ELTON

TECHNICAL WHITEPAPER

Rating Medical Device Vulnerabilities on the FDA-Qualified CVSS Rubric

How ELTON applies the MITRE Rubric for Applying CVSS to Medical Devices (MDDT Q171974) to every CVSS v3.1 rating, and carries the same discipline into CVSS v4.0 with MDCVJ

KEY TAKEAWAYS

FDA qualified the MITRE Rubric for Applying CVSS to Medical Devices under the Medical Device Development Tools program as Q171974, for patient-centric impact and urgency in time-sensitive postmarket disclosures supporting CVSS v3.0. A rubric-based rating rests on a method FDA has already reviewed.

The rubric replaces judgment with decision questions. For each CVSS base metric it asks a short sequence of yes or no questions about the device as deployed, and the metric value falls out of the answers. Ten analysts stop producing ten scores.

The questions need facts about the device: which interfaces reach the vulnerable component, which trust boundaries the exploit crosses, which assets sit on the component and beyond it. ELTON reads those facts from the digital twin built from the manufacturer's own submission documents.

Every rating ELTON produces carries its trail: the metric value, the rubric questions asked, the answer to each, the twin element that supplied it, and the intelligence that backs it, down to the file and line in the finding.

Preconditions and postconditions tie the finding to the twin. What the exploit needs before it runs is checked against the path in. What it leaves behind drives the path out, which is where the Subsequent metrics of CVSS v4.0 come from.

CVSS v4.0 asks questions the v3 rubric never did. ELTON MDCVJ extends the qualified rubric's logic to v4.0 and is under review in the MDDT program. Every finding carries both scores through the transition.

Executive summary

A CVSS score is only as defensible as the reasoning behind it. The National Vulnerability Database publishes a generic score for a CVE that assumes the worst case: network reachability, no controls, the most sensitive deployment. A medical device is rarely that case. The same CVE on a service port behind a technician login, or on a component that never touches patient data, deserves a different number, and FDA reviewers know it.

The problem has never been the arithmetic. The problem is that the inputs to the arithmetic were opinions. Hand the same finding to ten experienced analysts and you get ten vectors, each defended in the room and none defensible on paper. A manufacturer with hundreds of findings across SBOM scans, code analysis and penetration tests cannot afford a method that changes with the analyst.

MITRE wrote the Rubric for Applying CVSS to Medical Devices to fix that. It turns each CVSS base metric into a sequence of concrete questions about the device as deployed, with clinical impact in view. FDA evaluated the rubric under the Medical Device Development Tools (MDDT) program and qualified it as Q171974. A rating produced by following the rubric rests on a method FDA has already reviewed for exactly this job.

This paper covers what FDA qualified and what it did not, how the rubric's questions work, where the answers come from in ELTON, what the rating trail looks like on one real finding, and how ELTON carries the same discipline into CVSS v4.0 with MDCVJ, the Medical Device Contextual Vulnerability Justification method now under MDDT review.

This paper is written for:

- Product security and regulatory affairs leads who rate vulnerabilities for premarket submissions and postmarket reports
- Reviewers and auditors who need to follow a rating back to the evidence that produced it
- Engineering teams migrating a portfolio from CVSS v3.1 to v4.0 before the July 2027 deadline

1. What FDA qualified, and what it did not

The MDDT program exists so industry does not have to re-argue methods inside every submission. A sponsor proposes a tool, FDA evaluates it against a specific context of use, and qualification means the data that tool produces can be relied on within that scope.

MITRE submitted the Rubric for Applying CVSS to Medical Devices to the program in November 2017 under submission number Q171974, and FDA qualified version 0.12.04, dated September 3, 2019. The authors are Penny Chase and Steve Christey Coley of MITRE. The work was done under contract with FDA because CVSS and its guidance were written for enterprise IT and do not reflect the clinical environment or patient safety. The rubric was written against the FIRST CVSS v3.0 base metrics. CVSS v3.1 kept the same base metrics and values, and ELTON applies the rubric to produce v3.1 vectors.

