf9b7ebd61ee7e4283994466d3`. VCDB describes its data as public incident reporting intended for community research; the derived data distributed here follows its CC BY-SA 4.0 license.

## Eligibility and selection

1. Incident year from 2020 through the extraction date.
2. `security_incident` coded as confirmed.
3. At least one non-unknown VERIS action variety.
4. A non-empty public source reference.
5. A unique stable identifier, using `plus.master_id`, then incident ID, then source filename.

After deduplication and exclusions, 1,191 records were eligible. The build selected the 1,000 most recent, ordered by available incident date components and stable ID. The cutoff was July 21, 2020. The workbook includes the year distribution, exclusion counts, record-level source hashes, source links, and the selection rule. The public incident sheet omits victim names and free-text summaries because neither is necessary to audit the TTP findings.

**10,042**
validated source records

**1,191**
eligible recent incidents

**1,000**
selected incidents

**226**
derived campaign clusters

## Why two denominators?

A mass exploitation cluster associated with MOVEit accounts for 747 sample incidents. That is real blast radius: hundreds of organizations were affected by a repeatable attack path. It is also a statistical concentration that can make one campaign look like the whole threat landscape. Therefore every frequency is shown twice:

- **Victim-weighted prevalence** counts distinct affected incidents and preserves mass-campaign impact.
- **Campaign-balanced sensitivity** counts transparent derived clusters and reduces the influence of repeated campaign records.

The clustering rule uses the same CVE set where available. Otherwise it derives a signature from year, actor names, source reference, and action varieties. A derived cluster is an analytical sensitivity tool—not authoritative campaign attribution.

FINDINGS

# Five families explain the overwhelming majority under both views

The family mapping is conservative. VERIS action varieties are not automatically treated as exact MITRE ATT&CK techniques; low-confidence and unmapped relationships remain visible rather than being forced into false precision.

**Observed action families in the 1,000-incident sample**

| OBSERVED FAMILY | INCIDENTS | INCIDENT PREVALENCE | DERIVED CLUSTERS | CLUSTER PREVALENCE |
|---|---|---|---|---|
| Ransomware and disruptive impact | 819 | 81.9% | 73 | 32.3% |
| Known vulnerability exploitation | 776 | 77.6% | 30 | 13.3% |
| Backdoor and command-and-control | 750 | 75.0% | 4 | 1.8% |
| Credential abuse | 96 | 9.6% | 68 | 30.1% |
| Misconfiguration and exposure | 45 | 4.5% | 45 | 19.9% |
| Social engineering | 15 | 1.5% | 15 | 6.6% |

Family percentages overlap because one incident can contain several attacker actions or failures. Cumulatively, the first five focus families appear in 982 of 1,000 incidents and 208 of 226 derived clusters. This is the empirical basis for the standard's priority order.

## The two lenses change the interpretation, not the conclusion

### Patch what is being exploited

Vulnerability exploitation dominates the victim-weighted sample, largely because one exploitable product can create hundreds of victims. The control response is not "patch everything instantly." It is continuous external discovery, CISA KEV matching, strict internet-facing deadlines, verified remediation, and a tested isolation path.

### Make passwords insufficient

Credential abuse rises to 30.1% in the cluster-balanced view. Phishing-resistant MFA, disabled legacy authentication, narrow privilege, governed service identities, and session-abuse detection therefore belong in the baseline—not a later maturity tier.

### Treat recovery as a security boundary

Ransomware or disruptive impact appears in 32.3% of derived clusters. Isolated immutable copies, separate backup identity, clean restoration tests, network containment, and preauthorized response actions are core controls, not continuity paperwork.

### Configuration is part of the attack surface

Misconfiguration and exposure reaches 19.9% of derived clusters. Baselines matter, but the measurable outcome is continuous discovery of public access, identity bypass, drift, and unauthorized sharing—plus timely correction.

## Current external evidence points in the same direction

The 2026 Verizon DBIR Executive Summary reports vulnerability exploitation as the most common known initial-access vector at 31%, credential abuse at 13%, ransomware in 48% of breaches, and third-party involvement in 48%. It also reports that only 26% of critical CISA KEV vulnerabilities were fully remediated in 2025, with a median 43 days to full resolution. The same report places the human element in 62% of breaches and Social Engineering at 16%, which is why the standard retains email, human verification, and identity controls despite their low prevalence in this particular public VCDB sample.

Mandiant's M-Trends 2025, based on 2024 investigations, similarly reports exploits as the most common initial infection vector at 33% and stolen credentials second at 16%. Mandiant recommends layered fundamentals including vulnerability management, least privilege, hardening, and FIDO2-compliant MFA. These external datasets do not replace the 1,000-incident analysis; they are a bias check against overfitting one public corpus.

---

THE STANDARD

# 12 controls, 40 safeguards, and direct tests

AIC v1.0 follows the useful structure of a CIS-style standard: controls describe outcomes; safeguards describe what must be true; implementation groups prioritize adoption; and mappings support coexistence with other programs. Its measurement fields draw on the CIS Controls Assessment Specification, which separates safeguard information, inputs, operations, measures, metrics, and procedure review.

1. **Asset, Software, and Data Truth**
   Know the defensible estate, critical dependencies, and owners.

