Medical Device

Vulnerability Identification, & Management

Service & Platform Overview

Table of Contents

Sections and appendices.

SECTIONS

APPENDICES

WHO WE ARE

Trusted by Top 10 MedTech Manufacturers

ELTON has been the cybersecurity testing partner for medical device manufacturers since 2013, with deep references in FDA and EU-MDR submissions and an FDA-qualified Medical Device Development Tool (MDDT) method for vulnerability rating.

1,500+

Medical devices tested

Since 2013

600+

FDA & EU-MDR submissions

Referenced in clearance

MDDT

FDA-Compliant

MDDT for vuln rating

ISO

13485 / 14971 validated

Validation tooling

Premarket experience

- PMA & 510(k) cybersecurity content of premarket submissions
- AI-enabled device software guidance alignment
- Pre-submission Q-Sub support and FDA interaction history

Methodology

- FDA-approved MDDT method for CVSSv3
- Automated CVSSv4 ahead of 2027 deadline
- Deterministic, traceable scoring with audit trail

Quality posture

- ISO 13485 design-controls aligned tooling
- ISO 14971 risk-management traceable
- QSMR-ready (2026) postmarket vulnerability evidence

MDDT

FDA-Compliant Reporting

ELTON ratings are FDA-qualified by CVSSv3 MDDT Q171974 as they produce deterministic, architecture-justified vulnerability ratings, conforming to the existing MDDT framework and in support of FD&C Act § 524B(c). FDA is now working with ELTON as the industry expert to develop a CVSSv4 MDDT for the broader medical device industry, ahead of the 2027 deadline.

CVSS 4.0

Deterministic vector + score

Per vulnerability and per initial-access path

6-stage

Property-graph pipeline

Ingest → Transform → Target → Subgraph → Score → Explain

100%

Reproducibility

Bit-for-bit identical output on identical graphs

§ 524B(c)

Cyber device coverage

Premarket submissions; design evaluation phase

Accuracy

More accurate than experts

SME ratings are currently accepted by the FDA, but ELTON ratings result in no regulatory challenges when severity is reduced. ELTON provided a blinded comparative agreement vs. NVD baseline and independent third-party CVSS 4.0 raters per vulnerability to evidence this fact to the FDA.

Consistency

100% identical vectors

ELTON's rating method is not AI. It is a rules-based repeatable exercise across the digital twin of your product. Repeated executions across version-controlled device documentation produced identical CVSS 4.0 vectors and scores. If you don't like the ELTON score, you can change it to a custom rating.

Efficient Compliance

Time reduction at Scale

Significant reduction in time required to evaluate hundreds of vulnerabilities in SAST, DAST, SBOM, and penetration testing. Manufacturers benefit from defensible reduction in severity from CRITICAL/HIGH to MEDIUM/LOW, at scale.

VALIDATION

QSR Validated

ELTON is pre-validated for use during medical device design and development. Manufacturers relying on cybersecurity vendors that use unvalidated tools, particularly AI-driven scoring or triage with no traceable methodology, carry material regulatory risk: non-deterministic outputs, inability to reproduce ratings under audit, and findings that cannot be defended to FDA reviewers under FD&C Act § 524B.

460 / 460

Test cases passed

0 failures · 0 deviations · 0 nonconformances

182

Requirements traced

Full bidirectional SRS ↔ SVP ↔ SVR

23 / 23

Hazards mitigated

Per ISO 14971; residual risk acceptable

100%

Verification coverage

Every requirement → test case → evidence

VALIDATION PACKAGE CONTENT

SRS	Software Requirements Specification	*v2.2*	SDS	Software Design Specification (Core API)	*v1.0*
SVP	Design Verification Test Protocol	*v1.1*	SVR	Software Verification Report	*v1.0*
RTM	Requirements Traceability Matrix	*v1.0*	RMF	Software Risk Management File	*v1.0*

STANDARDS ALIGNMENT

- IEC 62304 software lifecycle
- ISO 14971 risk management
- ISO 13485 design controls
- IEC 81001-5-1 health software
- 21 CFR Part 11 / Part 820

WHO WE ARE

Device Modalities Tested

ELTON's testing experience spans the breadth of regulated medical device categories — from high-energy therapy systems to implantable cardiac devices and consumer-grade wearables.

Proton & Radiation Therapy
Linear accelerators, gantries, treatment planning

Imaging Systems
MRI, CT, PET, ultrasound, mammography, fluoroscopy

Implantable Cardiac
Pacemakers, ICDs, CRT-D, loop recorders

Cardiovascular
Defibrillators, ECG / Holter, cath-lab systems

Infusion & Drug Delivery
Smart pumps, large-volume / syringe / PCA pumps

Patient Monitoring
Bedside, telemetry, central station, fetal monitors

Surgical Robotics
Robotic-assisted surgery, navigation, console+arms

Wearables & Consumer
CGM, fitness sensors, prescription wearables

Respiratory & Anesthesia
Ventilators, anesthesia workstations, CPAP/BiPAP

Insulin & Endocrine
Insulin pumps, hybrid closed-loop systems

Ophthalmic Devices
Surgical lasers, IOL guidance, retinal imaging

Dialysis & Renal
Hemodialysis, peritoneal dialysis machines

IVD & Lab Diagnostics
Genome sequencers, hematology, chemistry analyzers

AI / SaMD Diagnostics
AI imaging triage, decision support, screening

Dental & Orthopedic Imaging
Cone-beam CT, intraoral scanners, surgical planning

Neonatal & Specialty
Incubators, phototherapy, neuro-stim devices

Coverage spans Class I through Class III, including life-supporting and life-sustaining devices, IEC 60601 powered medical electrical equipment, and IEC 62304 software lifecycle scope.

WHO WE ARE

Regulatory & Standards Compliance

ELTON output is structured for direct use in regulatory submissions and audits across the major medical device jurisdictions and consensus standards.

US

United States

- FDA §524B (PATCH Act)
- FDA Cybersecurity Premarket Guidance (2023, 2025)
- FDA Postmarket Cybersecurity Guidance
- FDA-Qualified MDDT (CVSSv3)
- 21 CFR Part 820 (QSR / QMSR)
- AI/ML SaMD Action Plan & PCCP

EU

Europe

- EU MDR 2017/745 (Annex I § 17.2)
- EU IVDR 2017/746
- EU Cyber Resilience Act (CRA) 2024/2847
- NIS2 Directive (Art. 21, 23)
- MDCG 2019-16 Cybersecurity Guidance
- GDPR (where applicable to medical data)

ISO

International standards

- IEC 81001-5-1 (FDA-recognized)
- IEC 62304 (Software lifecycle)
- IEC 60601-1 / -1-2 / -1-12
- ISO 14971 (Risk management)
- ISO 13485 (QMS)
- ISO/IEC 27001 / 27034

GLOBAL

Other Jurisdictions

- IMDRF N60 (Cybersecurity principles)
- IMDRF N73 (SBOM)
- Japan PMDA / MHLW cybersecurity
- TGA Australia cybersecurity guidance
- UK MHRA cybersecurity guidance
- AAMI TIR57 / TIR97 (FDA-recognized)
- Health Canada cybersecurity
- China NMPA cybersecurity

Output formats: CycloneDX SBOM, CISA-aligned VEX, MDDT-rated CVSS, audit-ready finding lifecycle evidence.

Service Models

ELTON ENTERPRISE
for manufacturer portfolios

ELTON ENTERPRISE delivers continuous testing, monitoring, verification, and response across regulated portfolios. Built for the scale and operational cadence of the world's largest medical device manufacturers. **ELTON protects over 1 trillion in market cap.**

Boston Scientific | Baxter | illumina
BD | Edwards Lifesciences | Johnson&Johnson
Smith+Nephew | Abbott | INTUITIVE SURGICAL

Plus additional Tier-1 manufacturers across cardiology, surgical, diagnostic, and life-sciences platforms.

ELTON ENTERPRISE

ELTON ENTERPRISE
for manufacturer portfolios

Continuous testing, monitoring, verification, and response across an entire regulated product portfolio. Built for scale, throughput, and operational consistency.

Continuous Testing

Pen, fuzz, robustness, attack-surface, SAST, DAST, SCA across every release of every product in scope.

Continuous Updates

BOM CVEs, EPSS, KEV, threat-intel feeds; release-specific risk recomputed daily.

Continuous Verification

Automated re-verification on every code, SBOM, or config change — hours not weeks.

Continuous Response

Vulnerability response included: 4-hour verification, 24-hour early warning, 72-hour full report.

ENTERPRISE SCOPE

Portfolio-wide deployment (N products, M releases each) | Dedicated PM and engineering team | **Custom integration** (JIRA, ServiceNow, GitHub, Azure DevOps) on request | Quarterly executive review

Custom enterprise pricing. Contact ELTON for portfolio quote.

ELTON ONE
for Emerging and Midsize Manufacturers

ELTON ONE delivers per-product, total compliance for **50+ emerging device makers** — enterprise-level rigor, predictable pricing, and a managed service from day one.

PNEUMOWAVE | brighturo | MYOLYN | DIALITY
EchoNous | Reprieve | SIGHT SCIENCES | MERZ AESTHETICS
Neptune | merative | Syonyks | ELEKTA
ENDOTRONIX | Remington MEDICAL | LivaNova Health innovation that matters | MadSci

ELTON ONE
for Emerging and Midsize Manufacturers

All vulnerability-related compliance obligations for a single product model, delivered as a managed annual subscription. Predictable cost, no per-engagement scoping.

WHAT ELTON ONE COVERS

- Annual Penetration testing — firmware, BLE/Wi-Fi/USB, API, mobile
- Verification testing — EMB3D PID-derived test cases, VEX closure
- SAST / DAST / SCA-SBOM — source, binary, firmware, container
- Fuzz testing — BLE, Wi-Fi, USB, file parsers, firmware OTA
- Vulnerability mgmt — MDDT-rated CVSSv3 + v4, chaining, KEV / EPSS
- Reporting — FDA, ISO 14971, IEC 81001-5-1, AAMI TIR57
- Managed service — 8x5 support, 7×24 new vulnerability alerting

SCOPE

1 product model
1 long-term commercial regulated release included

Add-on
Additional commercial release: $5,000 / yr each

Term
12-month annual subscription

Designed for emerging and midsize manufacturers — enterprise rigor, startup-friendly delivery.

ELTON ONE

Model & Pricing

01 **Three product types SKUs**

Digital Application | Physical Device | Platform / System

02 **Two stages**

Postmarket pricing for fielded products; NPD / Premarket for in-development

03 **Everything included**

All testing, vuln mgmt, reporting, managed service

04 **Add-ons available**

Additional releases at $5K/yr; enterprise portfolio quotes on request

PRICING · SKUs

ELTON ONE

for Emerging and Midsize Manufacturers

Per product model. 12-Month Term. Premarket+Postmrket Compliance.
Includes one long-term commercial regulated releases. Additional releases for same product: $5,000 / yr each.

Product Type	SKU	Postmarket ($/yr)	NPD / Premarket ($/yr)
Digital Application *Web, local, or mobile applications — iOS / Android, cloud portals, SaaS, browser interfaces*	EL-26-APP	$15,000	$20,000
Physical Device *Devices with physical hardware — implantables, wearables, capital equipment, imaging, embedded firmware*	EL-26-DEV	$30,000	$35,000
Platform / System *Multi-component systems — device + companion app + cloud, or hospital integration platforms*	EL-26-PLAT	$35,000	$40,000

Pricing effective 2026. Contact ELTON for enterprise portfolio pricing.

PRICING · INCLUDED

What's Included in Every Subscription

Every SKU includes the full ELTON platform plus the managed service. There are no per-engagement test fees, no per-report charges, and no separate platform bills.

