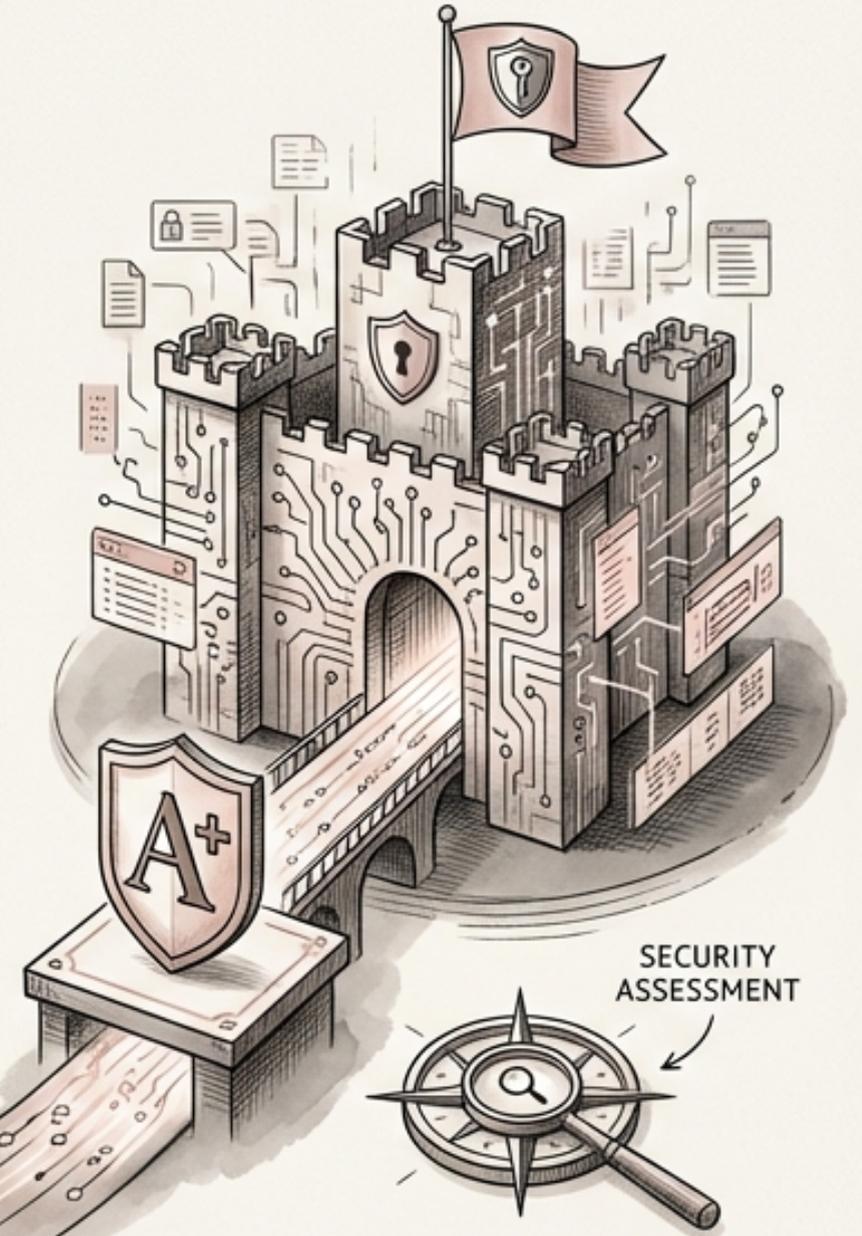


Securing System Design:

Architecture Security Assessment Defined

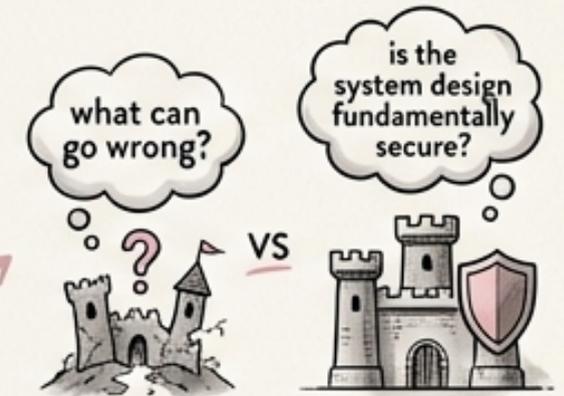


80px

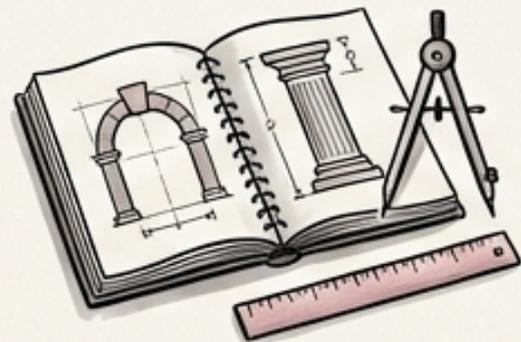
SECURING SYSTEM DESIGN: ARCHITECTURE SECURITY ASSESSMENT DEFINED



- Architecture Security Assessment (ASA) evaluates the fundamental soundness of a system's design before code is written.
- Unlike threat modeling (which asks "what can go wrong?"), ASA asks, "is the system design fundamentally secure?".

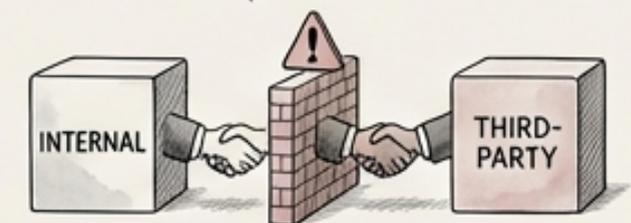
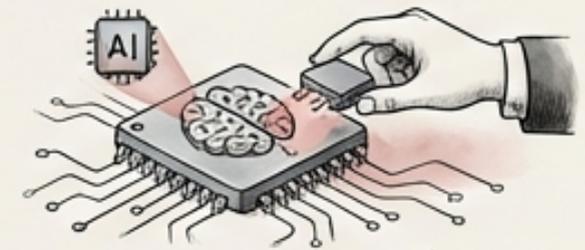


- ASA proactively identifies vulnerabilities and weaknesses early in the development lifecycle, reducing rework and costs.
- For AI-augmented teams, ASA includes evaluating how AI components integrate securely into the overall architecture.
- Focus is on design principles and patterns, ensuring a robust and resilient system from the outset.



When to Condiuct an Architecture Security Assessment

- **New Application/Major Service:** Always perform an ASA for new applications or significant service deployments.
- **Significant Architecture Change:** Reassess the architecture after any major modifications or refactoring.
- **New AI Component Integration:** Evaluate the security implications of adding new AI components to the system.
- **Cloud Migration:** Assess the architectural changes and security controls required for migrating to the cloud.
- **Third-Party Integration:** Evaluate the security risks associated with integrating third-party systems or APIs.



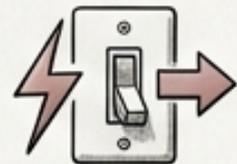
ASA Cadence: Integrating Security into the Development Lifecycle



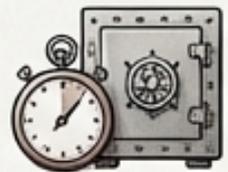
- **Initial Assessment:** Conduct a thorough ASA during the initial design phase of a project.