The qualified context of use is specific, and worth quoting. FDA's summary of evidence states that the rubric "is qualified for the evaluation and justification of patient-centric, situational impact and urgency characteristics in time-sensitive postmarket vulnerability disclosures of medical devices, when supporting the FIRST CVSS V3.0 standard," and that the vector string should always be published together with the score.

The same summary is candid about what the evidence did not yet support. FDA wrote that the rubric "is not suitable for estimating the impact and urgency of a 'chained' vulnerability attack," where a series of individual vulnerabilities is used to lower a security architecture's defenses until a hidden route appears. It also declined to qualify the premarket justification use it had started to see in submissions, saying that an architecture-based approach to the cyber resilience of a design would be needed before premarket reliance made sense.

One thing we will not do is inflate any of that. Qualification covers the rating methodology, for postmarket disclosures, on CVSS v3.0. FDA has not endorsed the ELTON platform, and no vendor can honestly claim that kind of blanket blessing. What a manufacturer gets is narrower and more useful. Every v3.1 score ELTON produces follows a method FDA has already reviewed. Every v4.0 score follows an architecture-based extension of that method, MDCVJ, which FDA is reviewing now. And the two gaps FDA named, chained attacks and premarket architecture, are the two things the dependency graph and the digital twin were built to address.

2. Why a rubric beats an expert

An experienced analyst carries the right instincts and the wrong record. The instinct says a buffer overflow in a DICOM listener behind a segmented VLAN is not a 9.8. The record says "Attack Vector: Adjacent" with no note of why, and a reviewer six months later cannot tell whether the analyst knew about the VLAN or guessed.

A rubric fixes three things at once.

Rules. The rubric asks concrete decision questions about vector, access, interaction and clinical impact. The score falls out of the answers, not out of seniority.

Traceability. Every answer is recorded. When a reviewer asks why the score is a 5.9, the response is the decision trail rather than a shrug and a resume.

Repeatability. The same finding rated next quarter, or by a different analyst, lands on the same number. Consistency is what makes a rating program auditable.

3. How the rubric works: questions, not judgment

For each CVSS base metric the rubric lays out a short decision tree. Each node is a yes or no question about the device and its deployment, and each path through the tree ends at a metric value. The analyst records every answer. An unknown answer defaults to the worst case. The answers are kept as an extended vector whose codes begin with X, published next to the CVSS vector, so anyone reading the score can see the path that produced it. Where an answer suggests a potential impact to patient safety, the rubric marks the point PIPS and asks for a hazards analysis.

Attack Vector is the clearest example, and the metric a reviewer challenges most often. The rubric's Attack Vector questions, as ELTON records them on every finding, are:

XAVN. Can the attacker utilize some type of network or communication protocol to exploit this vulnerability?

XAVT. Does the network use OSI layer 3 or 4 protocols, e.g. IP, TCP/IP, or UDP?

XAVW. Is the communication over a wireless channel?

XAVR. Is the range approximately 10 feet or less?

XAVP. Must the attacker have physical contact with the device?

XAVPA. Is the device accessed through a human-user interface, i.e. a user interface intended for manual operation by device users?

A vulnerability reachable over TCP/IP answers yes to both network questions and the trail stops at Network. A Bluetooth Low Energy flaw answers yes to the wireless question and the range question, and the trail ends at Local, the rubric's value for short-range radio where proximity implies a degree of authorization. A Wi-Fi flaw with building-wide range lands at Adjacent. A debug header behind the enclosure answers yes to physical contact and lands at Physical, with XAVPA recording whether the interface is an intended human interface or one the attacker had to break into. The questions the trail never reaches are recorded too, so a reviewer can see that wireless, range and physical contact were considered and ruled out rather than skipped.

The other exploitability metrics and Scope each take one to three questions:

Attack Complexity, XACL. Can the attacker attempt to exploit the vulnerability at will, without requiring any special circumstances, configurations, or use of other vulnerabilities or attacks before attacking this vulnerability? Yes is Low, no is High. The rubric's examples of High include shared secrets the attacker must obtain, rare non-default configurations, a man-in-the-middle position, and a race window.