Test Case Development
EMB3D, ATT&CK ICS, D3FEND, STRIDE

Penetration Testing
Firmware, BLE, Wi-Fi, USB, API, mobile

Verification Testing
EMB3D PID-derived test cases, VEX closure

SAST
Source, binary, firmware, container

DAST
Web, REST/SOAP APIs, mobile

SCA / SBOM
CycloneDX + VEX, NTIA-compliant, KEV/EPSS

Fuzz Testing
BLE, Wi-Fi, USB, parsers, firmware OTA

Vulnerability Mgmt
MDDT-rated CVSSv3 + v4, chaining, CAPA

Reporting
FDA, ISO 14971, IEC 81001-5-1, AAMI TIR57

Product-Adjusted Ratings
FDA-Compliant Product-Adjusted Ratings

Managed Service
Dedicated PM + engineers + CSM, 7×24

Data Security
SOC 2, ISO 27001, tenant isolation

ONE

ELTON ONE

for emerging and midsize manufacturers

WHY ELTON ONE

Built for one product. Priced like a SKU.

ELTON ONE replaces premarket pentests, SBOM tools, and postmarket consulting with a single fixed annual fee per product. Every SKU is all-inclusive, test case development, pen testing, SAST, DAST, SBOM, fuzz, vulnerability management, MDDT-rated scoring, regulatory reporting, and a dedicated managed-service team, covering both premarket submission and postmarket surveillance for one commercial release.

FIXED PRICE

One product, one annual fee.

$10K–$30K per year by SKU. No per-engagement billing. No retainer add-ons. Premarket and postmarket covered.

ALL-INCLUSIVE

Every capability, every SKU.

Pen test, SAST, DAST, SCA/SBOM, fuzz, verification, MDDT-rated vulnerability management, and a dedicated managed-service team — all bundled.

COMPLIANCE-READY

FDA, ISO, IEC — built in.

CycloneDX SBOM + VEX, MDDT CVSSv3/v4, audit-ready reporting against FDA 524B, ISO 14971, IEC 81001-5-1, and AAMI TIR57.

ONE

ELTON ONE

for emerging and midsize manufacturers

WHY ELTON ONE

One-off pentests fail the regulators. It also fails the product.

01

REGULATORY

Legacy: one annual test. ONE: SPDF for the full year.

FDA's Secure Product Development Framework requires testing across premarket, release, and postmarket. ONE bundles all of it — premarket submission evidence plus 12 months of postmarket surveillance — into one fixed annual SKU.

02

KNOWLEDGE

Legacy: starts from zero. ONE: a digital twin that compounds.

Every pentest, every SBOM CVE, every disclosure feeds the same product-specific digital twin. Architecture, trust boundaries, and prior findings persist — so every new test starts where the last one ended.

03

DEPTH

Legacy: a pentest. ONE: the full testing stack.

Pen test, SAST, DAST, SCA/SBOM, fuzz, and EMB3D-derived verification testing — running on the same product, all year, for one price.

FDA Premarket Cybersecurity Guidance (2023) · IEC 81001-5-1 · ISO 14971 · AAMI TIR57

ELTON ONE

for emerging and midsize manufacturers

WHY ELTON ONE

Pentest model: rebook a quarter out. ONE: verified overnight.

04 · TESTLINK · AUTOMATED VERIFICATION

Critical CVE? Verified overnight.

TestLink runs continuously against the product's digital twin. SBOM CVEs, SAST/DAST findings, fuzz hits, and disclosed vulnerabilities are auto-triaged, mapped to the architecture, and verified at L1 (twin), L2 (code), or L3 (hardware) — included in the annual SKU.

Seconds to verify a CVE

05 · INCIDENT RESPONSE · INCLUDED

CRA 24-hour. NIS2 72-hour. Already covered.

The digital twin is already loaded and the L3 harness is already configured. AI verification in 4 hours, early warning in 24, full report in 72 — no separate retainer.

06 · LIVING REPORT

Pentest PDFs expire. ONE reports stay current.

VEX, CVSSv3/v4 ratings, finding lifecycle, and submission evidence are regenerated as the product changes — ready for eSTAR, CRA, and NIS2 audits any day of the year.

Postmarket CVEs · SBOM monitoring · SAST/DAST findings · disclosure intake — all routed through one continuous pipeline for the product.

ELTON ONE

for emerging and midsize manufacturers

WHY ELTON ONE

One fixed price. Every test the regulators expect.

07 · FIXED-PRICE SKU

No quotes. No engagements. No surprises.

Three SKUs by product type: Digital App, Physical Device, Platform/System. One annual fee covers premarket plus 12 months of postmarket — every test, every report, every CVE.

08 · FULL TESTING STACK

Everything you'd otherwise contract separately.

Pen testing, SAST, DAST, SCA/SBOM, fuzz, EMB3D-derived verification, MDDT-rated triage, VEX reporting, and a dedicated managed-service team — bundled, not bolted on.

09 · WHAT THE SKU INCLUDES

One commercial release

Premarket + 12-mo postmarket

Managed service, 7 × 24

ONE net effect: regulatory-grade cybersecurity for one product, at one predictable annual cost — without standing up an internal program.

ELTON ONE

for emerging and midsize manufacturers

HOW IT WORKS

The ONE model in five steps.

01	**SELECT**	Pick the SKU that matches the product: Digital App, Physical Device, or Platform/System.
02	**ONBOARD**	ELTON builds the product's digital twin — SBOM, architecture, trust boundaries, EMB3D PIDs.
03	**PREMARKET**	Full testing stack runs on the release: pen test, SAST, DAST, SCA/SBOM, fuzz, verification testing.
04	**SUBMIT**	eSTAR-ready cybersecurity package: threat model, SBOM/VEX (CycloneDX), MDDT-rated findings, test evidence.
05	**POSTMARKET**	12 months of continuous CVE triage, verification, CAPA tracking, and incident response — same fixed fee.

Workflow

Proven Compliance Process

HOW ELTON WORKS

Platform Workflow

1 NEW RELEASE

Software / firmware change

2 DISCOVER

Architecture, SBOM, threat model

ELTON VIEW CREATED

Point-in-time snapshot of the product

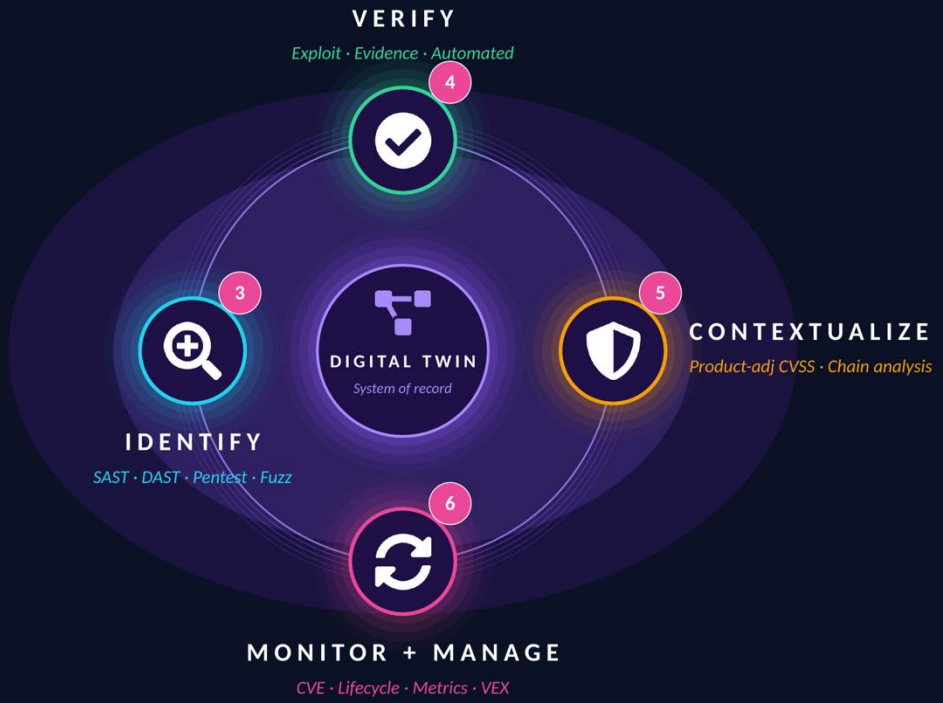

HOW ELTON WORKS

Platform Workflow

After NEW RELEASE creates a view and DISCOVER captures the digital twin, the four flywheel phases run continuously. Every stage produces evidence consumed by the next, and outputs that satisfy FDA §524B, EU MDR, CRA, and IMDRF traceability.

1 NEW RELEASE

Software / firmware change

2

DISCOVER

Context

- Components/Dataflows
- Materials (SBOM)
- Security Design & Assets

3

IDENTIFY

Vulnerabilities

- Penetration testing
- Test case results
- SCA/SBOM CVE
- SAST
- DAST
- Fuzzing

4

VERIFY

Test Case-Based

- Verify exploitability
- Document evidence
- Track lifecycle

5

CONTEXTUALIZE

Rate

- Exploitability
- CVSSv3 + CVSSv4
- Control effectiveness
- Cyber impact

6

MONITOR + MANAGE

Keep State

- New CVEs
- Software updates
- New vulnerabilities
- Release posture
- VEX export

→ *continuous loop. every test. every new CVE* →

HOW ELTON WORKS

Portfolio-wide Visibility

Real-time portfolio view across every flywheel phase — evidence in one place, current at every release.

DISCOVER | IDENTIFY | VERIFY | CONTEXTUALIZE | MONITOR + MANAGE

PORTFOLIO RISK MATRIX (illustrative)

	Critical	High	Medium	Low
Robotic v2	8	24	41	12
Proton v5.1	2	9	18	7
Imaging v2.x	4	18	35	19
Imaging v3.x	1	7	22	11
Cardiac v0.1	0	3	10	6
Ventilator v1	3	11	17	9
Wearable v2	1	5	12	4

ACTIVE NOTIFICATIONS

CRITICAL Imaging v3.x
New KEV-listed CVE in OS package; reachable via LAN

HIGH Surgical Family
BOM diff: libcurl bumped, 4 new CVEs to triage

MED Cardiac Implant
Verification re-run scheduled for v2.4 release

INFO Ventilator
Quarterly portfolio review packet ready

Digital Twin: Architecture, Materials, Controls, Impacts

Every product release is modeled as a typed digital twin. The twin is the canonical anchor for every test case, every vulnerability assignment, rating attribute, and remediation tracking downstream.

Architecture

Component graph, interfaces, dataflows, trust zones, initial-access points. Imported from architecture diagrams; refined through scanning.

Materials

Machine-readable SBOM / HBOM with NTIA-compliant fields. Continuous CVE matching against NVD, EPSS, KEV.

Controls

Implementation details for each security control; mapped to threats and required privileges. Effectiveness scored.

Impacts

Asset C/I/A ratings linked to clinical SME-defined harms; productized impact severity drives risk acceptability.

Output: a continuously-updated digital twin that answers "what does this CVE actually mean for this device?" in seconds.

Vulnerability Identification

Real test cases against the actual device, plus all candidate findings from automated tooling. Every result is bound to a twin resource and ranked.

Pen Testing

Software / Hardware / Cloud

Test Case Verification

EMB3D PID-derived cases

SAST

Source / binary / firmware

DAST

Web, REST/SOAP, mobile

SCA / SBOM

CycloneDX + VEX, KEV/EPSS

Fuzz

BLE, USB, parsers, OTA

TestLink™ — Remote / On-Demand Hardware Bench

Where ELTON lab access isn't possible TestLink is a portable testing appliance with LTE/5G uplink. Engineers can re-run any test against the device on demand without enterprise-network access. Enables continuous testing while the customer retains the device.