- **Annual Reassessment:** Perform a comprehensive ASA at least annually to identify emerging threats and vulnerabilities.



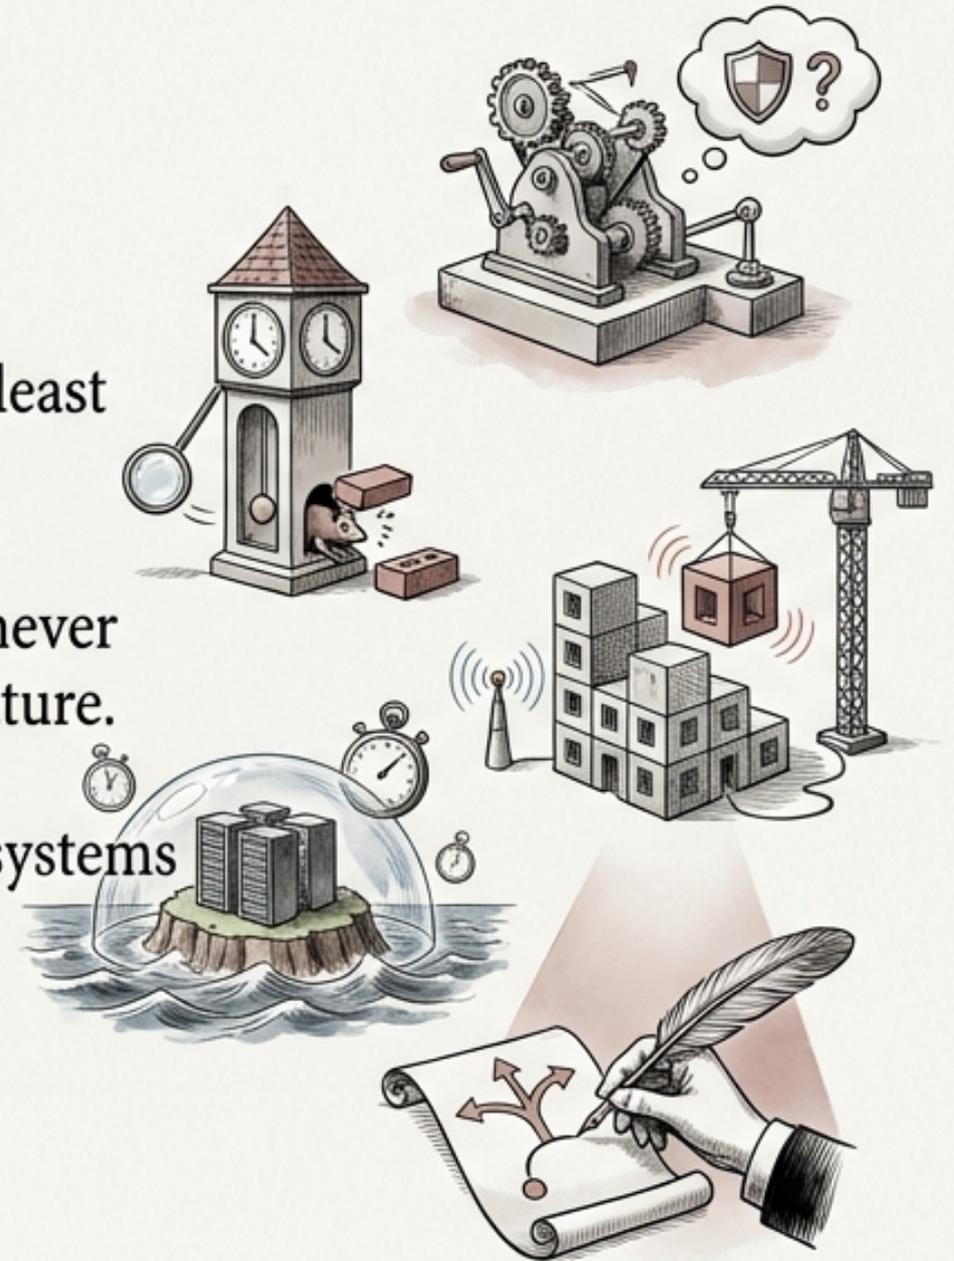
- **Event-Driven Reassessment:** Trigger a reassessment whenever there is a major architectural change or significant new feature.



- Consider more frequent reassessments for highly critical systems or systems with rapidly evolving threat landscapes.



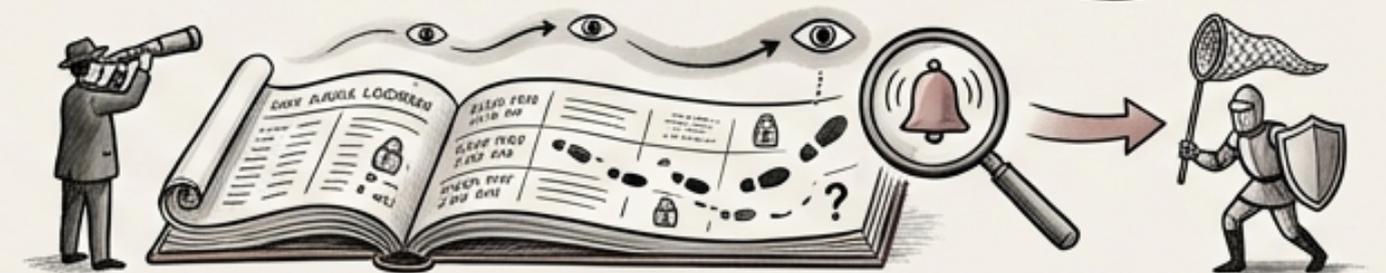
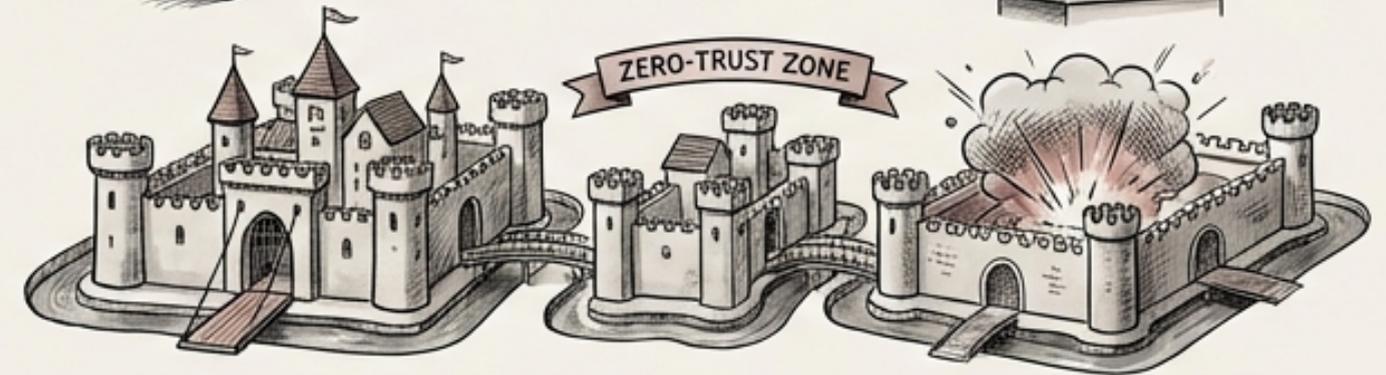
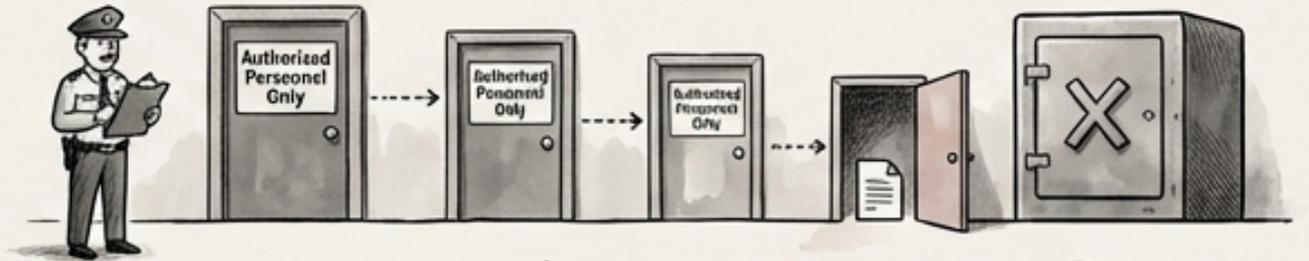
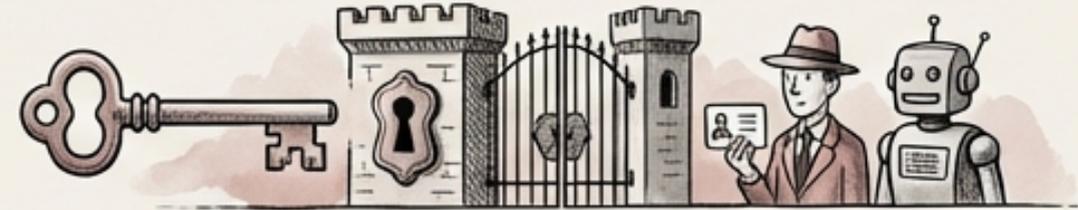
- Document the rationale for your chosen cadence and any deviations from the recommended schedule.



