Operation Quiet Rollback.

A five-inject red-team tabletop exercise for hybrid PQC deployments.

SERIES	TOOL	VERSION	LICENSE
QUANTUM · TOOLS	04 OF 05	0.1.0 · 2026	CC BY 4.0

THE SETUP

An organization has just deployed hybrid TLS 1.3 (X25519 + ML-KEM-768) on its externally facing API gateway. The legacy internal load balancer does not support ML-KEM and silently negotiates classical-only cipher suites with the gateway. A network-position attacker begins passively recording all traffic. The blue team has been told 'the deployment is complete' and has not yet been informed of the load-balancer gap.

The exercise covers five injects over a single working session. Each maps to a structural pitfall from the security-architects briefing. The objective is not to 'win' — it is to surface which gaps the team can detect, which they would miss in production, and which controls (logging, MTU testing, PKI validation, IDS coordination) need to be added before deployment.

TIME BUDGET — 90 MINUTES

0 – 10 min	Facilitator reads the setup. Blue team is told only what's in the setup paragraph; the loa d-balancer detail is withheld.
10 – 75 min	Five injects, twelve minutes per inject. Pause after each for blue-team response and red-t eam rebuttal.
75 – 90 min	Debrief. Score against the five questions. Identify the three highest-priority gaps to fix before the next deployment.

Roles & logistics.

Four roles. One facilitator. No live systems required.

FACILITATOR (1)
Runs the clock. Reveals injects. Scores the debrief. Holds the inject pack and the time budget.

RED TEAM (1–2)
Represents the network-position attacker. Reads inject content aloud. Pushes back when blue claims a detection that isn't actually wired in production.

BLUE TEAM (4–6)
Security operations + network engineering + identity + PKI representation. Their job: respond to each inject with what would actually happen today.

OBSERVER (OPTIONAL)
Takes notes for the post-exercise writeup. Does not participate in the discussion. Useful for the post-mortem document.

MATERIALS

- This PDF, printed double-sided.
- Five inject cards (cut from pages 3–7 of this PDF, or cite the inject numbers verbally).
- One debrief scoring sheet per blue-team member (page 8).
- A whiteboard or large notebook for the running incident log.
- A timer or visible clock.

No live systems are required. This is a discussion exercise, not a fire drill.

INJECT 01 · CIPHER-SUITE DOWNGRADE

The load balancer is silently classical-only.

RED TEAM REVEALS

Wireshark capture shows that the TLS cipher suite negotiated between the API gateway and the internal load balancer is ECDHE-only on 40% of sessions. The hybrid ML-KEM exchange is not being used on those sessions. openssl s_client -cipher output confirms the load balancer offers only classical groups.

BLUE TEAM MUST

Identify whether cipher-suite negotiation logging exists in the gateway. Determine whether anyone would notice this in production today. Articulate what alert or dashboard would have surfaced the gap.

PITFALL **DOWNGRADE ATTACK INVISIBILITY**

INJECT 02 · HNDL CAPTURE DURING THE GAP

Adversary begins recording classical-only sessions.

RED TEAM REVEALS

Beginning at the moment of the partial deployment, the attacker is HNDL-capturing the classical-only traffic on the internal segment. The captures include API tokens, customer PII, and signed JWTs from the identity provider. The captures persist for the indefinite future.

BLUE TEAM MUST

Identify which sessions are at HNDL risk. Determine whether session-level cipher-suite metadata is retained in logs for retroactive analysis. Decide whether to issue a retroactive risk advisory to privacy and legal.

PITFALL **SHADOW CRYPTOGRAPHY ON UNPROTECTED PATHS**

INJECT 03 · MTU FRAGMENTATION

Vendor patch increases handshake beyond 1,500 bytes.

RED TEAM REVEALS

The load-balancer vendor releases a PQC patch that raises the handshake size past the standard 1,500-byte Ethernet MTU on certain client paths. After the patch, intermittent connection failures appear for a subset of mobile clients. Stateful DPI in the path is silently dropping fragmented IP packets.

BLUE TEAM MUST

Identify whether MTU was tested at every boundary in the path before the patch was deployed. Determine the rollback plan — and whether the rollback would re-introduce the classical-only condition from Inject 01.

PITFALL **MTU FRAGMENTATION UNDER DPI**

INJECT 04 · IDS / IPS FALSE POSITIVES

SIEM flags PQC traffic as anomalous.

RED TEAM REVEALS

The IDS — which has not been updated since before the hybrid deployment — flags the PQC handshakes from a newly-rolled department as anomalous. The volume of alerts is high enough that security operations is now systematically ignoring the alert class, including a small number of genuine probe attempts hidden in the noise.

BLUE TEAM MUST

Identify the change-ticket trail for the IDS signature update — was it part of the deployment checklist or omitted? Determine the path to whitelist legitimate PQC traffic at the correct signature level.

PITFALL **FIREWALL / DPI / IDS INCOMPATIBILITY**

INJECT 05 · PKI CHAIN DEPENDENCY

Code-signing CA cannot issue ML-DSA.

RED TEAM REVEALS

An emergency software release needs to ship in the next four hours. The internal code-signing CA does not yet support ML-DSA. Hot-fix is now blocked behind a PKI dependency nobody on the architecture team identified during the deployment planning.

BLUE TEAM MUST

Identify whether the CA dependency was on the inventory before deployment. Determine the exception process — does it require classical signing, dual-cert, or a delay to the release? Document the decision for post-exercise review.

PITFALL **PKI CHAIN FAILURE**

Debrief.

Five questions. Three actions. Honesty matters more than the score.

SCORING — CIRCLE Y OR N

01.	Was the classical downgrade detected within the scenario timeframe?	Y	N
02.	Was a rollback plan executed without full session loss?	Y	N
03.	Did the HNDL exposure window have a defined response?	Y	N
04.	Were IDS rules updated as part of the deployment checklist?	Y	N
05.	Was the CA dependency identified during inventory?	Y	N

INTERPRETATION

- 3 or more Y answers — the deployment is reasonably defensible. Schedule the next exercise.
- 2 or fewer Y answers — the deployment should not ship without the missing controls. Block production rollout.
- Either way — the three lowest-scoring areas become named work items for the next sprint with owners, dates, and a follow-up at the next quarterly review.

THE THREE NAMED ACTIONS (FILL IN)

1. _____

Owner: _____ Date: _____ Review at: _____

2. _____

Owner: _____ Date: _____ Review at: _____

3. _____

Owner: _____ Date: _____ Review at: _____