Professional Working Guide

AI-Assisted Application Risk Assessment

A practical prompt library and workflow for Quantum-Aware Cyber Risk

PURPOSE

Use this guide to turn the course's AI reflection questions into an authorized, application-specific risk assessment. It is designed for a professional to run in their own approved AI environment and convert the discussion into a practical set of findings, owners, and next actions.

This guide contains all 61 AI reflection prompts from the course, organized by the original course location. The prompts are for defensive analysis only: use them to identify potential exposures, not to test, exploit, or access systems.

Start Safely

- Use an AI service approved by your organization. Follow its data-handling, retention, and privacy requirements.
- Do not paste passwords, API keys, private certificates, customer records, production logs, source code, or unredacted architecture diagrams into a public AI tool.
- Describe systems at a useful but non-sensitive level. Replace names with labels such as [Customer API], [Identity Provider], or [Cloud Storage].
- Treat AI output as a hypothesis list. Validate each finding with the accountable engineering, security, privacy, or vendor owner.

What to Prepare

Before using the prompt library, prepare a short application profile. This gives the AI enough context to help without exposing sensitive operational detail.

Topic	Safe information to provide
Application purpose	What the application does, who uses it, and the business outcome it supports.
Data categories	Examples: customer account data, payments, health data, documents, logs, backups, metadata. Do not include records or identifiers.
Architecture	High-level components: web/mobile client, APIs, databases, cloud services, identity provider, queue, analytics, backup/archive.
Trust boundaries	Where users, administrators, services, vendors, and third parties exchange

Topic	Safe information to provide
	data or access.
Cryptography	Known protocols, certificate use, key management approach, encryption at rest/in transit, and unknown areas to investigate.
Retention	How long data, logs, backups, exports, and vendor-held copies may persist.
Current concerns	Recent migrations, legacy systems, acquisitions, audit findings, or areas that are not well documented.

Reusable Context Prompt

Paste this first, then add the application profile above. Follow it with one or more prompts from the library.

COPY / PASTE

You are supporting an authorized defensive risk assessment for an application. Use only the high-level, non-sensitive information I provide. Do not suggest exploitation steps, credential collection, or unauthorized testing. Identify assumptions, questions to validate, likely risk themes, and practical next actions. For each finding, state: (1) why it matters, (2) evidence or information needed to validate it, (3) likely owner, and (4) a priority of High, Medium, or Low. If information is missing, ask focused clarifying questions rather than inventing details.

Five-Step Assessment Workflow

1. Map the application. Share the application profile and ask the AI to identify missing information, trust boundaries, long-life data, and likely dependencies.
2. Run the relevant prompt set. Start with the course-start prompt, then choose the prompts most relevant to data, cryptography, architecture, vendors, migration, visibility, and governance.
3. Turn responses into evidence requests. For each potential risk, ask what diagram, inventory, configuration, contract, log source, or owner interview would confirm or refute it.
4. Create a findings register. Capture the risk statement, affected asset or data, validation evidence, owner, priority, and target date.
5. Review with people who own the system. Use the AI output to guide the discussion; do not make risk decisions solely from generated answers.

Suggested First Conversation

COPY / PASTE

Using the application profile below, help me identify where Harvest Now, Decrypt Later risk could exist. Focus on long-life data, identity, backups, vendors, cryptographic dependencies, and visibility gaps. Return a prioritized list of hypotheses, the evidence needed to validate each one, and the role that should own the next action.

[Paste my sanitized application profile here]

Findings Register

Use one row per validated or plausible finding. This can be copied into a spreadsheet, ticketing system, or risk register.

Field	What to capture
Risk statement	A plain-language statement of the exposure or decision gap.
Affected asset / data	Application component, data category, trust relationship, or vendor.
Why it matters	Business, regulatory, operational, or long-term confidentiality impact.
Validation needed	Inventory, diagram, configuration review, contract, log source, or interview.
Owner	Named team or accountable role.
Priority and next step	High/Medium/Low, a concrete action, and a target date.