Seven Key Domains of Architecture Security Assessment



-  **Authentication & Identity:** Securely verifying user and system identities
-  **Authorization & Access Control:** Enforcing least privilege and controlling access to resources
-  **Data Protection** (at rest, in transit, in use): Protecting data confidentiality, integrity, and availability across its lifecycle
-  **Network Security:** Segmenting networks and implementing zero-trust principles to limit the blast radius of attacks
-  **Logging & Monitoring:** Capturing and analyzing security events to detect and respond to threats



Maturity Scoring: Measuring Security Posture Within Each Domain

-  Each of the seven security domains is scored on a maturity scale to assess the current security posture.



Maturity levels help identify areas that need improvement and track progress over time.



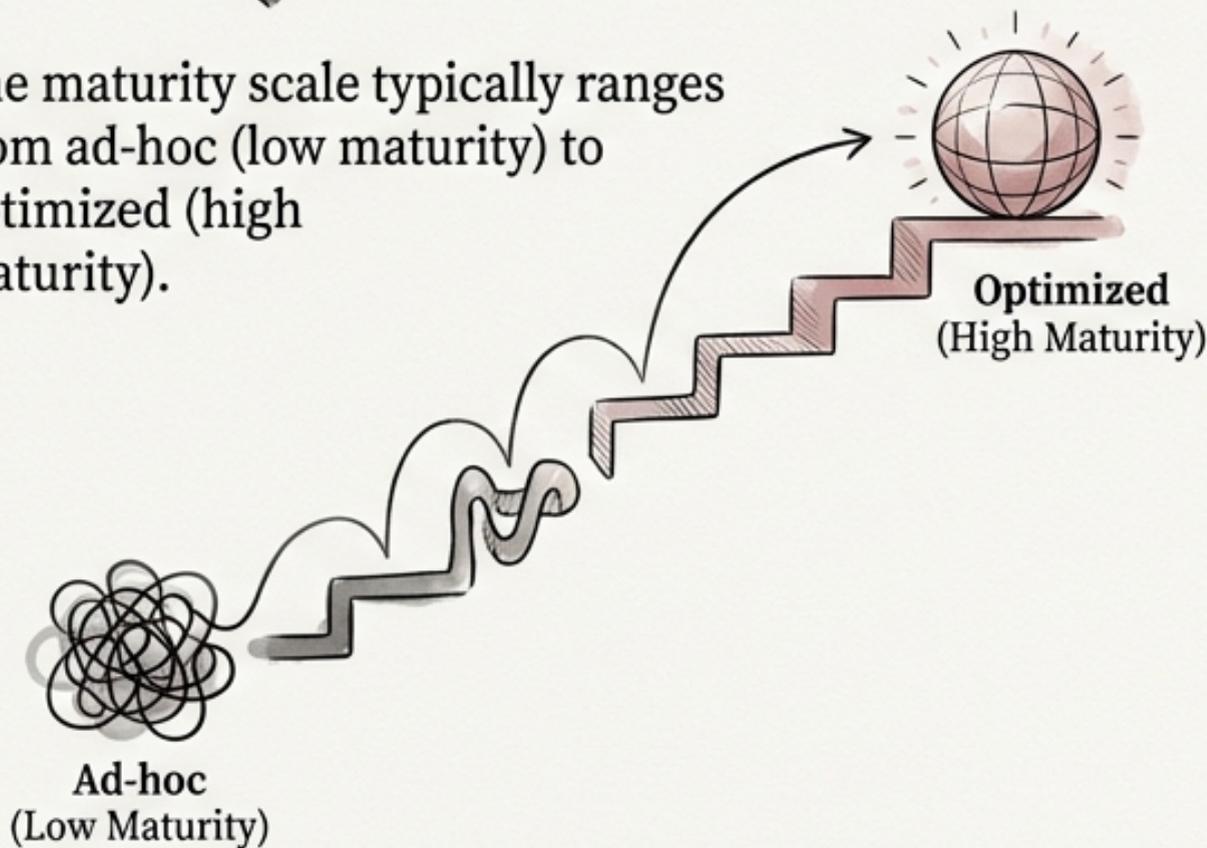
Define clear criteria for each maturity level to ensure consistent and objective assessments.

- Example maturity levels:



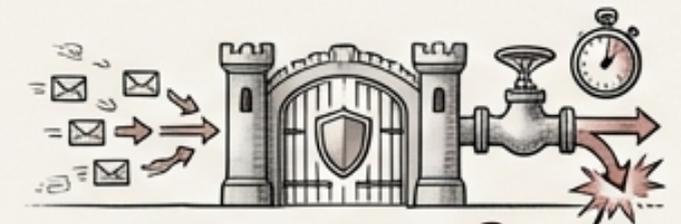
(Customize to org standards.)

- The maturity scale typically ranges from ad-hoc (low maturity) to optimized (high maturity).



