How to Conduct a Threat Model

End-to-end STRIDE threat modeling guide

Step 1: Scope and Prepare

Define what you're threat modeling (a feature, a service, an API). Gather architecture diagrams, data flow documentation, and the team.

- Time-box the session: 60-90 minutes for initial model
- Include: developer, architect, security engineer, product owner
- Bring: whiteboard or diagramming tool, architecture docs, data classification

Step 2: Create the Data Flow Diagram

Draw the system showing processes, data stores, data flows, and external entities. Mark trust boundaries where data crosses security domains.

- Processes: circles — your code, services, functions
- Data stores: parallel lines — databases, caches, file systems
- Data flows: arrows — API calls, messages, file transfers
- External entities: rectangles — users, third-party services, AI tools
- Trust boundaries: dashed lines — where authentication/authorization is enforced

Step 3: Apply STRIDE to Each Element

For each element on the DFD, systematically ask the six STRIDE questions. Focus on trust boundary crossings — this is where most threats exist.

- Spoofing: Can someone pretend to be this entity?
- Tampering: Can data be modified in transit or at rest?
- Repudiation: Can actions be performed without accountability?
- Information Disclosure: Can sensitive data leak from this component?
- Denial of Service: Can this component be overwhelmed?
- Elevation of Privilege: Can a user gain unauthorized access?

Step 4: Rate and Mitigate

Score each threat (DREAD or CVSS), then define mitigations. Not every threat needs immediate mitigation — accept low-risk threats with documentation.

- Create a tracking ticket for each threat requiring mitigation
- Link mitigations to specific code changes or architecture decisions
- Review the threat model at each major feature change or quarterly