CONTEXTUALIZE · RATING

Product-Adjusted Ratings

Initial vulnerability CVSS scores are context-free. ELTON re-rates every CVE through the device's actual architecture, controls, and trust boundaries producing a vector that reflects real risk for that release.

Attack-Surface Reachability

Trace each CVE from every initial-access point to the affected component. Unreachable = Not Affected.

Control Effectiveness

Cumulative effect of multiple controls (auth, segmentation, encryption) on AC and AT metrics.

Trust-Boundary Traversal

PR and Scope (v3) / SC-SI-SA (v4) derived from the privilege deltas an exploit must cross.

EXAMPLE — SAME CVE, DIFFERENT CONTEXT

NVD: 9.8 Critical (context-free) → **ELTON (Release A, TLS-protected):** 5.4 Medium → **ELTON (Release B, exposed):** 9.4 Critical

Chains & Trust Boundary Analysis

Single-CVE rating is not enough. ELTON evaluates how vulnerabilities chain together and how trust boundaries amplify or contain blast radius — changing the exploitability picture release-by-release.

Vulnerability Chains

Cascade analysis: how does the exploitability of CVE-X change when CVE-Y is also present?

- Tests chains for feasibility and exploitability
- Re-evaluates AT, PR, Scope when chained
- Produces a per-chain canonical rating

Trust Boundaries

Trust zones define the privilege deltas an exploit must cross to reach an asset.

- Per-component trust level (External, User, Admin, Service)
- Identifies low-trust single points of failure
- Blast-radius bounded by zone analysis

Aligned with 2023 FDA Guidance § Vulnerability Chaining (Page 27).

MONITOR · MANAGE

Living Reports with Continuous Updates

Once a test ships, the relevant question is whether yesterday's test picture is still valid today or in 6 months. ELTON monitors against global feeds and re-evaluates against every existing in scope to keep results current for fresh reporting.

Vulnerability Monitoring

Continuous BOM monitoring against NVD, EPSS, KEV, and proprietary feeds. New CVEs against any in-device component are auto-triaged within hours.

Severity Alerting

Findings from monitoring (or new test results) are automatically rated on ingest, customers are alerted on a daily basis and provided weekly summaries.

Monitoring SLA — CRITICAL/HIGH: 30-min triage | Daily portfolio recompute | CISA / CRA / NIS2 vulnerability metrics tracked automatically

RESPOND · OPTIMALLY

Smart Fix and Prescriptive Remediation

Not every vulnerability needs a fix; some need a chain of fixes; some need a design change. ELTON helps engineers triage with context and model remediation impact before code is touched.

Smart Fix

Prioritize issues that disrupt vulnerability chains with broad impact. Identify fragile single-points-of-failure. Distinguish simple updates from intensive design changes.

What-If Analysis

Model the effect of a proposed change — update a library, harden a control, segment a network — against the digital twin before committing to it. See the new CVSS picture.

Optimized Release

Consolidate findings, threats, vulnerabilities, and risks into one decision view. Output: a prioritized work-list bound to the next release.

Net result: regulatory-defensible remediation decisions that limit unnecessary fixes.

APPENDIX 1

Autonomous Testing & Verification

Combating the AI-Vulnerability Discovery Threat

AI

AI Vulnerability Discovery Avalanche

What once required years of expertise can now be done by anyone with an LLM and a fuzzer.

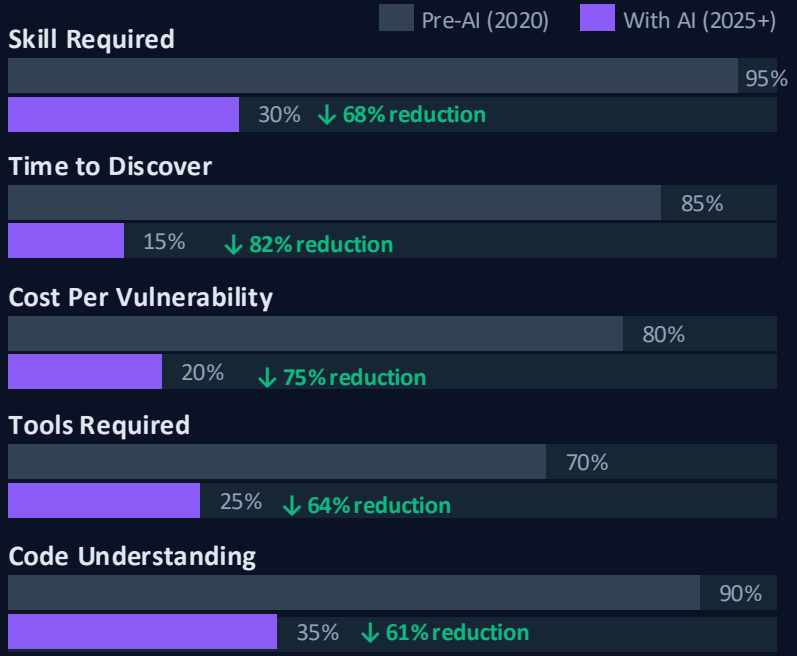

What This Means for Manufacturers

More researchers finding more bugs

The skill barrier has collapsed. A junior dev with an AI code auditor can find vulnerabilities that previously required deep expertise.

Faster discovery cycles

AI tools don't sleep. One researcher with AI tooling can audit an entire codebase in hours, not weeks.

Variant explosion

One CVE published triggers AI variant finders across every related library. One finding cascades into dozens.

Lower-quality disclosures, higher volume

As the bar drops, disclosure quality drops. More reports that are incomplete or wrong, but still require CRA-timeline response.

WHY

Autonomous Testing Is No Longer Optional

Adversary tooling is autonomous and accelerating. Defender testing must run at the same cadence — and produce verified findings, not raw alerts. Open-box is the methodology that makes that possible.

Adversaries are autonomous

AI-assisted discovery (XBOW, Big Sleep, public fuzzers) compresses weeks of reverse-engineering into hours. External researchers operate at machine speed.

Findings come from many sources

SAST, DAST, SCA, fuzzing, pen tests, CVE feeds, and researcher disclosures all produce candidate findings. Manual triage cannot keep pace.

Verification must be continuous

Every code change, SBOM update, or new CVE can shift exploitability. Periodic engagements cannot answer the question on the day it matters.

Open-box is the architectural prerequisite for all three. Without it, autonomous testing is constrained to the surface only an external attacker could see — the wrong constraint for the manufacturer.

WHY

Why Autonomous Testing Is Required

Four operational pressures push manufacturers toward autonomous testing. Each is examined in the slides that follow, with the open-box dependency made explicit.

01

Adversaries Run at Machine Speed

Public AI tooling already finds and reproduces vulnerabilities autonomously. Manual defender pace is no longer competitive.

02

CVE Volume Exceeds Human Capacity

Hundreds to thousands of CVEs per device per year. Manual triage at this volume is mathematically infeasible without automation.

03

Findings Arrive From Many Sources

SAST, DAST, SCA, fuzz, pen tests, CVE feeds, researcher reports. Each source needs the same verification — against the same architecture.

04

Verification Must Be Continuous

A finding's exploitability changes with every code, config, and SBOM change. Periodic point-in-time tests cannot answer this.

W H Y

Adversaries Are Already Autonomous

Public benchmarks 2024–2026 document the collapse of human-pace reverse-engineering as the bottleneck for vulnerability discovery. Defenders without autonomous testing fall behind on every release.

85×

XBOW v. principal pentester

40-hour assessment in 28 minutes

20

OSS 0-days from Big Sleep

One evaluation cycle, no human

6×

AI agent capability growth

Steps completed, 18 months

< 1 hr

ELTON autonomous test depth

Equivalent of 5–8-yr pentester

Implication: external researchers reach internal device behavior at machine speed. Closed-box defender testing cannot operate at that pace. Autonomous defender testing can — **but only when given the same internal access the device already grants its own runtime.**

Sources: XBOW 2025 benchmarks; Google Project Zero "Big Sleep" (CVE-2025-6965); ELTON internal testing.

WHY

CVE Volume Exceeds Human Capacity

Manual triage time per finding has not changed. CVE volume per device has. The two curves crossed several years ago for medium-complexity devices.

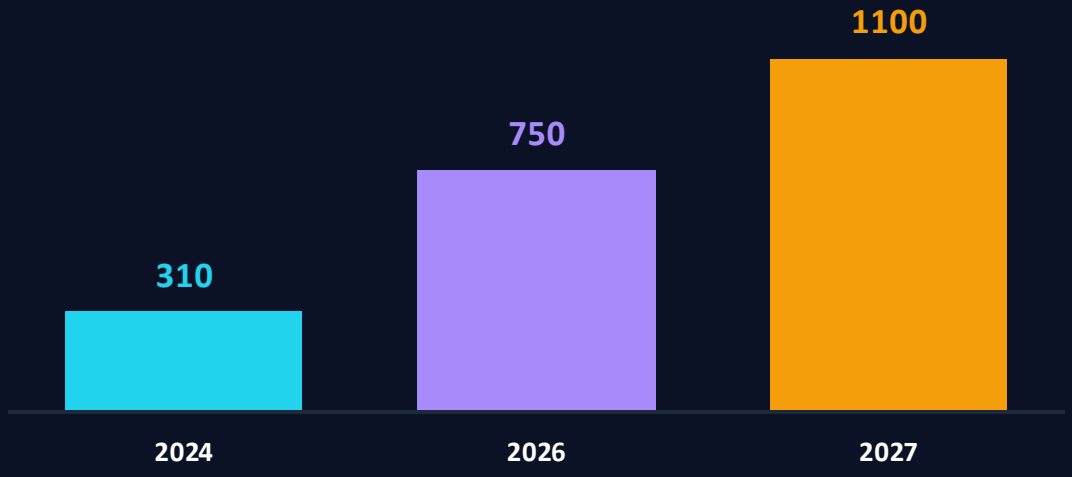

CVEs per device per year (mid-complexity medical device)

MANUAL TRIAGE MATH

~4 hr

Median engineer time per CVE to triage, justify, and document

1,100 CVEs × 4 hr =

4,400 hours / device / year

= ~2 full-time engineers per product, doing nothing but triage.

Autonomous verification reduces this 90%+ — but only when the engine has source, architecture, and SBOM context to compute exploitability without an engineer in the loop.

VERIFICATION

Unified Verification Across Sources

Modern programs run multiple discovery streams in parallel. Every finding lands as a candidate; verification determines whether it is real for this device, this release, this configuration. The verification logic is identical across sources — and depends on the same internal artifacts.

SAST

Source-code static analysis

DAST

Runtime dynamic testing

SCA / SBOM

Component CVE matching

Pen Test

Manual + autonomous probing

Fuzz

Input mutation campaigns

External

CVE feeds, researcher reports

AUTONOMOUS VERIFICATION ENGINE

Inputs: source, architecture, threat model, SBOM, credentials → Output: confirmed exploitability per finding

VERIFICATION

Verification Must Be Continuous, Not Periodic

A finding's exploitability is a function of the current code, current config, and current SBOM. Any of those change and the prior verification result expires. Annual or per-release pen tests cannot keep up.

Code change

Patch, refactor, or feature in any service or firmware module

SBOM update

Library upgrade, OS patch, or transitive dependency bump

New CVE

Public disclosure against any component already in the device

Researcher report

Coordinated disclosure or external test result

Config / deployment

Trust boundary, network topology, or default credential change

Time elapsed

EPSS / KEV signals shift even with no code change

Re-verification frequency target: every change. Manual closed-box methodology cannot meet this. Autonomous open-box can — by re-running the verification pipeline against the same internal artifacts that were used the first time.