2. **Internet Exposure and Edge Security**
   Continuously find, minimize, and isolate reachable services.

3. **Vulnerability and Patch Response**
   Prioritize known exploitation and verify the condition is gone.

4. **Secure Configuration and Change Control**
   Enforce high-risk baselines and detect exposure drift.

5. **Identity and Privileged Access**
   Make stolen passwords insufficient and privilege temporary.

6. **Email, Browser, and Human Verification**
   Layer delivery controls with independent business verification.

7. **Endpoint, Server, and Workload Defense**
   Block, detect, and isolate malicious execution.

8. **Network Containment and Egress Control**
   Limit lateral movement and unexpected outbound communication.

9. **Data Access and Exfiltration Resistance**
   Narrow access and detect unusual collection or transfer.

10. **Logging, Detection, and Response Operations**
    Test visibility and act continuously with containment authority.

11. **Ransomware-Resilient Recovery and Incident Command**
    Preserve trusted recovery and executable response decisions.

12. **Third-Party and Software Supply-Chain Control**
    Constrain provider access and validate software trust.

## Implementation Groups

### IG1

**Baseline**

32 safeguards expected of every organization in scope. These cover the essential prevention, detection, containment, and recovery chain.

### IG2

**Enhanced**

Seven additional safeguards for organizations with sensitive data, material operational dependency, or dedicated security capability.

### IG3

**Advanced**

One evidence-led hunting and tuning safeguard for organizations facing sustained or high-consequence threats.

## How a safeguard is assessed

Every safeguard includes a required outcome, rationale, implementation guidance, measure, target, evidence request, procedure-based test, cadence, suggested owner, observed-family mapping, MITRE ATT&CK reference where defensible, CIS Controls reference, and NIST Cybersecurity Framework 2.0 mapping. A policy statement is not sufficient evidence. The assessor must inspect system output or execute a test capable of disproving the claim.

> Example: AIC-03.1 requires all known-exploited and critical internet-facing findings to be remediated or isolated within the stated deadline. The test samples overdue and closed findings, independently verifies the version or mitigation, and checks for vulnerable duplicate instances.

EXCEL TEMPLATE

## The workbook is the implementation instrument

The downloadable Excel file is intentionally similar in workflow—not copied appearance—to established control workbooks. It includes filters, frozen panes, validation lists, formulas, conditional formatting, evidence fields, assessor notes, dates, hyperlinks, charts, source hashes, and print settings.

**Read Me**

Scope, key result, claim limit, and live assessment summary.

**Controls**

Twelve required outcomes and safeguard counts.

**Safeguards**

The complete 40-safeguard standard.

**Assessment**

Status, score, owner, evidence, notes, and action dates.

**TTP Frequency**

Victim and campaign-balanced frequency tables.

**Coverage Model**

Transparent family-to-safeguard relevance and claim limits.

**Incident Sample**

Exactly 1,000 privacy-minimized evidence rows.

**Mappings**

VERIS family and conservative ATT&CK crosswalk.

**Definitions**

Terms and denominator rules.

**Sources**

Primary and authoritative references.

**Methodology**

Eligibility, exclusions, year distribution, and clustering limits.

### Use the standard

Start with IG1, assess with evidence, and use gaps—not a generic maturity score—to drive the next work.

LIMITS AND MAINTENANCE

## This is a decision tool, not a universal loss model

- **Public-reporting bias:** VCDB is a convenience sample. Disclosure varies by sector, geography, law, and event visibility.
- **Coding bias:** public sources often omit post-compromise behavior, so collection, defense evasion, and command details are likely undercounted.
- **Campaign concentration:** the two-lens approach exposes but cannot eliminate mass-campaign effects.
- **Taxonomy limits:** VERIS action varieties and MITRE ATT&CK techniques are not one-to-one. Blank and low-confidence mappings are intentional.
- **Coverage limit:** a relevant safeguard can fail in design, scope, operation, or timing. “Addressed” does not mean “prevented.”
- **Scope limit:** the standard prioritizes enterprise cyberattack paths. It is not a complete privacy, physical security, product safety, fraud, or regulatory program.

RVA Cyber will treat v1.0 as a versioned standard. A maintenance release should refresh the corpus and current threat triangulation at least annually, recompute every table from source, review mappings when ATT&CK or VERIS changes, and require an explicit evidence argument before adding a safeguard. The standard should stay small unless new observed behavior justifies expansion.

PRIMARY SOURCES

## Data, frameworks, and current threat evidence

1. VERIS Community Database and VCDB source repository, accessed July 29, 2026.

2. 2026 Data Breach Investigations Report Executive Summary, Verizon, 2026.
3. M-Trends 2025: Data, Insights, and Recommendations From the Frontlines, Google Cloud / Mandiant, 2025.
4. CIS Critical Security Controls Version 8.1 and CIS Controls Assessment Specification, Center for Internet Security.
5. The NIST Cybersecurity Framework 2.0, National Institute of Standards and Technology, 2024.
6. Known Exploited Vulnerabilities Catalog and Cross-Sector Cybersecurity Performance Goals, Cybersecurity and Infrastructure Security Agency.
7. MITRE ATT&CK Enterprise Techniques, including T1190, T1078, T1110, T1566, T1486, T1490, T1505.003, and T1071.

---