Microservice Security Patterns: Protecting Distributed Systems

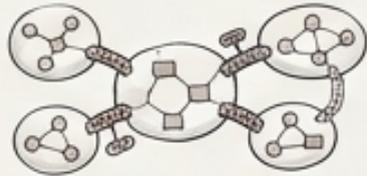
- **Service Mesh with mTLS:** Enforce mutual TLS (mTLS) for secure communication between microservices using a service mesh like Istio or Linkerd.
- **API Gateway with Rate Limiting:** Implement an API gateway to control access to microservices and protect against denial-of-service (DoS) attacks using rate limiting.
- **Distributed Tracing for Security Events:** Use distributed tracing tools like Jaeger or Zipkin to correlate security events across microservices and identify malicious activity.
- **Implement robust input validation and output encoding** within each microservice to prevent common vulnerabilities like SQL injection and cross-site scripting (XSS).
- **Leverage container security best practices**, such as using minimal images and regularly scanning for vulnerabilities.



Zero Trust Architecture: Principles for Secure Access



- **Identity-Based Access:** Verify and authenticate every user and device before granting access to resources.



- **Microsegmentation:** Divide the network into small, isolated segments to limit the blast radius of attacks.



- **Continuous Verification:** Continuously monitor and assess the security posture of users, devices, and applications.



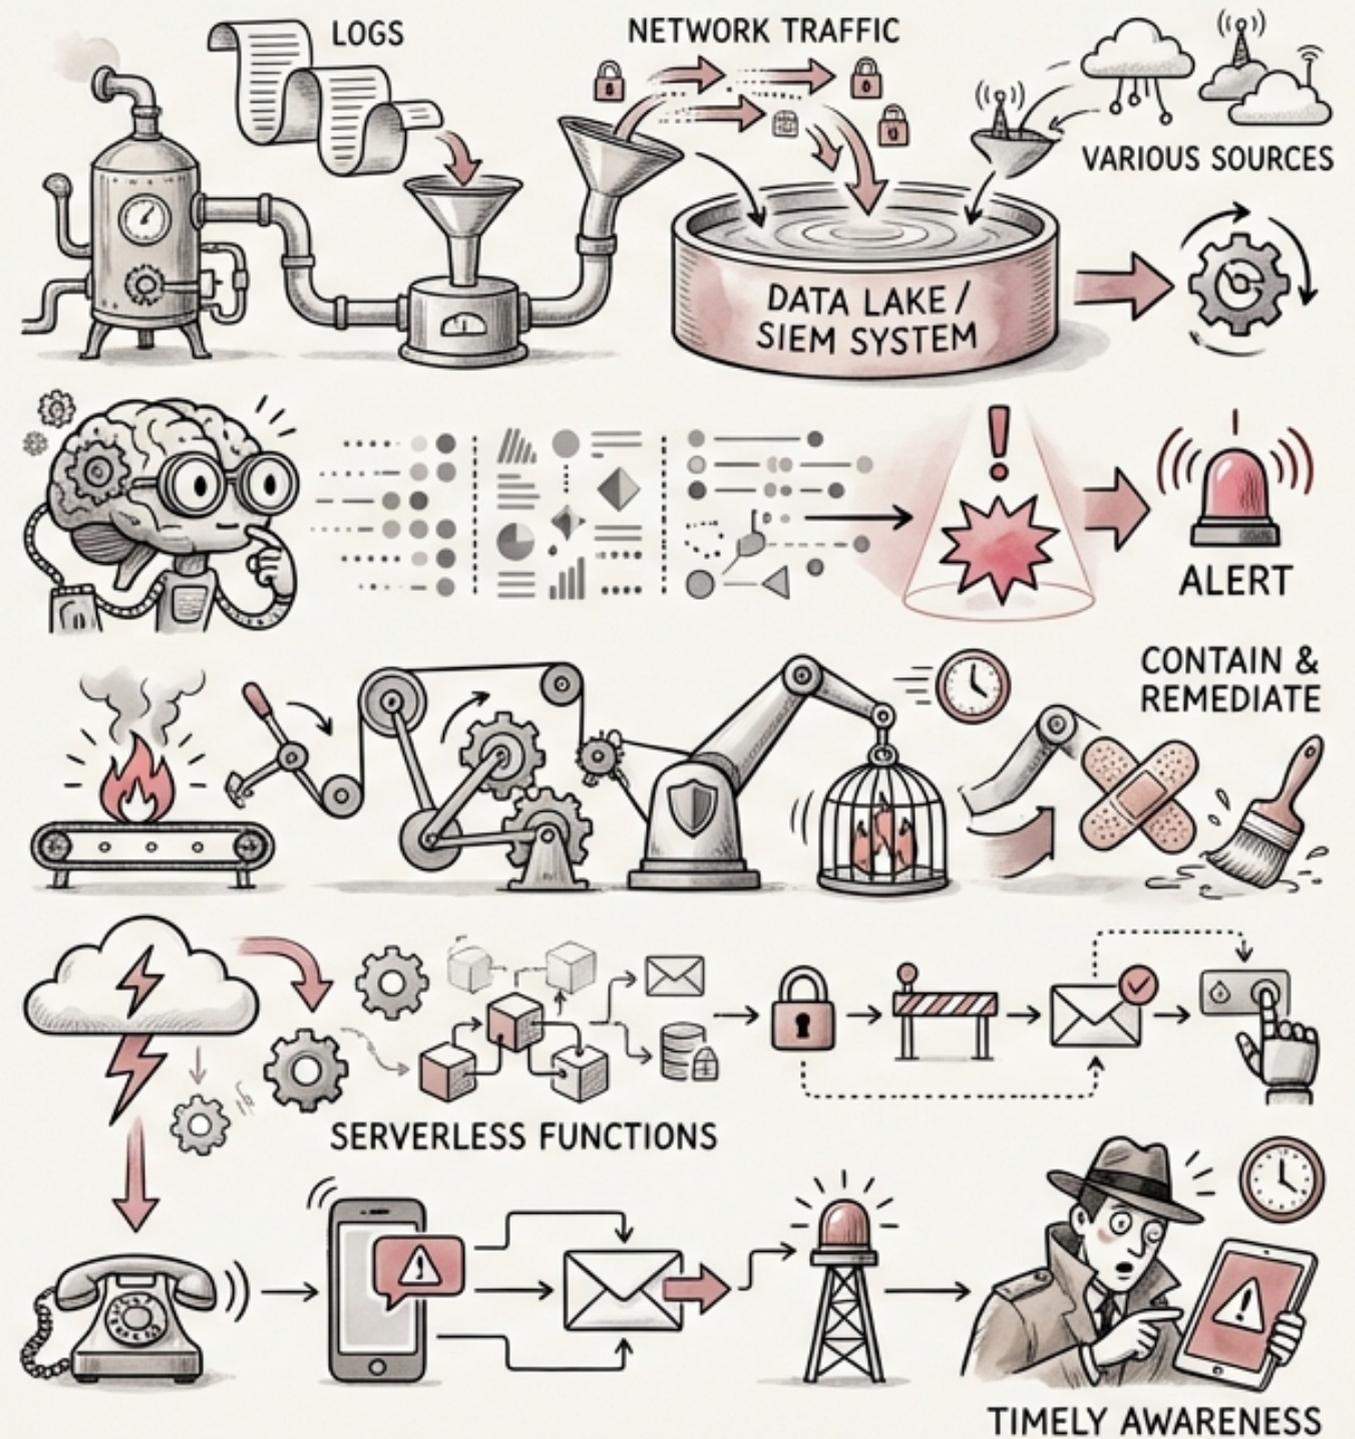
- **Assume Breach:** Operate under the assumption that the network has already been compromised.