Privileges Required, XPRL, XPRZ, XPRS. Does the device or component use an authorization model that supports login for multiple users or roles with different privilege levels? Before attempting to exploit the vulnerability, must the attacker be authorized to the affected component? Must the attacker have administrator, maintainer, or other system-level privileges? Yes to all three is High, yes to XPRL and XPRZ with no to XPRS is Low, anything else is None. A legacy device with a single implicit user scores None.

User Interaction, XUI. To successfully exploit the vulnerability, must the attacker depend on another user or victim to perform an action or otherwise interact with the system? The victim has to be someone other than the attacker.

Scope, XS. Can the attacker affect a component whose authority, its authorization scope, is different than that of the vulnerable component? The rubric's examples include other systems that depend on the vulnerable component for essential performance and other devices the vulnerable device connects to.

The impact metrics are where the rubric does its clinical work. For Confidentiality it asks the analyst to walk six categories of data the attacker could read and rate each one High, Low or None: PHI or PII (XCP, with XCPM asking whether 500 or more patients are affected), data related to diagnosis or monitoring (XCD), data that affects delivery of therapy (XCT), data about clinical workflow (XCW), private system or system-user data such as passwords and private keys (XCS), and any other critical data (XCO). If any category is High or Unknown, Confidentiality is High (XCH). If any is Low, it is Low (XCL). Otherwise it is None. Integrity (XIP through XIO) and Availability (XAP through XAO) walk the same six categories for data the attacker could modify and functions the attacker could deny. The rubric also carries the Environmental group, where a hospital's compensating controls such as segmentation or locked rooms adjust the modified metrics without touching the base score.

Every one of those questions is about the device rather than the CVE: which interface, which authorization model, which component, which data. Generic scores skip them, so the rubric needs a source of truth about the product.

4. Where the answers come from: the digital twin

A rubric is only as good as the facts fed into it. In ELTON those facts come from the digital twin, a property graph of the device's security architecture built from documentation the manufacturer already produces for FDA review: cybersecurity architecture views and data flow diagrams, the SBOM, the threat model, the cybersecurity risk assessment and risk management report, the security controls documentation, and the cybersecurity testing reports. No new paperwork, no questionnaires.

The twin holds the elements the rubric questions ask about:

Components, each a node: operating systems, applications, services, processes and software dependencies, bound to a real location in the architecture.

Interfaces, typed Network, Adjacent, Local or Physical, with a configurable flag for interfaces that are optional or disabled in some deployment models.

Data flows, directed edges between components across interfaces, carrying the assets that move on them.

Assets on each component, with confidentiality, integrity and availability requirements expressed as High, Low or None.

Trust levels and trust zones, which define the security hierarchy the exploit has to climb.

Countermeasures recorded on the components and interfaces they protect.

Initial access points from the threat model: the hospital network, Wi-Fi, Bluetooth, USB, the touchscreen, a service port, a debug header.

Every element carries a tag naming the document it came from, and every manual edit is logged. When a rubric question asks which interface reaches the vulnerable service, the answer is a named edge in the graph with a citation, not a recollection.

5. One finding, the trail a reviewer reads

Here is how the pieces fit on a finding rated in ELTON. The vulnerable component is the authentication handler of a FastAPI service. The handler logs the submitted API or Bearer token in cleartext on both the normal and the failure path. The token crosses a trust boundary into the log tier, and any principal who can read those logs holds a live credential for the data service.

ELTON produced CVSS:3.1/AV:N/AC:L/PR:L/UI:N/S:C/C:H/I:H/A:L, a 9.9, and CVSS:4.0/AV:N/AC:L/AT:N/PR:L/UI:N/VC:H/VI:H/VA:L/SC:H/SI:H/SA:L, a 9.4. Both are Critical. Context can lower a score, and it can also confirm one. Here the twin confirmed the worst case, and the record shows why.

The assumptions on the record, which are the preconditions the exploit needs:

The attacker has read access to the central log sink where the credentials land in cleartext.

The logged tokens are live credentials the data service accepts on replay.

The network endpoints on the service carry authentication requests with the token in the password field.

How each v3.1 metric was answered:

Attack Vector, Network. The vulnerability is exploitable through the network-facing HTTPS endpoints on the component, the Archive API endpoint and the Meta API endpoint, both over TCP/IP.