VERIFICATION

Three-Tier Verification Model

Higher Tiers = Higher Confidence = More Not Affected Findings

LEVEL 1

Digital Twin Metadata

- SBOM + architecture docs
- Component mapping
- Reachability analysis
- Security profile overlay

Findings resolved Not Affected:

30-40%

Baseline Confidence

LEVEL 2

Code & Firmware Analysis

- Source code / firmware image
- Custom test harnesses
- Static analysis
- Emulation where possible

Findings resolved Not Affected:

50-65%

Strong Confidence

LEVEL 3

Full Product Runtime Execution

- Sample device in lab
- Live exploit verification
- Dynamic attack surface testing
- Real-world conditions

Findings resolved Not Affected:

70-85%

Highest Confidence

VERIFICATION

Automated Verification

Less Noise. Less Triage. More Signal.

A coordinated system of verification agents, deterministic validators, and the digital twin, enabled by ELTON TestLink

VERIFICATION

Proof over Probability

Creative AI discovers. Deterministic logic decides what's real.

This separation, exploration from verification, is what enables depth, scale, and trust simultaneously.

EXPLORATION

Autonomous agents explore creatively, probing every attack surface with human-like reasoning.

Network Interfaces	DICOM, HL7 FHIR, MQTT protocols probed in parallel
Firmware & Binaries	Component-level analysis of compiled device images
Authentication	Credential management, session handling, access control
Vulnerability Patterns	Known CVEs + novel attack chains explored creatively
Clinical Workflows	Device behavior tested under adversarial conditions

VERIFICATION

Findings are only accepted when exploitability is confirmed through controlled, deterministic validation.

L1: Digital Twin — 30–40%
Architecture reachability confirms component not present or code path unreachable

L2: Code & Firmware — 50–65%
Static + emulated analysis proves vulnerable code not in execution path

L3: Hardware Exploit — 70–85%
Real exploitation attempt on actual device hardware — the definitive proof

If it can't be proven, it's not a vulnerability.

APPENDIX 2

ELTON TestLink Appliance

Overnight Automation

ELTON TESTLINK

On-Premise Cybersecurity Testing Appliance

TestLink is a portable hardware appliance shipped to the manufacturer's site, paired with the device under test, and connected to ELTON SaaS over 5G/LTE. The customer keeps the product where it lives; ELTON gets the test access it needs.

On-Demand Verification

Re-run any test against any released firmware or build at any time. Verify a CVE in minutes instead of scheduling a new pentest engagement.

Continuous Autonomous Testing

Runs ELTON's autonomous test engine against the live device 24/7, not just during scheduled engagements. New CVE → verified result in hours.

Safe Remote SME Access

5G/LTE uplink lets ELTON engineers evaluate the target remotely with no enterprise-network access required. No VPN, no firewall changes, no IT tickets.

Full Hardware Test Surface

Direct connection into USB, network, JTAG, UART, serial, GPIO, and other physical interfaces — the same surface a sophisticated attacker would use.

TestLink is included with every ELTON ONE subscription. ENTERPRISE deployments scale TestLink across the regulated portfolio.

Direct Hardware Test Access

TestLink connects to the device's actual physical and logical interfaces with a customized cable and connector kit. Tests run against the real hardware — not a simulator, not a virtual representation.

USB

USB-A / USB-C / micro / mini

Network

Ethernet (RJ45), 1/10G, SFP

JTAG

Standard 20-pin and custom pinouts

UART

TTL / RS-232 / RS-485

Serial

SPI, I2C, CAN bus

GPIO

Pin headers + breakout cabling

Wireless

BLE, Wi-Fi, Zigbee (via SDR)

Debug

Console, single-board access

Customized to the product — remotely, or in a single on-site visit

Cables, connectors, and adapters are built per device. After commissioning, TestLink can stay connected indefinitely or be unplugged and reconnected when testing is needed.

TESTLINK · DEPLOYMENT

Persistent Access for Real-World Response

TestLink stays paired with the customer's device on the customer's premises. The engineering team can swap firmware, switch releases, or reproduce a customer-reported issue — and ELTON can verify against any of those states in real time.

Continuous Pentest

Autonomous test engine runs 24/7. New CVE drops at 3am? Verified result is waiting in the morning, against the actual current build.

Incident Response

Researcher discloses a vuln, or a customer reports an anomaly. Engineering loads the affected firmware on the device; ELTON verifies and triages within hours.

Multi-Release Triage

Same vulnerability behaves differently across v2.1, v2.4, and v3.0. Engineering can stage each release on the device; TestLink lets ELTON evaluate every one.

Plug-and-leave or plug-when-needed: *TestLink can persist on the device indefinitely, or be disconnected and reconnected only when testing is required — customer's choice.*

Why TestLink Beats Legacy Testing

Shipping a unit to ELTON's lab works for one-shot, point-in-time engagements. It does not work when the question is what the engineering team built last week, or targeting the release a hospital is actually running.

Capability	Legacy Testing	TestLink On-Site
Engineering team support during testing	No	Yes
Triage same finding across releases	No	Yes
On-demand verification of new CVE	Schedule + ship	Hours
Continuous testing posture	No — point-in-time	Yes — 24/7
Incident-response turnaround	Days to weeks	Hours
Customer retains the device	No	Yes
Enterprise-network access required	N/A	No — 5G/LTE
Suitable for clinical-environment testing	No	Yes

Lab shipment remains valid for one-time premarket engagements. Postmarket and continuous-testing programs need TestLink.

APPENDIX 3

Product-Adjusted Ratings (CVSSv3 + CVSSv4)

FDA-Compliant & Defensible Vulnerability Severity

The Unmanageable Triage Burden on Manufacturers

Per-product vulnerability counts are scaling with overall CVE growth. Without automation, manual triage becomes physically impossible.

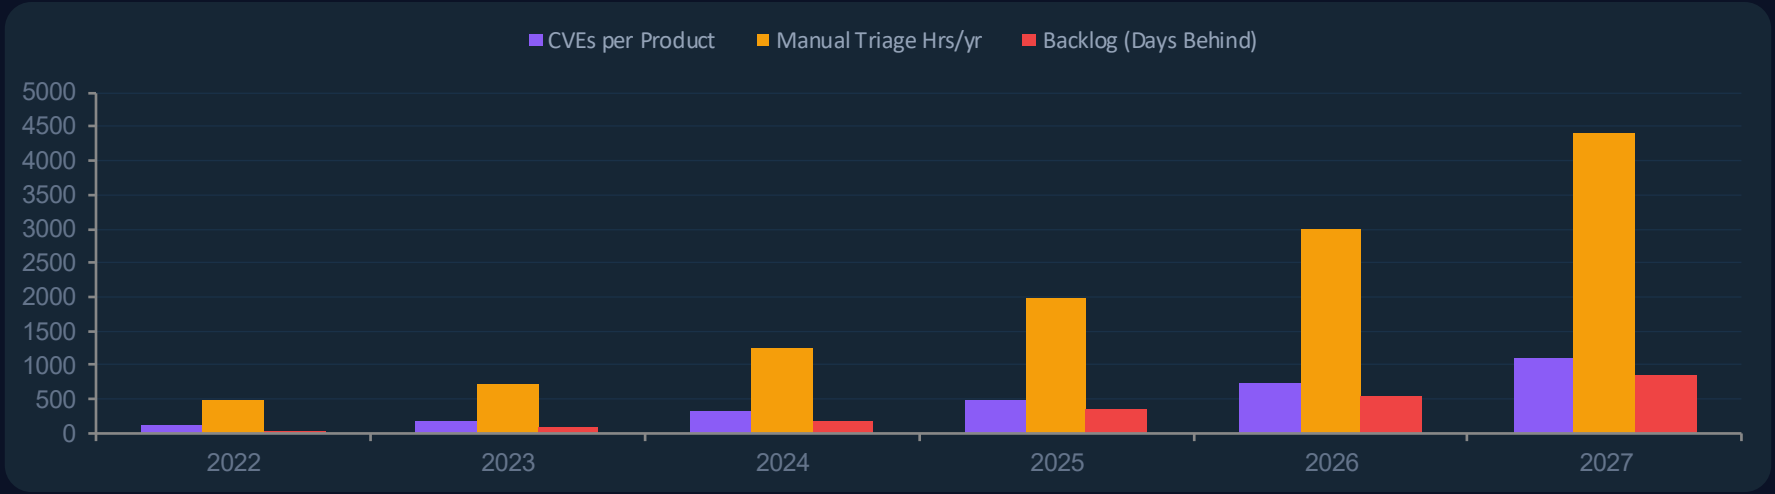

Without ELTON (Manual Triage)

2024: 310 CVEs × 4 hrs = **1,240 hrs/yr**

2026: 750 CVEs × 4 hrs = **3,000 hrs/yr**

2027: 1,100 CVEs × 4 hrs = **4,400 hrs/yr**

2+ FTE security engineers per product, just for triage.

With ELTON (Automated Triage)

2024: 310 auto-triaged, 45 human = **180 hrs/yr**

2026: 750 auto-triaged, 85 human = **340 hrs/yr**

2027: 1,100 auto-triaged, 110 human = **440 hrs/yr**

90% reduction. Scales with CVE volume without adding headcount.

PRODUCT-ADJUSTED

Digital Twin Changes Everything

Legacy testing engagements start from zero every time. ELTON's knowledge compounds with every assessment.

Traditional Pentest Model

Each test starts fresh. No persistent context.

ELTON Compound Model

Digital twin enriches. Every CVE builds context.

Pre-Loaded Context

Agents start with full architecture, SBOM, and security profile. Zero recon time.

Graduated Confidence

Three tiers of increasing evidentiary weight, not binary exploit/no-exploit.

Continuous Loop

New CVE → auto-triage → agent re-verification. Not periodic re-testing.

Regulatory Evidence

Every test generates VEX, audit trail, and compliance artifacts automatically.

PRODUCT-ADJUSTED

FDA-Compliant Ratings

ELTON uses the **FDA-Qualified MITRE CVSS Rubric for Medical Devices (MDDT)** to generate product-adjusted vulnerability scores. The methodology is rules-based, device-contextual, and natively embedded in the platform, replacing inconsistent SME opinion with defensible, traceable, FDA-recognized ratings.

FDA-Qualified MDDT Methodology

- MITRE rubric qualified as MDDT by FDA (Oct 2020)
- Structured decision trees elect each CVSS vector value
- Scoring reflects device context — not generic IT assumptions
- Every decision is traceable and suitable for FDA submissions

Native to ELTON

- Digital twin maps components, data flows, trust boundaries
- MDDT logic applied automatically to every CVE and finding
- Eliminates SME variability — consistent, reproducible scores
- Audit-ready: every rating backed by verifiable device artifacts

HOW ELTON PRODUCT-ADJUSTS

Ingest CVE
NVD base score enters as generic CRITICAL / HIGH

→

Map to Twin
Locate affected component in product graph

→

Apply Rubric
MDDT decision trees elect each CVSS vector value

→

Adjusted Score
Defensible, traceable product-specific rating

PRODUCT-ADJUSTED

Example Product-Adjusted Rating

A CVE enters as **CRITICAL (9.8)**. ELTON maps to the product graph, identifies a countermeasure, and adjusts the AttackVector to produce **LOW (2.4)**.