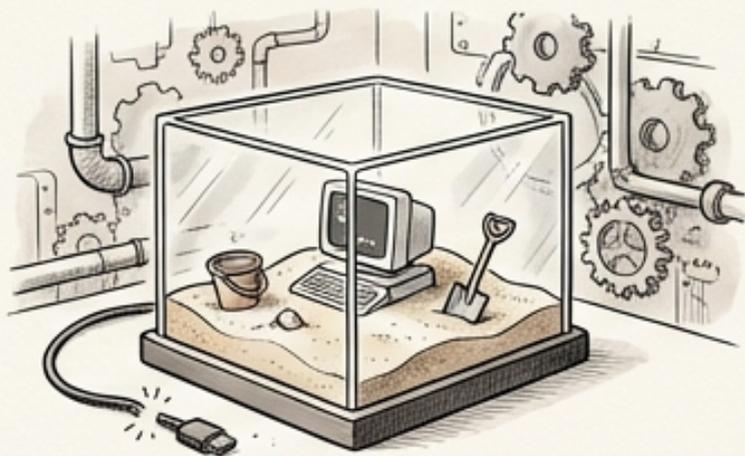
- **Least Privilege Access:** Grant users only the minimum level of access required to perform their job duties.

Event-Driven Security: Responding to Threats in Real-Time

- **SECURITY EVENT STREAMING:** Stream security events from various sources (e.g., logs, network traffic) to a central data lake or security information and event management (SIEM) system.
- **REAL-TIME ANOMALY DETECTION:** Use machine learning algorithms to detect anomalous behavior in real-time and trigger alerts.
- **AUTOMATED INCIDENT RESPONSE:** Automate incident response workflows to quickly contain and remediate security incidents.
- **Leverage serverless functions to process security events and trigger automated actions.**
- **IMPLEMENT ROBUST ALERTING AND NOTIFICATION MECHANISMS:** Ensure timely awareness of security incidents.



AI Integration Patterns: Securing AI-Powered Applications



- **Sandboxed AI Execution:** Run AI models in sandboxed environments to isolate them from the rest of the system and prevent malicious code execution.

- **Output Validation Layers:** Implement validation layers to check the output of AI models for accuracy, consistency, and security.



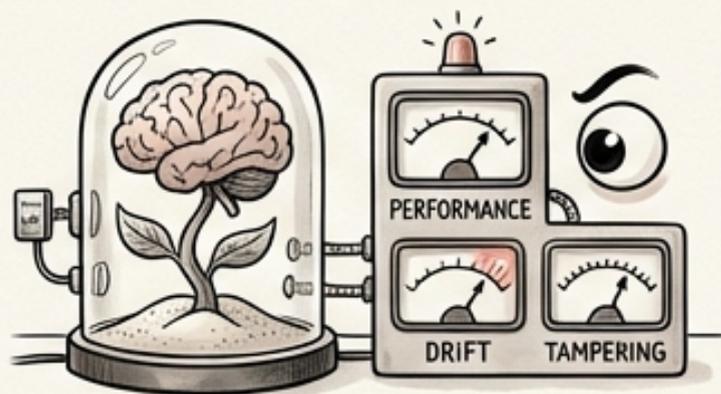
- **Human-in-the-Loop Gates:** Incorporate human review steps into AI-powered workflows to prevent errors and biases from propagating.



- **Input Sanitization:** Thoroughly sanitize and validate input data to prevent adversarial attacks on AI models.

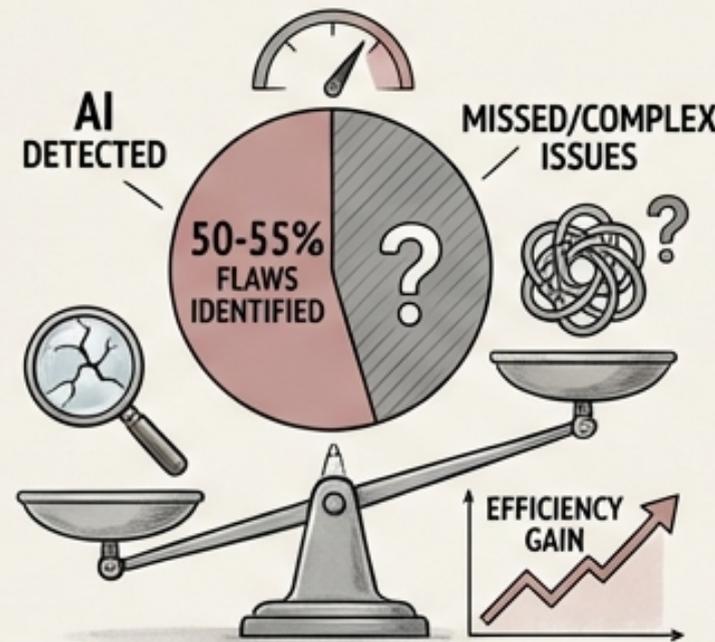
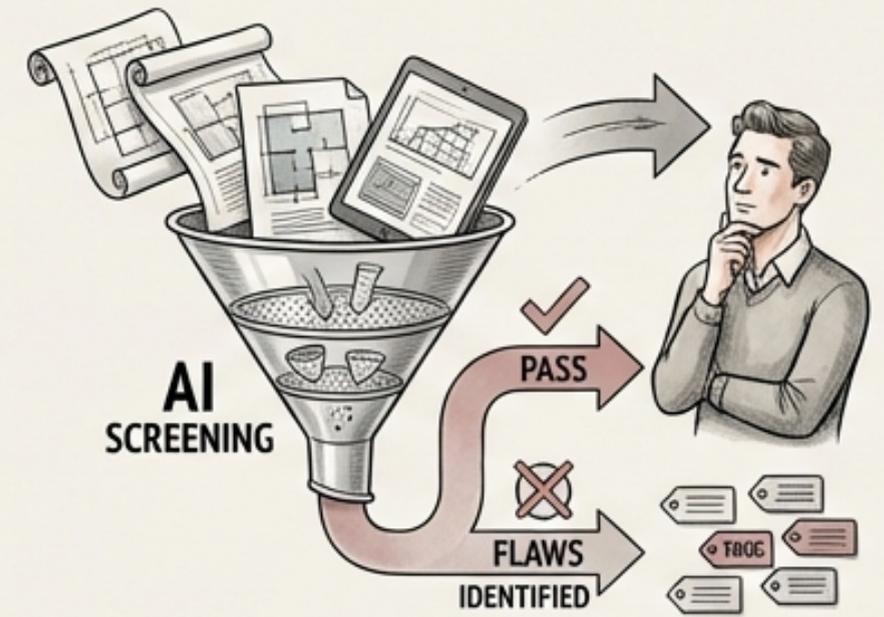
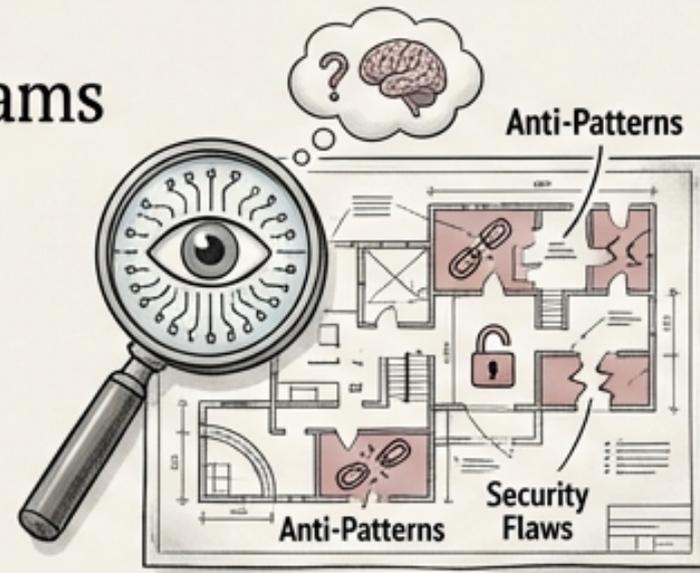


