BLUE GOAT CYBER

Medical Device Cybersecurity · FDA Specialists

THREAT MODELING STARTER KIT

Medical Device Threat Modeling Starter Kit

STRIDE-per-element + ANSI/AAMI SW96 methodology, a device-specific checklist, attack-tree template, and the traceability matrix reviewers want.

bluegoatcyber.com

© Blue Goat Cyber. For educational use. Not legal or regulatory advice.

WHY THIS MATTERS

Medical device threat modeling starter kit.

The FDA's 2026 final premarket guidance and ANSI/AAMI SW96:2023 both require a structured, documented threat model. Reviewers reject generic 'STRIDE applied to a SaMD' write-ups. This kit gives you a repeatable methodology - decompose the system, enumerate threats, rate them, mitigate, trace - that you can run on your own device today and present to a reviewer with confidence.

Key takeaway: A threat model is not a one-time deliverable. It is a living artifact that gets updated every time the architecture changes - new interface, new dependency, new deployment target. Build it for maintenance, not for a single submission.

1. The five-step methodology

Step 1 · Scope	Define the system boundary, in-scope components, assumptions, and out-of-scope items. Identify primary and supporting assets (PHI, device function, calibration data, audit logs).
Step 2 · Decompose	Build a Data Flow Diagram (DFD) showing processes, data stores, external entities, and trust boundaries. One DFD per major use case if needed.
Step 3 · Enumerate	Apply STRIDE-per-element. For each element/flow, walk the six categories and document credible threats. Add attack trees for high-value flows.
Step 4 · Rate	Score with CVSS v3.1 + clinical impact (patient safety, data integrity, availability). Map to ISO 14971 risk levels.
Step 5 · Mitigate & trace	For each accepted threat, define a mitigation, trace it to a security requirement, design control, and V&V test case.

2. STRIDE quick reference

S - Spoofing	Attacker pretends to be a legitimate user, device, or service. Mitigations: strong auth, mutual TLS, device certificates.
T - Tampering	Unauthorized modification of data, firmware, or configuration in transit or at rest. Mitigations: signed firmware, integrity checks, secure boot, write-protected partitions.

R - Repudiation	User/system denies an action. Mitigations: tamper-evident audit logs, signed log entries, NTP-synced timestamps.
I - Information disclosure	Confidential data exposed (PHI, credentials, IP). Mitigations: encryption at rest + in transit, least-privilege, memory scrubbing, anti-debug.
D - Denial of service	Device or critical function rendered unavailable. Mitigations: rate limiting, watchdogs, fail-safe defaults, resource quotas.
E - Elevation of privilege	Attacker gains higher-level access than authorized. Mitigations: privilege separation, RBAC, secure update process, sandbox.

3. The medical-device-specific threat checklist

Walk through every interface and data flow in your device. For each, ask:

- Can the user be impersonated? (companion app, clinician portal, OEM service tool)
- Can the device be impersonated to its peers (cloud, gateway, paired devices)?
- Can a man-in-the-middle modify therapy commands, sensor readings, or firmware updates?
- Can wireless protocols (BLE, Wi-Fi, cellular, NFC, Zigbee) be jammed, replayed, or downgraded?
- Is the JTAG/SWD/UART debug interface disabled in production builds?
- Can USB enumeration cause unsafe behavior (BadUSB, HID injection, mass-storage)?
- Can a malicious mobile app on a paired phone send commands the user did not authorize?
- Can cloud API tokens be exfiltrated from the mobile app or device?
- Can the firmware update mechanism be abused to install unsigned or downgraded firmware?
- Are cryptographic keys protected (HSM, secure element, OS keystore)?
- Can audit logs be erased, suppressed, or forged?
- Are there denial-of-service paths to safety-critical functions (battery drain, watchdog reset loops)?
- For AI/ML-enabled SaMD: data poisoning, model inversion, adversarial inputs, prompt injection (LLM features)

4. Attack tree template (high-value flow)

For each high-value data flow (e.g., 'deliver therapy command'), build an attack tree with the goal at the root and progressively concrete attacks as children. Example skeleton:

Goal: Deliver unauthorized therapy command to device

1. Compromise companion app on patient's phone
  1.1 Exploit OS vulnerability (jailbreak/root)
  1.2 Trojanized side-loaded build
  1.3 Stolen device, no PIN
2. Spoof BLE pairing
  2.1 Replay captured pairing
  2.2 Downgrade to Just-Works
3. Compromise cloud back-end and push command
  3.1 Stolen API key in mobile binary
  3.2 SSRF/IDOR in operator portal

Key takeaway: For each leaf node, document existing mitigations and residual risk. Leaf nodes with no mitigation become threats that must be accepted (with justification) or remediated.

5. Traceability: the artifact reviewers actually want

Every threat in your model should trace through to a verifiable test. Maintain a single traceability matrix with these columns:

- Threat ID
- STRIDE category
- Affected element / data flow
- Pre-mitigation risk (likelihood × clinical impact)
- Mitigation (technical control or compensating control)
- Security requirement ID
- Design control reference
- V&V test case ID
- Post-mitigation residual risk
- Risk acceptance signoff

6. Common mistakes that get flagged in review

- **Generic STRIDE list** with no system-specific elements - reads as boilerplate
- **No DFD** or a DFD that doesn't show trust boundaries or external entities
- **Threats without mitigations** (or worse, mitigations without threats)
- **No clinical impact** in risk scoring - cybersecurity threats must connect to ISO 14971
- **Static document** with no evidence it's been updated since the architecture changed
- **Missing AI/ML threats** for AI-enabled SaMD - the 2026 guidance specifically calls these out
- **Mitigations not traced** to design controls or V&V test cases

READY FOR THE NEXT STEP?

Talk to a senior medical device security engineer.

We've contributed cybersecurity documentation, threat models, SBOMs, and pen test reports to **250+ FDA submissions** across 510(k), De Novo, and PMA pathways - with a **100% submission success rate**. If you're building a cyber device and want to go in with airtight artifacts on the first try, we should talk.

Book a threat modeling working session
30 minutes. No slide deck. Bring your device profile and your top three questions.

bluegoatcyber.com/contact

WHAT WE DO

- **FDA premarket cybersecurity** - SPDF, threat model, SBOM, pen test, security risk assessment, eSTAR-ready package
- **FDA cybersecurity deficiency response** - rapid remediation when reviewers send a letter
- **Postmarket programs** - vulnerability monitoring, VEX, coordinated disclosure, patch validation
- **Legacy device remediation** - bring older devices up to current FDA expectations