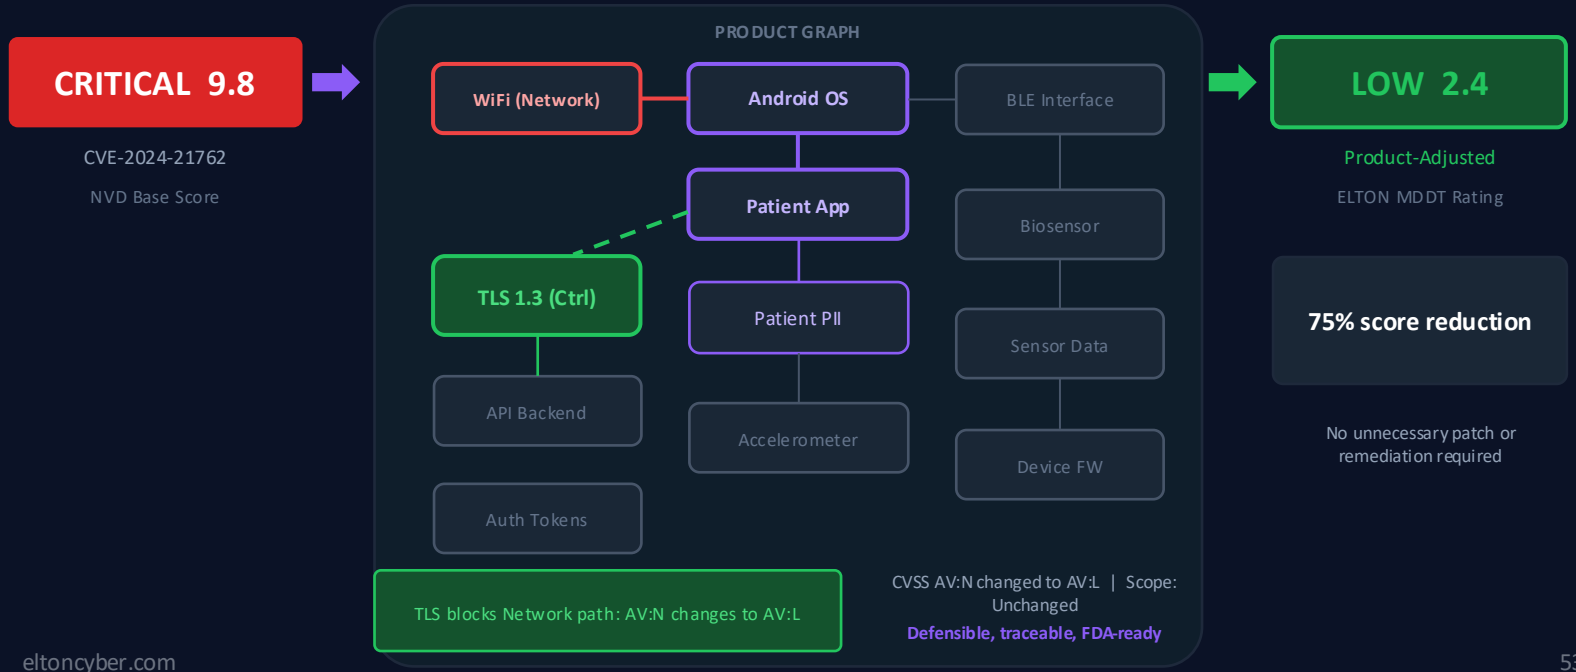

CVSSv4

Why CVSSv4 Matters

FDA phases out CVSSv3 by December 2026. CVSSv4 requires architecture-level analysis that cannot be computed manually at scale.

Dec 2026

FDA CVSSv3
Phaseout Deadline

4 New

Metrics Requiring
Architecture Data

100s

Vulns Per Product
Need Rescoring

0%

Can Be Done
Manually at Scale

CVSSv4 introduces Attack Requirements (AT), Subsequent System Impact (SC/SI/SA), and granular User Interaction (Active/Passive/None) — all requiring deep knowledge of the device architecture. Only a digital twin approach can automate this.

CVSSv4

CVSSv3 vs. CVSSv4: What Changed

Metric	CVSSv3.1	CVSSv4.0
Attack Vector (AV)	N / A / L / P	N / A / L / P (unchanged)
Attack Complexity (AC)	Low / High	Low / High (unchanged)
Attack Requirements (AT)	Does not exist	**None / Present (NEW)**
Privileges Required (PR)	None / Low / High	None / Low / High (unchanged)
User Interaction (UI)	None / Required	**None / Passive / Active (expanded)**
Scope (S)	Unchanged / Changed	**Removed (replaced by SC/SI/SA)**
Vulnerable System CIA	C / I / A	VC / VI / VA (renamed)
Subsequent System CIA	Does not exist	**SC / SI / SA (NEW)**

CVSSv4 adds 4 new dimensions that require product architecture context to compute. ELTON handles all of them automatically.

CVSSv4

Why CVSSv4 Can't Be Scored (Easily) Manually

CVSSv3 metrics could be answered from the CVE description alone. CVSSv4 introduces metrics that require a complete, architecture-aware understanding of the product's internal topology.

Attack Requirements (AT)

"Can the attacker reach it without relying on any deployment-specific element?"

- Enumerate every route from entry point to the vulnerable component
- Remove configurable nodes; re-evaluate surviving paths
- One non-configurable route surviving = AT None; all depend on config = AT Present

Combinatorial path analysis over a directed graph — routes grow exponentially with product complexity.

Subsequent Impact (SC/SI/SA)

"What can the attacker reach after exploitation — beyond the corridor?"

- Map the blast radius: components reachable via post-compromise trust augmentation
- Subtract the corridor (already scored under Vulnerable System CIA)
- Take highest severity per CIA dimension across remaining assets

Blast radius is conditional per vulnerability — misclassifying one node skews the entire assessment.

CVSSv4

Both Metrics Require a System Model

Without graph infrastructure, analysts are left manually tracing paths through architecture diagrams — introducing inconsistency at every step.

✗ Manual Scoring

- Requires architecture diagrams
- Judgment calls on every node
- Inconsistent across analysts

⚠ At Scale

- Paths grow exponentially
- Blast radius changes per CVE
- One missed edge changes the score

✓ ELTON Digital Twin

- Full system graph built-in
- Automated path & blast-radius analysis
- Consistent, auditable results

CVSSv3 → Read the advisory, pick values from a dropdown.
CVSSv4 → Build & query a full system model, or get it wrong.

ELTON's digital twin is that system model — purpose-built for exactly this class of computation.

COMPLIANT RATINGS

What the Scoring Engine Does

Given a product's architecture diagram, an attacker entry point, and a target vulnerability, the engine automatically computes CVSS 3.1 and CVSS 4.0 scores that reflect the specific context of that product.

Incoming Paths (The Corridor)

Can the attacker reach the vulnerability?

All feasible routes from the attacker's entry point to the vulnerable component. The engine prunes irrelevant branches and enforces trust constraints automatically.

Combinatorial path analysis that scales with graph size — impossible to do manually for non-trivial products.

Outgoing Paths (Blast Radius)

What can the attacker do after exploiting it?

All components reachable from the target after exploitation. Conditional trust augmentation based on the vulnerability's technical impact classification.

Re-evaluating forward reachability under different trust assumptions for each vulnerability — manual analysis cannot scale.

COMPLIANT RATINGS

How ELTON Computes Each Metric

Exploitability Metrics — Derived from Incoming Path Analysis

Attack Vector (AV)

Network > Adjacent > Local > Physical

ELTON traces the attack path from the entry point to the vulnerable component. It picks the most accessible interface along that path (e.g., network beats local). If there is no remote interface, it defaults to physical access.

Attack Complexity (AC)

Low or High

ELTON checks whether a security control (e.g., firewall, authentication) sits between the attacker and the vulnerability. If yes, complexity is High. It also factors in real-world exploit data: if EPSS is above 10% or active exploitation is confirmed, complexity is Low.

Attack Requirements (AT)

None or Present (CVSSv4 Only)

ELTON identifies which components on the attack path are configurable (e.g., can be disabled or hardened by the user). It then removes those components and checks: does a viable path still exist? If yes, AT = None. If every route depends on something configurable, AT = Present.

COMPLIANT RATINGS

Exploitability & Scope Metrics

Privileges Required (PR)

None / Low / High

ELTON maps the trust levels across every zone on the attack path, from the entry point to the vulnerable component. It compares what privileges the attacker starts with against what each zone requires, then picks the highest privilege level needed to traverse the full path.

User Interaction (UI)

v3.1: None/Required
v4.0: None/Passive/Active

ELTON reads the user-interaction attribute set on the entry point. In CVSSv3.1, any interaction was simply “Required.” CVSSv4 splits this into Passive (e.g., opening a file) vs. Active (e.g., entering credentials), so ELTON preserves that distinction when scoring for v4.

Scope (S) — v3.1 Only

Unchanged / Changed

ELTON checks whether a successful exploit lets the attacker break out of the vulnerable component's trust zone and affect other parts of the system. If the attacker can cross a trust boundary and compromise downstream components, Scope is Changed. Otherwise, it's Unchanged.

COMPLIANT RATINGS

Impact Metrics: Vulnerable & Subsequent Systems

Vulnerable System CIA

CVSSv3.1: C / I / A | CVSSv4.0: VC / VI / VA

Collect all assets on the target component. Each asset has independent C/I/A sensitivity ratings (High, Low, None).

For each dimension, the engine selects the highest severity rating among all assets. No assets = None.

Repeated across many vulnerability-scenario combinations with different asset profiles — bookkeeping is the primary source of human error.

Subsequent System CIA

CVSSv4.0 Only: SC / SI / SA (NEW)

Take post-compromise blast radius and subtract the corridor. Remaining nodes are components reachable only after exploitation.

Collect all assets from these subsequent components; select highest severity per dimension.

Misclassifying corridor vs. blast-radius-only nodes skews the assessment. This multi-step dependency chain is where manual analysis most commonly introduces errors.

APPENDIX 4

Vulnerability Lifecycle Management

Automating MTTT, MTTV, and MTTM

LIFECYCLE TRACKING

Validate. Rate. Track.

Three-tier deterministic validation feeds device-specific CVSS ratings. Every finding enters a VEX-aligned lifecycle with continuous compliance evidence.

VALIDATION

L1: Twin Analysis
Architecture reachability
30–40% Not Affected

L2: Code Verification
Static + emulated analysis
50–65% Not Affected

L3: Hardware Exploit
Real device exploitation
70–85% Not Affected

ADJUSTED RATINGS

CVSSv3 Re-Rating FDA-recognized MDDT methodology

CVSSv4 Automation
AT/SC/SI/SA metrics from twin

Device Context Overlay Clinical use + architecture adjust

Evidence-Backed Scores
Traceable from CVE to rating

Portfolio Aggregation Cross-product risk roll-up

LIFECYCLE TRACKING

VEX Generation
CycloneDX VEX as natural byproduct of every assessment

Finding Lifecycle
CISA-aligned states: Draft → Triage → Affected/Not Affected

Continuous Compliance
FDA 524B, CRA, NIS2, IMDRF audit-ready at all times

MTTT / MTTR Metrics
Time-to-triage and time-to-resolution tracked per finding

LIFECYCLE TRACKING

What is a View?

A View is ELTON's frozen snapshot of a product's architecture and software — the context every vulnerability is scored, verified, and tracked against.

A View captures a product state — architecture, SBOM, configuration, interfaces, trust boundaries and code — so ELTON can **contextualize every vulnerability** against the real software, not a generic component list. Findings, ratings, verification results and metrics are all bound to a View.

DRAFT **Floating during development**

The architecture and software are still evolving. Findings resolve fast in the same view — CI/CD-velocity vulnerability management against a moving build.

RELEASE **Locked to a shipped version**

The architecture and build are frozen. Findings accrue cross-view metrics (MTTR / MTTM) and generate defensible postmarket evidence per SKU in the field.

LIFECYCLE TRACKING

Every finding has a version. Every metric has a view.