- **Model Monitoring:** Continuously monitor AI models for performance degradation, drift, and signs of tampering.



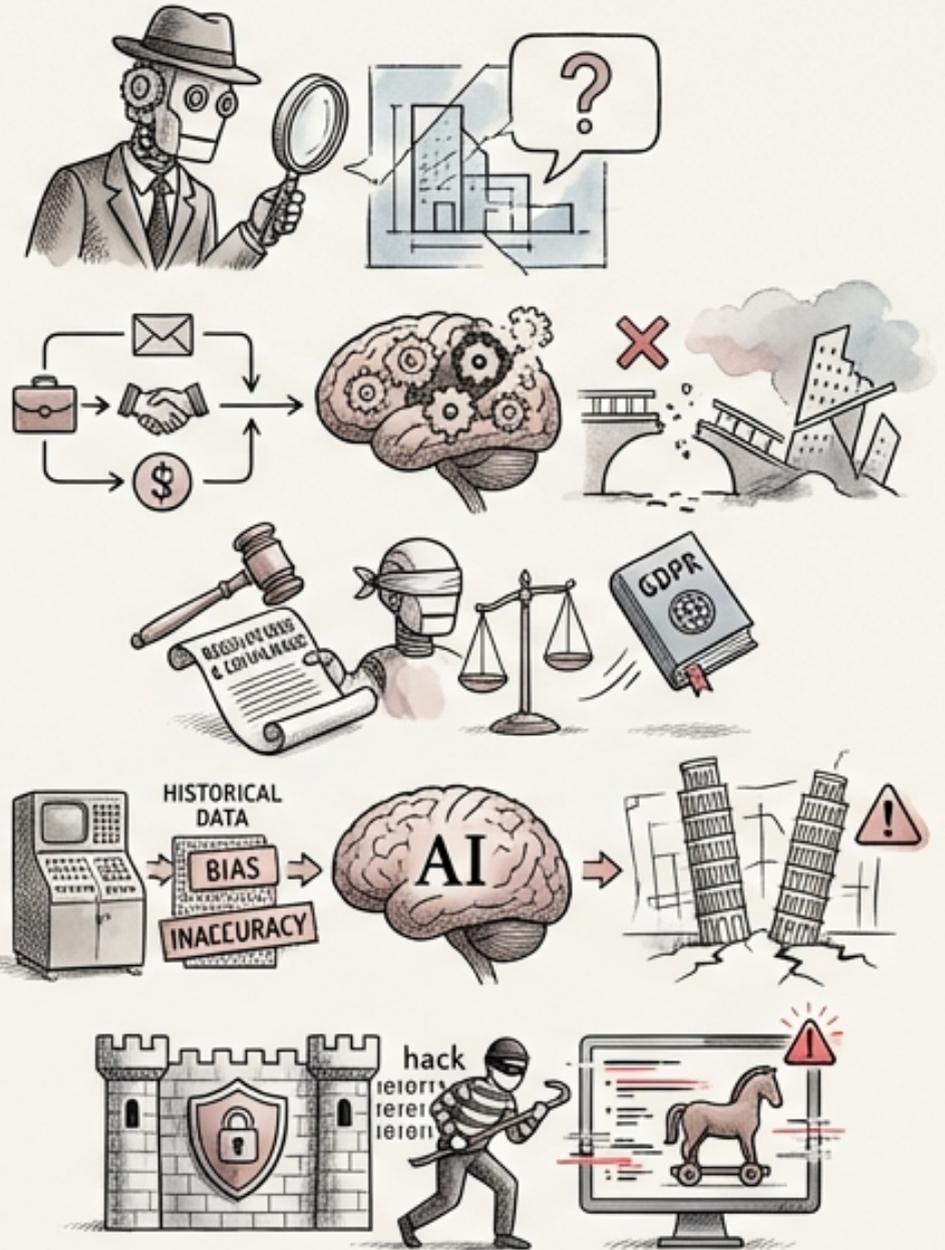
AI-Assisted Architecture Review: Enhancing Efficiency

- AI can analyze architecture diagrams and descriptions to identify common security anti-patterns.
- AI tools can serve as an initial screening pass before human review, increasing efficiency.
- AI-assisted reviews can identify approximately 50-55% of design flaws.
- This detection rate is lower than that of experienced architects, necessitating human oversight.



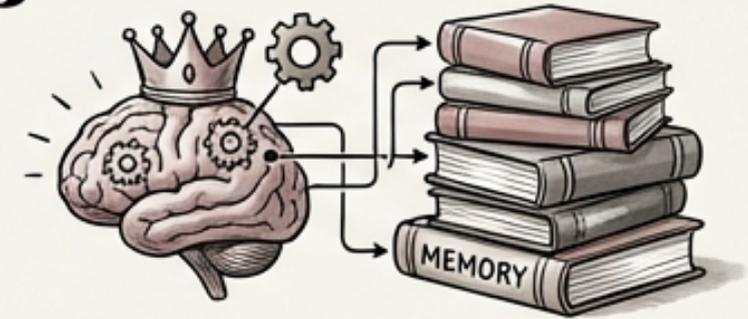
Limitations of AI-Powered Architecture Review

- AI lacks organizational context, limiting its ability to understand specific business requirements.
- AI may not fully grasp the intricacies of business logic, leading to inaccurate or incomplete assessments.
- AI tools may not be fully aware of relevant regulatory requirements or compliance standards.
- AI models are trained on existing data, which may contain biases or inaccuracies that can propagate into the assessment.
- AI is susceptible to adversarial attacks that can manipulate its output and compromise the assessment.



Architecture Decision Records (ADRs): Documenting Security Choices

AP ADRs document every security-relevant architecture decision, creating institutional memory.



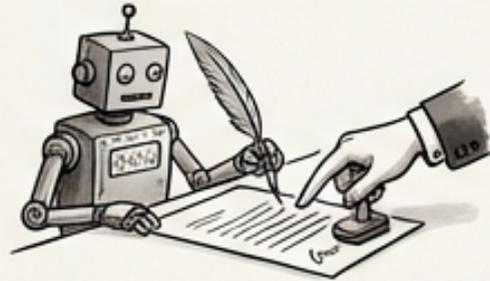
AC Each ADR should include the context, decision, consequences, and alternatives considered.