Attack Complexity, Low. A request carrying a token in the password field to the Archive API endpoint is logged unredacted, with no special conditions or additional steps.

Privileges Required, Low. The attacker needs a valid API or Bearer token submitted to the authentication handler, which enforces authorization before granting access.

User Interaction, None. The crafted request alone triggers the logging. No operator step is on the path.

Scope, Changed. The token moves from the FastAPI component into the log tier, a separate component accessible to a wider set of principals.

Confidentiality, High. On the vulnerable system the modelled assets are the OIDC identity and access tokens and the case archive. On the subsequent system they are patient PHI, case metadata and the bone models and STL files. All are modelled High for confidentiality.

Integrity, High. The case archive on the vulnerable system, and the case metadata and models on the subsequent system, are modelled High for integrity.

Availability, Low. The same assets are modelled Low for availability on both systems.

The Attack Vector trail in full. XAVN, yes: the request carrying the token travels over the Archive API, Meta API, Pre-op PDF API or Save API HTTPS endpoints, all Network category interfaces on the component that hosts the vulnerable handler. XAVT, yes: the vulnerable interfaces are HTTPS endpoints over TCP/IP. The evidence attached is the finding description itself, which reads in part: "The handler takes the credential arriving in credentials.password, aliased in the code as password_as_token, and treats it as either a permanent API token or a Bearer token at app.py:144 to 150." XAVW, XAVR, XAVP and XAVPA were asked and answered no, with the reason for each recorded: wired network interfaces, not short range, no physical contact, not a human-user interface.

A reviewer reads that record and sees the value, the questions, the answers, the outlined twin elements that supplied them, and the line of code behind the finding. Nothing in it depends on who ran the rating.

6. Preconditions and postconditions

The link between a finding and the twin is a pair of conditions.

Preconditions are what the exploit needs before it runs: a reachable interface, a privilege level, a configuration state, sometimes an earlier finding that opens the door. ELTON checks each one against the corridor from the initial access point to the vulnerable component. If the corridor does not supply a precondition, the scenario is not exploitable from that entry point, and the rubric questions on the exploitability side answer accordingly.

Postconditions are what the attacker holds afterward: a credential, a writable asset, a shell on a component, a path that was closed before. ELTON carries them into the traversal out of the component. In v3.1 that traversal answers Scope and informs the impact metrics. In v4.0 it answers the Subsequent System metrics directly. It also feeds the dependency graph, where one finding's postconditions are another finding's preconditions. FDA's summary says the rubric alone is not suitable for chained attacks, and ELTON does not ask it to be: the rubric rates one vulnerability on one component, and the graph carries the chain.

> A finding is a vulnerability only while every precondition it needs is produced upstream. Fix the root and the findings beneath it become weaknesses, and their scores say so.

7. The same method, carried into CVSS v4.0

FDA has recognized CVSS v4.0 as a consensus standard, and FIRST CVSS v4.0 becomes the required severity standard for submissions by July 2027. The new specification asks questions the v3 rubric never did, and they are harder to answer by hand.

Attack Requirements (AT) asks whether exploitation depends on deployment conditions an attacker cannot simply create. Answering it honestly needs a model of the deployed device: segmentation, required pairing, physical access assumptions, configurable interfaces that ship disabled.

Subsequent system impact replaces Scope. v4.0 scores confidentiality, integrity and availability on the vulnerable system and on the systems beyond it separately, as VC, VI and VA against SC, SI and SA. Determining subsequent impact needs a graph that knows how the privileges and assets compromised by one vulnerability can be reused to reach others.

Threat metrics replace Temporal. Exploit maturity now moves scores meaningfully and changes over time, so ratings need revisiting rather than archiving.