Medical device software is built with iterations of release and ultimately ships in discrete releases, but vulnerability management must be continuous across both phases of the product lifecycle. ELTON bridges the two by binding every finding to a view, allowing a finding to iterate through its lifecycle during development, and evidence which release closed a vulnerability in postmarket.

1 : 1

finding binds to exactly one view

3

metrics tracked per finding:
MTTT · MTTV · MTTM

4

compound states combining lifecycle × verification

LIFECYCLE TRACKING

Every finding has a version. Every metric has a view.

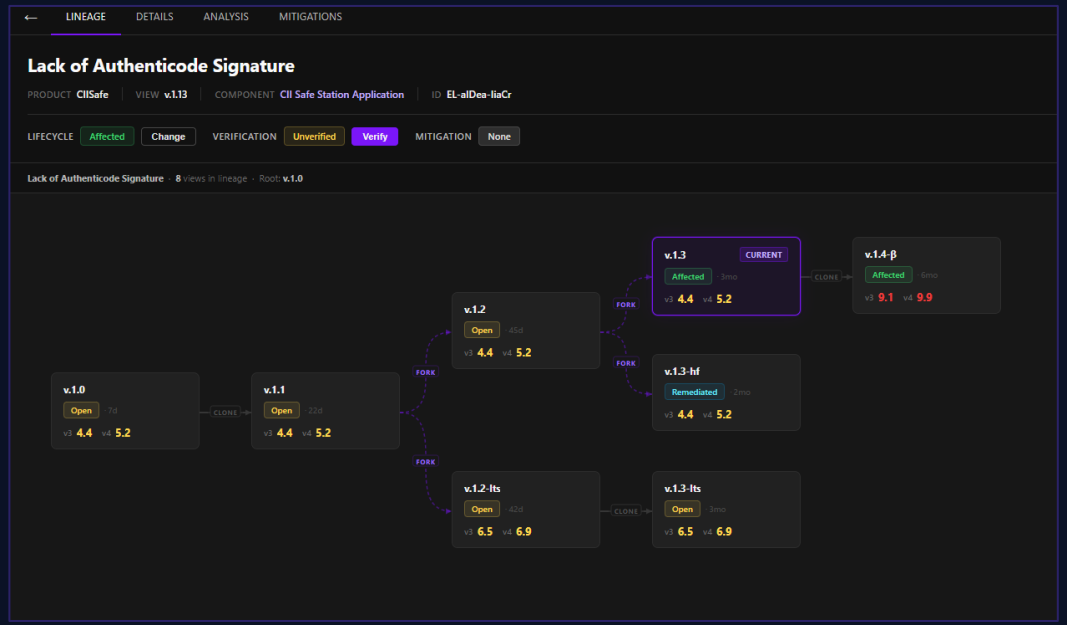

The same finding can be open in view v1.1, patched in view v1.3, and under review in view v1.2, each with their own accurate, independently-calculated metrics.

Findings Are View-Specific Entities

Same finding, same asset — scoped to the view it lives in.

One finding, one view

Every instance of a finding belongs to one view.

Its lifecycle status, rating, verification evidence and VEX reason describe that view only.

Enables per-release finding traceability.

Movement across views is a clone

When software evolves, findings propagate by cloning and linking, not moving.

The source finding is preserved. The clone starts fresh in the target view: Needs Triage (New), Not Verified.

History links the chain. Metrics start new.

This is why the same CVE can show different status, different ratings, and different metrics in two views at the same time — and that is correct.

LIFECYCLE TRACKING

Three Modes. Three Jobs.

View status is a property of the view — it governs every finding inside it.

DRAFT

Active development

- All operations permitted
- Findings continuously rated
- Architecture / twin editable
- All finding status changes allowed, such as “Mitigated” in same view
- Same-view MTTM metrics for active development

Every commit can introduce new findings without creating new views.

RELEASE

Shipped version

- Digital twin frozen
- Findings continuously rated
- Triage / Affected / Not Affected finding status
- No “Mitigated” transition
- Finding closure through clone to new view

Guarantees you know in which future release (view) a finding was later resolved

ARCHIVED

Historical record

- Read-only. No modifications
- No finding updates
- Outbound clones permitted
- Valid MTTM resolution endpoint
- Audit & compliance preserved

EOL releases (views) versions stay anchored in the audit trail.

Draft vs. Release: What Each View Mode Allows

View Mode	“Draft”	“Release”
Architecture / Digital Twin Changes	Fully editable	Locked — frozen software version
New Findings Opened	Allowed	Allowed
Status: Triage/Affected/Not Affected	Allowed	Allowed
Affected → Mitigated	Same-view (active development)	Must close finding in a new release (view)
Clone Operation	Allowed	Allowed
MTTM Calculation	Mitigated in the same dev view	Mitigated in a future view
Primary Use Case	Active CI/CD development	Shipped release / field compliance

Archived views are read-only. They remain valid resolution endpoints for MTTM lookups from other views.

LIFECYCLE TRACKING

Three Metrics. One Finding. Per View.

Every metric is calculated from the perspective of a specific view.

MTTT

Mean Time to Triage

START

Needs Triage (New)

END

First disposition

SCOPE

Same-view. Same calc in Draft and Release.

From 'we know about this' to 'we decided what it is'.

MTTV

Mean Time to Verify

START

Exits Needs Triage

END

Verification Status → Verified

SCOPE

Same-view. Cross-views get fresh Retesting Tests.

From 'triaged' to 'backed by test evidence'.

MTTM

Mean Time to Mitigate

START

Affected in this view

END

Mitigated in this view or by a clone in a new view

SCOPE

Draft: same-view. Release: cross-view.

From 'open and real' to 'closed and deployed'.

How MTTM Is Calculated in Each Mode

Mean Time to Mitigate — always view-relative, calculated differently by view status.

DRAFT VIEW · Same-view fast path

MTTM = Mitigated − Affected (this view)

- MTTM Start: verified as Affected in this view.
- MTTM End: verified as Mitigated in this view.
- Single value per finding.
- No cross-view mitigation while in draft status.

Optimized for premarket development velocity — commit-triggered scans can close findings without a new release.

RELEASED VIEW · Cross-view compliance

MTTM = Mitigated(clone) − Affected(this)

- MTTM Start: Affected in THIS view (not an ancestor).
- MTTM End: a clone is Mitigated downstream.
- Multiple values — one per resolution path.
- Aggregated as min / mean / unresolved.

Gives each shipped version an accurate exposure window and tracking from discovered-in view to fixed-in view.

Verification Testing: Three Tiers, One Outcome

Every finding can be linked to one or more test cases. PASS / FAIL simultaneously sets lifecycle and verification.

Verification Tiers

L1 — **Static Reachability**
Digital-twin analysis. Component presence, architecture, and reachability — no runtime required.

L2 — **Binary / Code**
Binary or source-code inspection. Confirms vulnerable code is (or isn't) compiled into the build.

L3 — **Live Exploit / Runtime**
Device or emulated runtime. Executes exploit, observes behavior. Highest evidence weight.

Outcome → Status

No Verification Possible
Affected · Not Verified

Verification Completed – Test Case Failed
Affected · Verified

Verification Completed – Test Case Passed
Not Affected · Verified

eltoncyber.com

The Six Lifecycle Statuses

Every finding in every view passes through one of these six lifecycle states.

Identification

Draft state for USER-created findings not yet ready for triage. Invisible to automation.

SET BY: User only **CLOCKS:** None

Needs Triage

Entry state for all findings. Sub-reasons: New / Updated.

SET BY: User or System **CLOCKS:** MTTT starts

Under Investigation

User-driven 'parking' status. Active review without rating change. Existing clocks continue.

SET BY: User only **CLOCKS:** MTTT Continue

Affected

Confirmed open and relevant. Strongest open-posture state.

SET BY: System or User **CLOCKS:** MTTT end · MTTM start · MTTV start

Not Affected

Closed as not affected. VEX-compliant reason REQUIRED.

SET BY: System or User **CLOCKS:** MTTT end · MTTM n/a · MTTMV start

Fixed

For views of status “release” must be Fixed in a subsequent release (view). Cannot be Fixed in the originating view. View of status “draft” allows same-view Fixed.

SET BY: User only **CLOCKS:** MTTR ends

Sub-Reasons, Closure Paths & Enforcement

Two triage sub-reasons. Two valid ways to close. Five rules the platform enforces.

Needs Triage → Sub-Reasons

NEW — First entry to the view from any source. Auto-selected for ELTON-AI and cloned findings.

UPDATED — Summary, technical detail, remediation, or assigned resource changed. Auto-selected by ELTON when metadata changes.

Closure Paths

FIXED — Resolved in a subsequent view only. Blocked in the originating (birth) view.

FALSE POSITIVE — Finding was invalid from the start. Available in any view (birth or cloned).

FIVE ENFORCED RULES

- Identification → Needs Triage is the only valid forward transition from draft.
- Not Affected requires a VEX reason. Inline mitigations requires Mitigation = Confirmed.
- Fixed is blocked in a view of status “release”. Closure happens in a subsequent view.
- Rating changes do not revert findings to Needs Triage.
- Once started, MTTT and MTTR clocks are never reset by re-assessment.

Not Affected — The Five VEX Reasons

Not Affected closure must carry a VEX-compliant reason. ELTON enforces the choice and exports it.

1	**Component not present**	Vulnerable component not present in the product's digital twin.
2	**Vulnerable code not present**	Component present but the specific vulnerable code is not included in the product build.
3	**Vulnerable code not in execute path**	Vulnerable code present but never reached during execution. Code path unreachable given control flow, config, or architecture.
4	**Vulnerable code cannot be controlled by adversary**	Code present and reachable, but the adversary cannot control inputs/conditions. Attack preconditions cannot be met.
5	**Inline mitigations already exist**	Vulnerability legitimate and code reachable, but a countermeasure prevents exploitation. Requires Mitigation Status = Confirmed. SPECIAL CASE

If verified by a passing test case, ELTON copies the VEX justification. No continued monitoring on Not Affected.

Verification & Mitigation Status

Two Verification & Mitigation statuses ride alongside lifecycle status. Each has its own rules.

Verification Status

Two states · Set by linking test case results

Not Verified

Default state. No test case result associated with the finding.

Verified

One or more test case results linked via the 'Verify' button. Carries lifecycle weight.

Mitigation Status

Three states · Set by ELTON Rating Analysis + user confirmation

None	No countermeasure on path. Default state. No gating effect on lifecycle.
Potential	Countermeasure identified on path but not confirmed. Does NOT satisfy 'Inline mitigations' VEX.
Confirmed	Countermeasure confirmed effective. Prerequisite for 'Inline mitigations' VEX reason.

LIFECYCLE TRACKING

How a Finding Propagates Across Releases

The clone operation is the only way a finding crosses views — and it resets the finding for the new context.

1 **Clone from source view**
User or system initiates

2 **Create clone in new**
Needs Triage (New)

3 **Pin Active Score**
Source rating carried forward

4 **Create Retesting Test**
Fresh test, new results

5 **Run auto-triage**
Rating + Verification run

6 **Record history link**
'Cloned From' / 'Cloned To'

Source is preserved

The original finding keeps its history, rating and metrics. It doesn't disappear when cloned — it remains the record of that software version.

Clone is a fresh start

The clone begins in Needs Triage (New), Not Verified, with an Active Score pinned from source. Its MTTT, MTTV and MTTM are its own.

eltoncyber.com

LIFECYCLE TRACKING

Four States Customers Report and Track

Every active finding sits in one of these. Reportable in the view dashboard and VEX export.

A + NV

Affected Not Verified

Open finding. Determination rests on analysis — no test evidence.

Minimize

A + V

Affected Verified

Open finding. FAIL test case linked as evidence of exposure.