AC ADRs enable future reassessment of architecture decisions based on changing threats and requirements.



AD AI can draft ADRs, but humans must validate the security reasoning and accuracy.

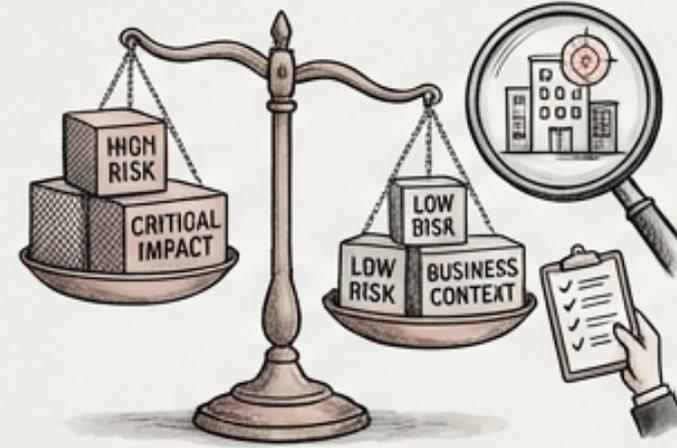


AF A consistent format for ADRs ensures clarity and ease of understanding.

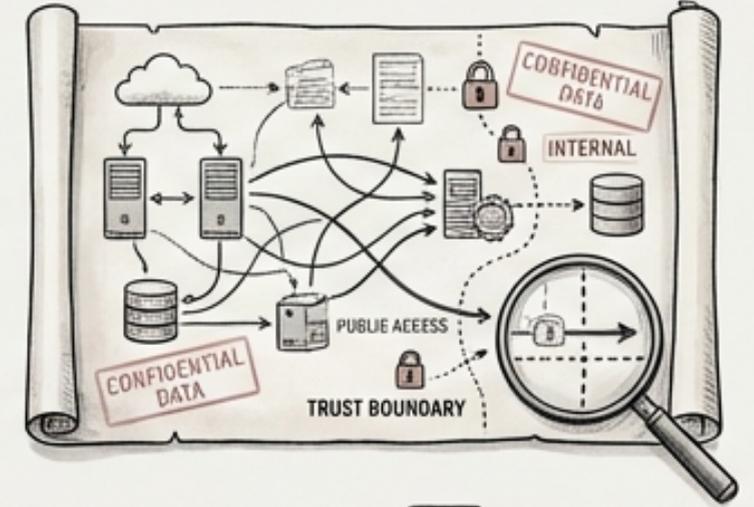


Key Deliverables of an Architecture Security Assessment

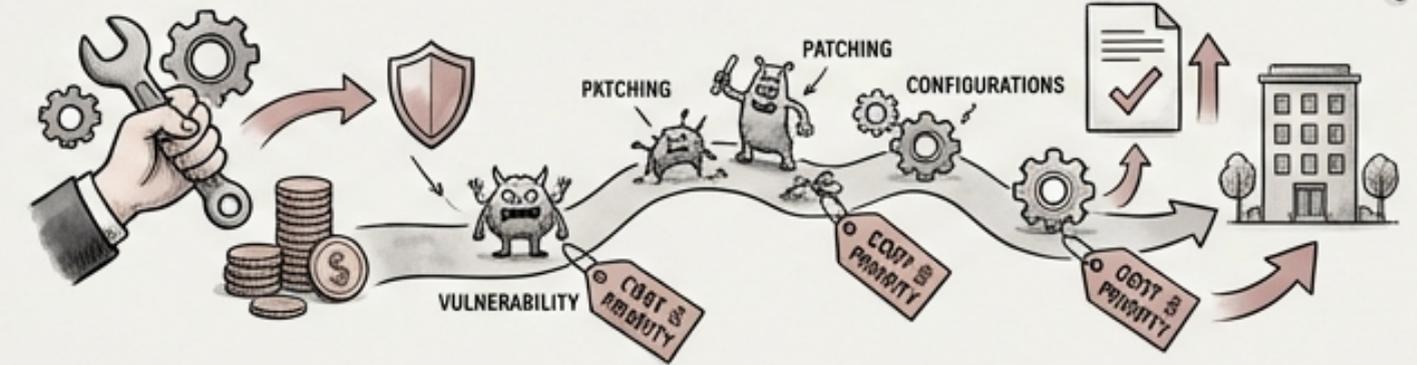
 **Risk-ranked findings with business context:** Prioritize vulnerabilities based on their potential impact and likelihood of exploitation.



 **Architecture diagrams annotated with trust boundaries and data classification:** Visualize the system's security architecture and data flows.



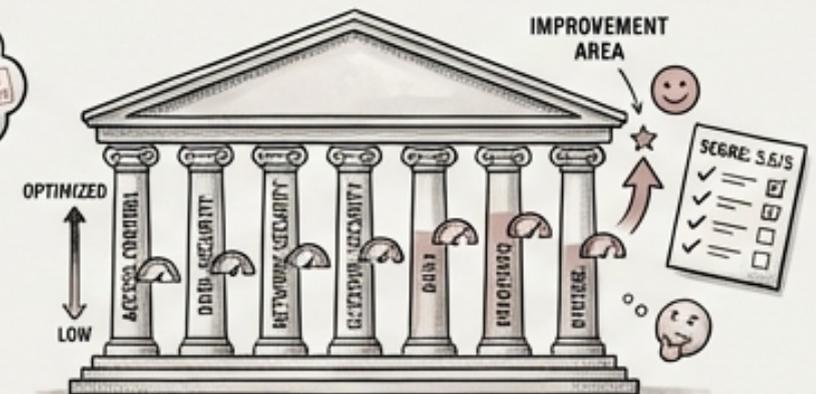
 **Remediation recommendations prioritized by risk and implementation cost:** Provide actionable guidance for addressing identified vulnerabilities.



 **Residual risk acceptance documentation:** Document any risks that are not fully mitigated and the rationale for accepting them.



 **Maturity scores** for each of the seven security domains, highlighting areas for improvement.

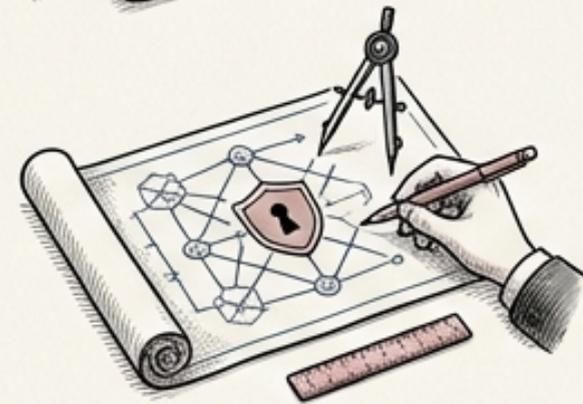


Actionable Insights: Leveraging ASA for Secure Development

Secure Development Lifecycle Optimization



- ASA helps identify and mitigate security vulnerabilities early in the development lifecycle, reducing rework and costs.