A v3.1 vector cannot be mechanically converted to v4.0, because the new metrics need information a score alone does not carry. ELTON MDCVJ, Medical Device Contextual Vulnerability Justification, extends the qualified rubric's logic to the v4.0 specification and reads the new answers from the twin. It proceeds in six stages:

1. **Ingest** the submission evidence: architecture views, SBOM, threat model, risk management report, controls documentation, testing reports.
2. **Build** the property graph. Components become nodes carrying trust level, assets with CIA sensitivity, countermeasures and configurability. Data flows become typed, directed edges. Every element is tagged to its source document.
3. **Bind** the vulnerability to a component, with its SSVC context and an on-component exposure indicator, so the unit of analysis is one vulnerability on one component in one device design.
4. **Trace** attack paths from every initial access point in the threat model: the pre-compromise corridor into the component and the post-compromise spread out of it. Traversal respects trust boundaries and directionality, so infeasible paths never inflate a score.
5. **Apply** deterministic CVSS 4.0 rules that extend the FDA-qualified rubric. Exploitability metrics read the entry interface, trust levels traversed, countermeasures, configurability and user interaction. Impact metrics read asset sensitivity on the vulnerable component and everything downstream. If no feasible path exists, the scenario is marked non-exploitable.
6. **Emit** a complete vector, score and metric-level justification for every entry point, citing the components, interfaces, trust levels and assets that drove each metric, back to the submission documents. The canonical score is the highest-severity feasible scenario; the others are retained as context.

Because each vector is computed deterministically from the graph, a score can only change when the underlying architectural inputs change. Any change to a vector therefore requires a traceable change to the change-controlled submission documentation the graph was built from. MDCVJ was submitted to CDRH in April 2026 and is under review in the MDDT program. Through the transition, every finding in ELTON carries both the v3.1 score and the v4.0 score side by side.

8. What this means in a submission

FDA's premarket cybersecurity guidance asks manufacturers to document the security architecture of the device in terms of its interfaces, interconnections and interactions, recommends that this take the form of architecture views, and asks for a plan to assess cybersecurity vulnerabilities. The twin is built from those same views. A rating that cites them is a rating a reviewer can check against the file already in front of them.

The practical result is a vulnerability record that holds up in three settings. In a postmarket disclosure, which is the rubric's qualified context of use, where the vector and the trail travel with the score and a re-rated CVE carries the reason it moved. In a premarket submission, where FDA asked for an architecture-based method before relying on ratings, which is what the twin supplies and what MDCVJ proposes. And in an audit, where the question is never "who decided this" because the record answers "what decided this".

Our position is simple. The number is the least interesting part of a rating. The questions, the answers, and the evidence that produced them are what a regulator reads, and they are what ELTON records on every finding, on both the v3.1 rubric FDA qualified and the v4.0 extension FDA is reviewing now.

MITRE Corporation, Rubric for Applying CVSS to Medical Devices, document PR 18-2208: https://www.mitre.org/sites/default/files/2021-11/pr-18-2208-rubric-for-applying-cvss-to-medical-devices.pdf

U.S. Food and Drug Administration, MDDT Summary of Evidence and Basis of Qualification, Rubric for Applying CVSS to Medical Devices, version 0.12.04, submission Q171974: https://www.fda.gov/media/143131/download

MITRE Corporation, Rubric for Applying CVSS to Medical Devices, overview page: https://www.mitre.org/md-cvss-rubric

FIRST, Common Vulnerability Scoring System v3.1 and v4.0 specifications.

U.S. Food and Drug Administration, Cybersecurity in Medical Devices: Quality Management System Considerations and Content of Premarket Submissions, February 3, 2026.

ELTON Cyber, ELTON MDCVJ MDDT Proposal v1.0: https://www.eltoncyber.com/assets/whitepapers/ELTON-MDCVJ-MDDT-Proposal-v1.0.pdf

ELTON Cyber, FDA MDDT rating methodology: https://www.eltoncyber.com/why-fda-mddt and CVSSv4 migration: https://www.eltoncyber.com/cvssv4-migration

About ELTON. ELTON is the exploitability management platform for medical devices, delivered as a managed program. ELTON builds a digital twin of each device, runs the pipeline across firmware, protocols, and real hardware every release, verifies which findings are exploitable on the device, and evidences every finding and every dismissal for regulatory review. ELTON meets FDA premarket (Section 524B) and postmarket cybersecurity vulnerability testing and management requirements. Talk to us at **eltoncyber.com**.