Strongest compliance posture for open findings

M + NV

Mitigated Not Verified

Closed, but no test confirms the fix works. Analysis-only closure.

Minimize

M + V

Mitigated Verified

Closed. PASS test case linked as evidence the fix is effective.

Maximize

LIFECYCLE TRACKING

One Test Result. Two Dimensions Updated.

Verification Testing outcomes simultaneously set Lifecycle and Verification statuses — and a VEX reason when applicable.

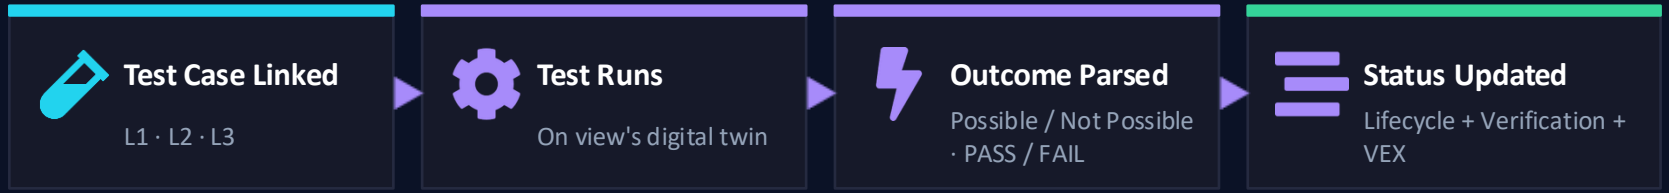

Not possible

Affected · Not Verified

No runtime. VEX reason not required.

FAIL

Affected · Verified

FAIL proves exposure. Strongest open-posture.

PASS

Not Affected · Verified

Auto-assigned best-fit VEX reason.

Which Metric Clock Runs on Which Transition

Once started, clocks never reset on re-assessment or rating updates.

Transition	MTTT	MTTM	MTTV
Identification → Needs Triage	Starts	—	—
Needs Triage → Under Investigation	Stops	—	Starts
Needs Triage → Affected	Stops	Starts	Starts
Needs Triage → Not Affected	Stops	—	Starts
Affected → Mitigated	—	Stops	—
Not Verified → Verified	—	—	Stops

Metric stability — once started, MTTT and MTTM clocks never reset. Verification resets on lifecycle change.

LIFECYCLE TRACKING

CVE-2025-1234 Across Three Views

Same CVE. Three views. Three different sets of metrics — each correct for its own version.

View A · v1.1 RELEASED

- Day 0: scan finds CVE
- Day 0+4m: Affected + Verified
- Day 180: view Released
- Day 270: cloned to v1.2/v1.3

View B · v1.2 DRAFT

- Day 270: clone → Needs Triage
- Day 270+6m: FAIL → Affected+V
- Patch not ready
- Remains Affected + Verified

View C · v1.3 DRAFT

- Day 270: clone → Needs Triage
- Day 270+5m: PASS → Not Affected
- Day 271: Mitigated + Verified
- Same-view MTTM = 1 day

Metric	View A (v1.1)	View B (v1.2)	View C (v1.3)
MTTT	4 min	6 min	5 min
MTTV	0 min (auto)	0 min (auto)	0 min (auto)
MTTM	271 days (via clone in C)	Unresolved	1 day (same-view)
Compound State	Affected + Verified	Affected + Verified	Mitigated + Verified

Appendix 5

Methodology & Approach

Regulatory Compliance with Open-Box Testing

Open-Box = Defensible Output

INPUTS

- **Architecture** — DFD, interfaces, zones
- **SBOM / HBOM** — Components, libs, firmware
- **Source / Binary** — Build env, dependencies
- **Test access** — Creds, debug, JTAG
- **External feeds** — CVE, KEV, EPSS, reports

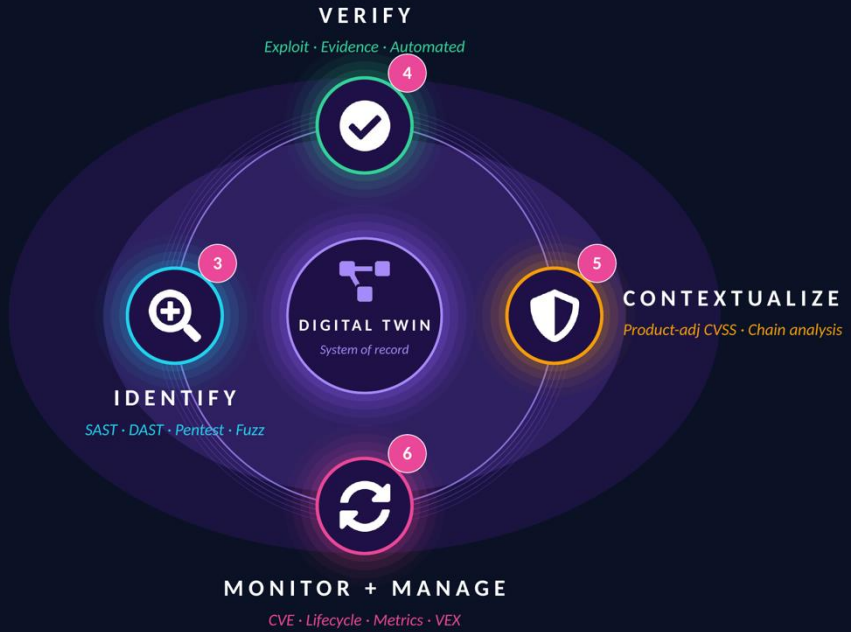

OUTPUTS

- **Verified findings** — True positives only
- **Attack graphs** — Per IA, all paths
- **Dual CVSS** — v3 + v4 + justify
- **VEX / CycloneDX** — Continuously emitted
- **Regulatory packet** — FDA, EU MDR, CRA

CONTINUOUS LOOP — *every code change, every SBOM update, every new CVE re-runs the pipeline against the same digital twin. Output regenerated automatically.*

WHY OPEN-BOX

Open-Box ≠ More Findings

Open-box access produces more test-case results, not more findings. The manufacturer chooses what is, and what is not, a finding in the regulatory report.

TEST CASES

what is exercised

- Open-box increases coverage
- Closed-box leaves gaps
- Per-IA, per-asset, per-component

RESULTS

what each test produced

- Pass · Partial · Fail
- Not applicable
- Evidence captured for each

TRIAGE

manufacturer-driven

- Reachability check
- Compensating control review
- False-positive removal

FINDINGS

what reaches the report

- Curated, justified
- Bound to twin resources
- Goes into regulatory submission

Volume of test results ≠ volume of findings. *More open-box access → more confidence in coverage. The manufacturer's triage decides what is reportable, what is risk-accepted, and what is not affected.*

DEFINITIONS

Testing Methodologies

Product security industry distinguishes three methodologies by the level of information provided to the tester. The distinction directly determines coverage, cost, and finding depth.

Closed-Box *(black-box)*

Tester receives no prior knowledge of the target. Simulates an external threat actor performing full reconnaissance. Can't verify findings automatically.

Access provided: none

Grey-Box

Partial information provided — typically credentials of an authenticated user. Simulates an insider or compromised account.

Access provided: partial

Open-Box *(white-box)*

Full documentation, source code, architecture diagrams, threat model, SBOM, and privileged credentials are provided. Automated CVE verification.

Access provided: complete

This deck addresses the methodology choice between closed-box and open-box testing for premarket and postmarket medical device assessments.

ANATOMY

Where a 4wk Assessment Spends Its Time

A standard penetration test is a fixed-time engagement. Five of the six phases have fixed duration. Phase 3, vulnerability identification, is the only variable. Methodology choice determines how much of that phase is consumed by reconnaissance versus exploitation.

1	2	3	4	5	6
Test Plan	**Attack Surface**	**Identify Weaknesses**	**Exploit Single Issues**	**Chain & Bypass**	**Reporting**
1 days	*1 days*	*Variable*	*2 days*	*2 days*	*Remainder*
Scope, constraints, threat focus	Interfaces, entry points, workflows	Reverse Engineering vs. Discovery split driven by access	Verify known vulnerabilities	Multi-step exploitation	Findings, evidence, recommendations

Phase 3 under closed-box: ~60–70% of hours spent on reverse engineering/reconnaissance before any vulnerability identification and exploitation begins.
Phase 3 under open-box: reconnaissance compressed; no reversing, remaining hours redirected to vulnerability identification, chaining, and bypass.

FIRST PRINCIPLES

The Cybersecurity Equation: Time vs. Money

Security posture is the balance between attacker resources and defender resources applied to the same target. The attacker operates under asymmetric constraints. Methodology choice determines whether the manufacturer narrows or widens that asymmetry.

Attacker / Researcher

Time budget: unbounded — weeks to years, major actors can have large finances that translate to unbound time

Starting data: whatever can be purchased, extracted, or reverse-engineered

Tooling: full access to AI-assisted discovery (XBOW, Big Sleep, public fuzzers)

Disclosure control: researcher decides severity, timeline, and public framing — **FDA/CISA side with researcher**

Manufacturer / Defender

Time budget: fixed; driven by release schedule and submission deadline, limited by willingness to pay a certain fee for a test

Starting data: complete — unless methodology choice withholds it from the test team

Tooling: same AI-assisted tooling available, applied with architectural context for proper ratings

Disclosure control: retained when findings surface internally first, disclosed to FDA when determined necessary

Closed-box testing voluntarily withholds the manufacturer's only structural advantage.

BENCHMARK

Peer Manufacturer Methodology Distribution

ELTON-observed pattern across the medical device manufacturers we support. The premarket/postmarket split shows the industry has already converged on open-box as the default for fielded products but lags on premarket adoption. Since 2020 nearly all premarket tests are now open-box and most postmarket are open-box.

PREMARKET

< 5%

of peer manufacturers rely on closed-box testing for premarket submissions.

The remainder provide source code, architecture, threat model, and privileged credentials. Decreased from 25% in 2020.

POSTMARKET

~85%

of manufacturer-led postmarket testing is executed open-box.

Once a device is fielded, cost of missed findings rises sharply — access defaults to full. Increased from 50% in 2020.

Observation: From ELTON execution across customer base, with specific trends increasing rapid since 2023.

PART 1

Premarket: Four Technical Arguments

The arguments below each derive from FDA guidance, FD&C §524B, or the cost-of-defect literature. Each point is examined in the following slides with regulatory citations.

01

FDA SPDF Requires It

The Secure Product Development Framework requires iterative, in-development testing integrated with the development process. Single-pass closed-box testing does not satisfy the framework.

02

Manufacturer Testing Scope

FDA expects the manufacturer to test its own product to the fullest extent. Closed or grey-box scope raises reviewer questions about completeness.

03

Test-Case Coverage Traceability

§524B premarket submissions require documented test cases mapped to the threat model. Without internal access, several threat categories cannot produce evidence.

04

Cost-of-Defect Curve

Finding a vulnerability in development is ~15× the baseline cost; finding it after release is 30–100×. Open-box shifts discovery left. Open-box enables automated verification of CVEs and hundreds of other findings.

PREMARKET 01

FDA SPDF: Continuous, In-Development Testing

FDA “Cybersecurity in Medical Devices” guidance (Final Sept 2023; Select Updates June 2025) makes the Secure Product Development Framework an expectation — not a recommendation.

Integrated

Penetration testing, fuzzing, and code review must be integrated into standard development processes — not a one-shot pre-submission engagement. Vulnerability volume is too high for one-shot.

Traceable

Security architecture views must establish traceability of architecture and security design elements to test cases executed.

Iterative/Automated

Multiple testing rounds before the final V&V lock, enables fixes and retests while remediation cost is still low. Manufacturers now require automated verification of known CVEs due to volume.