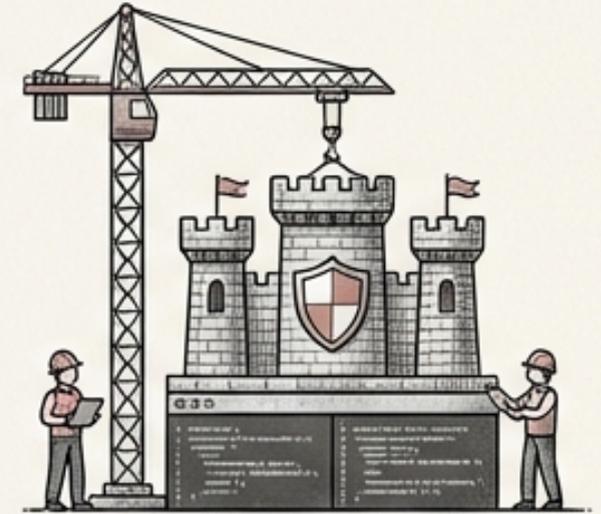
- By incorporating ASA into the development process, developers can build more secure and resilient applications.

- Understand and apply the security architecture patterns discussed to improve the security of your designs.



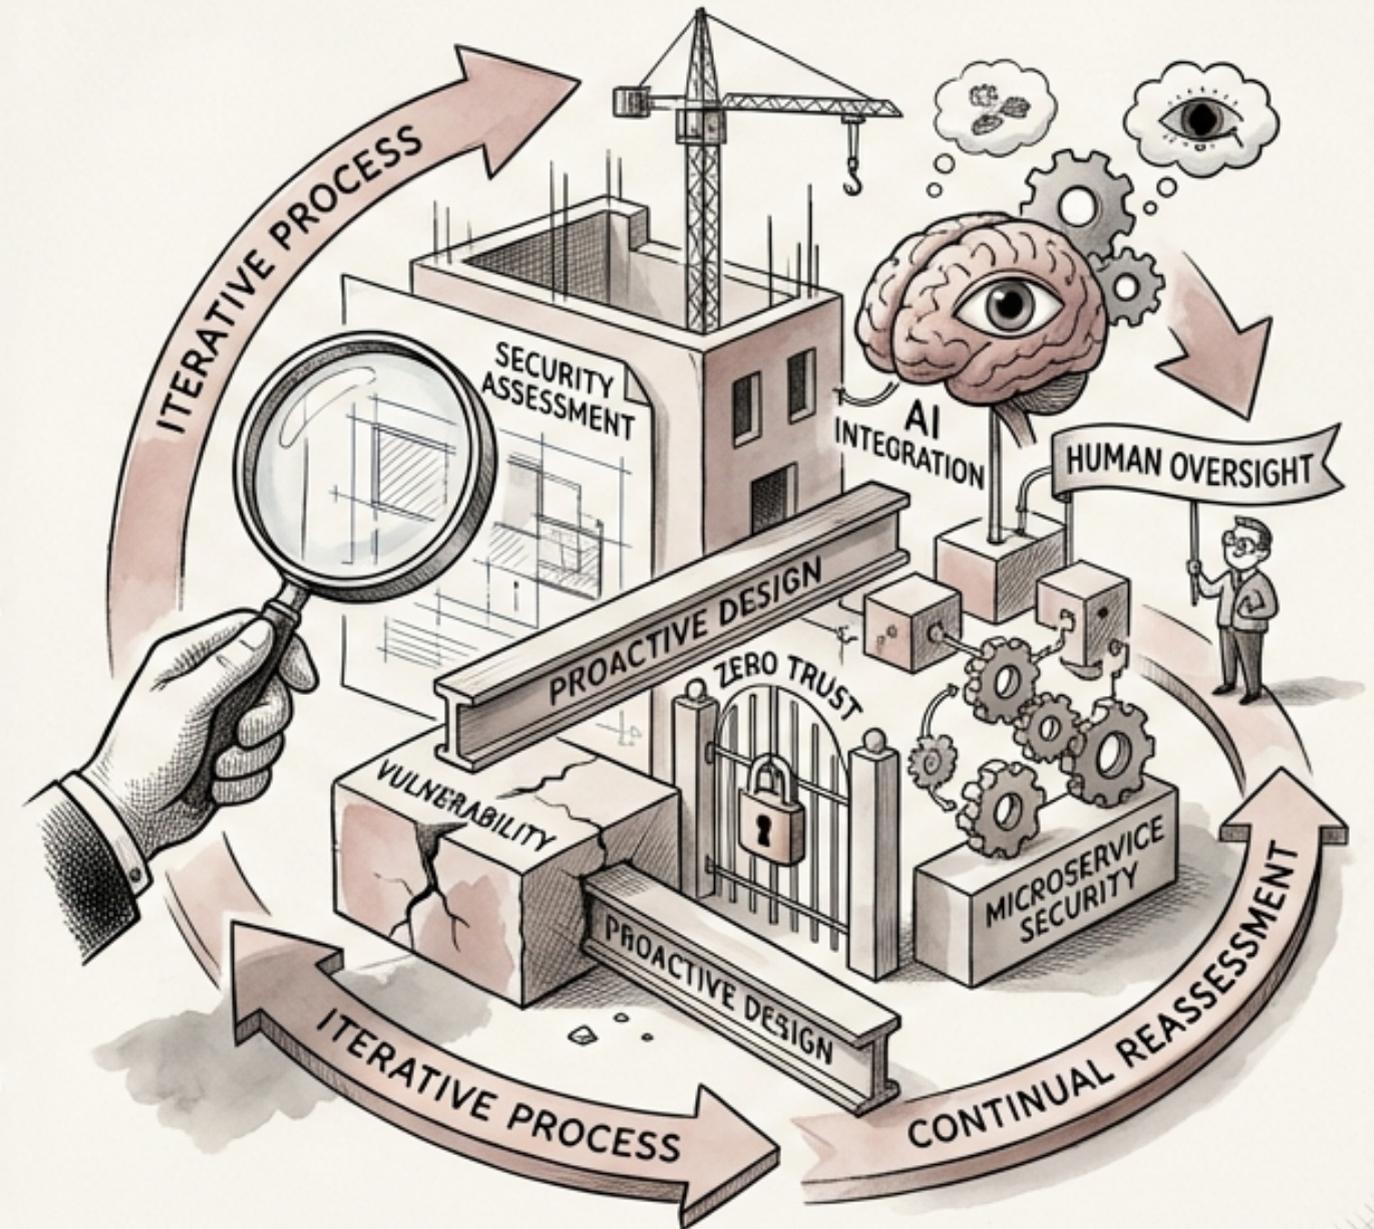
- Document all security-relevant architecture decisions in ADRs to facilitate knowledge sharing and future reassessment.

- Collaborate with security experts to ensure that architecture decisions are aligned with security best practices.



Conclusion: Building a Security-First Architecture

1. Architecture Security Assessment is a critical step in building secure and resilient systems.
2. By proactively evaluating system design, we can identify and mitigate vulnerabilities before code is written.
3. Leveraging patterns like Zero Trust, Microservice Security, and AI Integration, helps to build robust architectures.
4. ASA is an iterative process; continual reassessment is vital to maintaining security posture as threats evolve.
5. Remember to document all relevant decisions and keep human oversight a key element when implementing AI.



Thank You

- Questions?