Complete AI Prompt Library

The prompts below are reproduced from the course. Use them in the context of your own authorized environment. Questions written from an adversary perspective are intended to surface defensive blind spots; ask for risk analysis and validation steps, not attack instructions.

Course Start: Orientation

1. Help me identify where Harvest Now, Decrypt Later risk could exist in my organization, focusing on long-life data, identity, backups, vendors, and visibility gaps.

Module 1: The HNDL Reality Shift

1. If you could not break encryption today, what would you collect anyway?

2. What data would still matter if decrypted in 10 years?

3. What could you collect without triggering detection?

4. If you were collecting quietly, where would you start?

5. What becomes more valuable after decryption?

6. What are you choosing not to see?

7. What assumption has never been validated?

8. What is the first uncomfortable question your team should ask?

Module 2: Cryptography Under Pressure

1. Where would you look for encryption that no one tracks?

2. If identity trust breaks, what else breaks with it?

3. Where is TLS terminated, and what happens after?

4. What would an adversary gain from your backups in 10 years?

5. What internal trust would you exploit first?

6. What crypto dependency would surprise you most if exposed?

Module 3: Data Lifecycle Risk

1. What data would still matter if decrypted in 10 years?

2. What could be collected today without being noticed?

3. What identity relationships would matter most in 10 years?

4. What would your backups reveal if decrypted?

5. What data would create the greatest future liability?

Module 4: Leadership and Organizational Risk

1. If teams disagree, where would you exploit the delay?
2. What risk exists because no team owns it?
3. Which decision in your organization requires too many approvals?
4. What risk grows because you are focused elsewhere?
5. What decision would move forward if alignment improved?

Module 5: Adversary Models

1. What data would you collect today to exploit later?
2. What could you collect without changing the system?
3. What would make your organization worth waiting for?
4. If you were the adversary, where would you start tomorrow?

Module 6: Architecture Failure Points

1. What connection would you exploit first?
2. If identity trust is compromised, what systems follow?
3. What internal communication would you intercept first?
4. Where does encryption stop in your architecture?
5. What third-party system holds your data longer than you do?
6. What trust relationship would create the greatest exposure if compromised?

Module 7: Post-Quantum Cryptography Reality

1. What breaks when organizations assume this is easy?
2. What system in your environment cannot be easily upgraded?
3. What dependency delays your entire migration?
4. What risk grows if you wait?
5. What system prevents you from moving forward?

Module 8: Supply Chain and Third-Party Risk

1. Which external system would you target to access the most data indirectly?
2. What system holds your data longer than you expect?
3. What vendor delays your entire migration?
4. What system trusts a system you have never evaluated?
5. If you were the adversary, which vendor would you target first?

Module 9: Migration Complexity

1. What system dependency would you exploit during transition?
2. What dependency is missing from your inventory?
3. What system prevents your migration from progressing?
4. What gap exists only during migration?
5. What dependency defines your entire migration timeline?

Module 10: Detection and Visibility Gaps

1. How would you collect data without creating a signal?
2. What would passive collection look like in your logs?
3. What risk exists even if your dashboards are green?
4. What risk exists even if nothing is detected?

Module 11: AI-Assisted Gap Discovery and Exposure Analysis

1. What would you assume is exposed if you had no visibility?
2. Where is trust assumed? What dependencies are undocumented? What systems exchange sensitive data? Where does encryption terminate?
3. What data would you store for future use?
4. Architecture: where are controls assumed? Data: what lasts beyond 10 years? Migration: what breaks first? Leadership: what is unvalidated?
5. What question are we avoiding because it would force us to change our roadmap?
6. What are we not looking at that we should be?

Module 12: Strategy, Execution, and Leadership Translation

1. What happens if the organization does nothing?

Close the Loop

A useful assessment ends with a short review meeting: confirm the highest-priority findings, assign owners, agree the evidence to collect, and decide which items belong in the architecture roadmap, vendor review, cryptographic inventory, or security backlog. Re-run the prompt sets when the application, vendors, or migration plan materially changes.