Citation: *FDA “Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions” (Final 2023, Select Updates 2025). Closed-box testing cannot satisfy SPDF's traceability and iteration requirements.*

PREMARKET 02

Manufacturer Testing Scope & §524B RTA Authority

FD&C Act §524B (effective March 29, 2023; RTA enforcement October 1, 2023) requires evidence of cybersecurity testing in premarket submissions. FDA guidance explicitly favors white-box testing for its traceability to internal controls.

> *"White-box penetration testing engagements involve testers receiving full system knowledge including architecture diagrams, source code, credentials, and design documentation."*
>
> — FDA Cybersecurity Guidance, 2023 (Section on Cybersecurity Testing)

§524B REFUSE-TO-ACCEPT (RTA) DEFICIENCY TRIGGERS

Scope	Penetration test scope not mapped by test cases to architecture and security design **(new for 2026)**
Traceability	Findings not traceable to test cases mapped to identified threats in model documentation
White-box evidence	Submission lacks source-code-level verification (exploitability verification) of SAST/SCA/SBOM results
Coverage	Results do not cover firmware, mobile, and cloud components (system-level or intra-system)

PREMARKET 03

Test-Case Coverage by Methodology

FDA review expects test cases mapped to every architectural element, trust boundary, and security requirement. Each “Not tested” row below is a deficiency risk in submission review.

Security Requirement / Threat Path	Closed-Box	Open-Box
Network services	Partial (external only)	Covered
Authentication / privilege boundaries / session logic	Partial	Covered
Firmware / bootloader / update process / secure boot / SBOM	Not tested (not accessible)	Covered
Inter-service/ process / system-level trust boundaries	Partial (some accessible)	Covered
Hard-coded secrets / cryptographic use	Partial (reversing required)	Covered
Source-level exploitability of SAST/SCA findings	Not tested	Covered
Fuzzing of key attack surface	Not tested (requires harness)	Covered

Source: FDA 2023 Final Guidance, SPDF traceability requirements; AAMI TIR57:2016 & Verification Evidence

PREMARKET 04

Cost of Defect by Discovery Phase

Open-box testing catches vulnerabilities during Development (15×) rather than during Late V&V / DVT (30×) or post-launch (100×).

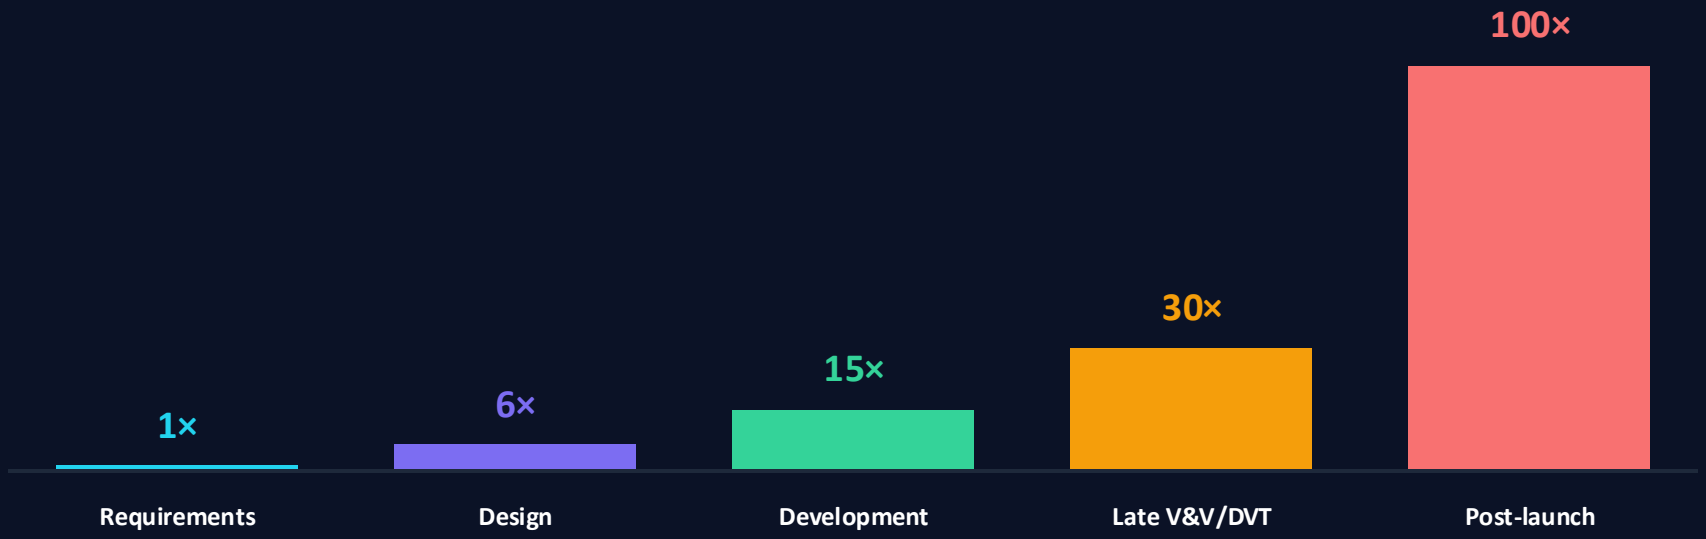

Cost multiplier relative to Requirements-phase baseline. ***A vulnerability caught open-box in development costs ~6.7× less to remediate than caught in late V&V, and ~6.7× to 20× less than caught post-launch.***

REGULATORY MAP

Open-Box Expectation Across Jurisdictions

Each regulation below either mandates or establishes traceability / iterative-testing requirements that are not achievable without manufacturer open-box access.

Regulation / Standard	Citation	Open-Box Implication
FDA "Cybersecurity in Medical Devices"	Final 2023; Updates 2025	SPDF requires iterative in-dev testing with threat-to-test traceability
FD&C Act §524B (PATCH Act)	RTA effective Oct 1 2023	Premarket submissions subject to RTA without white-box evidence
AAMI TIR57:2016 (FDA-recognized)	§ Verification Evidence	Mandates traceability: threats → controls → test evidence
AAMI TIR97:2019 (FDA-recognized)	Postmarket risk mgmt	Postmarket testing must trace to premarket threat model
IEC 81001-5-1:2021 (FDA-recognized)	§ 9.6, § 9.7	Pen testers require threat-model access; scope derived from model
EU MDR 2017/745	Annex I § 17.2; MDCG 2019-16	"State of the art" verification of controls against threat model
EU Cyber Resilience Act (CRA)	Reg 2024/2847, Art. 10, 13	Prove absence of "known exploitable vulnerabilities" at placement
IMDRF N60 / N73	TPLC framework	Threat models verified through testing at each development phase
NIS2 Directive	Art. 21, 23	Testing is a mandatory risk-management measure; 24h/72h reporting

Sources: FDA.gov; AAMI; IEC; EUR-Lex; IMDRF.org; NIS2 Directive official text. NIST SP 800-115 referenced for methodology terminology.

PART 2

Postmarket: Open-Box Value

Postmarket arguments are driven by disclosure timelines imposed externally. Each point below is examined in the following slides.

05

Manufacturer Is the Lowest-Cost Discoverer

Internal open-box testing consumes much less engineering time. External discovery triggers coordinated disclosure, patching urgency, and regulatory reporting — multi-order cost escalation.

06

Clock Control Is Retained Only by Finding First

CRA Art. 14 (24h / 72h), NIS2 Art. 23 (72h), CISA CVD (45–90d), and FDA coordinated disclosure all start externally. Internal discovery is the only path to retaining fix-timeline authority.

07

AI-Assisted Discovery Has Lowered the External Bar & Increased Vulnerability Volume

Autonomous pentesting (XBOW, Big Sleep) compresses weeks of manual RE into hours. External researchers already use these tools. Open-box is required to match them internally and automated the same type of identification/verification at scale.

POSTMARKET 06

Externally-Imposed Disclosure Timelines

Once discovery occurs outside the manufacturer, the applicable clock begins. The manufacturer cannot extend these windows unilaterally.

CRA Art. 14

24h

Early warning to ENISA/CSIRT for actively exploited vulnerabilities

CRA Art. 14

72h

Full incident notification with severity assessment

NIS2 Art. 23

72h

Incident report for essential/important entities

CISA CVD

45–90d

Coordinated vulnerability disclosure public window

FDA §524B: *Coordinated disclosure policy required; remediation must be timely. No fixed hour-clock — but delayed patches after known disclosure create enforcement exposure.*

POSTMARKET 07

AI-Assisted Vulnerability Discovery

Public benchmarks 2024–2026 document step-function reductions in discovery cost for external researchers. The same tools, applied internally, require open-box access to match that depth.

85×

XBOW matched a 40-hour manual assessment in 28 minutes. Manufacturers must now verify at that speed.

20

Critical OSS vulns autonomously found by Google Big Sleep in one eval

6×

AI agent steps-completed growth in 18 months (1.7 → 9.8)

< 1 hr

ELTON-observed AI testing replicates 5–8-yr pentester depth per device

External researchers no longer require years of reverse-engineering to reach device internals. Open-box methodology preserves parity between internal assessment depth and the external threat environment.

Sources: XBOW 2025 public benchmarks; Google Project Zero "Big Sleep" (CVE-2025-6965); ARTEMIS (2025); ELTON internal testing.

OPERATIONAL

Artifacts Required for Open-Box Testing

Open-box methodology defines the minimum set of artifacts provided to the test team. Provisioning is standard across FDA-recognized standards (AAMI TIR57, IEC 81001-5-1) and NIST SP 800-115.

Source & Build

- Firmware, mobile, and cloud service source
- Build environment and dependencies
- Cryptographic material usage (non-production)

Architecture & Documentation

- Threat model and trust boundary diagrams
- Data-flow and interface specifications
- SBOM and third-party component inventory

Privileged Test Access

- Admin and service credentials on test rigs
- Debug interfaces (JTAG, serial, console)
- Non-production firmware signing keys only

Out of scope: *production signing keys, patient data, fielded device credentials. Open-box is a test-environment posture; it does not loosen production access controls.*

COMPARISON

Closed-Box vs. Open-Box Methodology

Dimension-by-dimension comparison.

Dimension	Closed-Box	Open-Box
Cost to remediate findings	30–100×	~15×
§524B RTA risk	High	Low
FDA review questions on scope	Expected	Rare
Disclosure clock control	CISA / CRA / NIS2	Retained
CVSS severity setting (public)	Researcher/regulator	Manufacturer
Firmware / boundary / crypto coverage	Minimal	Full
SPDF expectation alignment	Inconsistent	Aligned
Peer manufacturer alignment (ELTON-observed)	< 5%	> 95%
Test-case → threat model traceability	Not achievable	Achievable

SUMMARY

ELTON’s Recommendation

1. Regulatory position	FDA SPDF, §524B, AAMI TIR57/97, IEC 81001-5-1, EU MDR, CRA, IMDRF N60, and NIS2 each require testing outcomes, traceability, or evidence that closed-box methodology cannot generate. Open-box is the methodology that satisfies these requirements.
2. Coverage position	Closed-box testing cannot produce evidence for for many FDA-required threat paths. All surface in §524B submission review. Open-box is required for test case coverage and evidence.
3. Economic position	Industry cost-of-defect multipliers show 30–100× remediation cost for post-release discovery versus ~15× for in-development. Open-box testing moves discovery to the lower-cost phase.
4. Threat-environment position	AI-assisted discovery tooling is now standard for external researchers. Internal testing without open-box access cannot match the depth available to attackers with or without AI.

RECOMMENDATION: Adopt open-box as the default methodology for premarket submissions and postmarket assessments.