Commvault

FRONTIER AI RESILIENCE CHECKLIST

FRONTIER AI CHANGED THE MATH. HAS YOUR RECOVERY PLAN KEPT UP?

Frontier AI models are able to identify vulnerabilities at a pace that traditional remediation programs were not built to absorb. When exploitation windows compress from weeks to minutes, prevention alone cannot define readiness. The four steps below are where resilience planning has to start.

1 EVALUATE RECOVERY RISKS

ASK HARDER QUESTIONS THAN "DO BACKUPS EXIST?"

- ☐ Adopt "assume the breach" as a standing operating principle, not a crisis posture.
- ☐ Be explicit: Can critical systems be restored end-to-end without reintroducing threats?
- ☐ Confirm recovery environments are isolated from compromised production systems.
- ☐ Validate recovery plans are mapped to current dependencies, not last year's architecture and systems.
- ☐ Document your Mean Time to Clean Recovery (MTCR). Establish a baseline if one does not exist, and track new findings.

Why this matters: The gap between "we have backups" and "we can recover" is where organizations can get hurt. Assessing that gap honestly, before an incident forces the issue, is where resilience planning must start. MTCR is the number that matters; it's not a theoretical estimate, but a measured, validated time to a known-good state.

2 MAKE ISOLATION THE BASELINE

AIR-GAPPED, IMMUTABLE RECOVERY IS A STANDARD REQUIREMENT – NOT AN ADVANCED OPTION.

- ☐ Assume that some vulnerabilities will outpace normal remediation cycles. Plan accordingly.
- ☐ Review or implement a 3-2-1 strategy: 3 copies, 2 different media types, 1 geographically separate.
- ☐ Confirm recovery environments are structurally separated from production identity, network, and management planes, and not just access-controlled.
- ☐ Test data isolation, vaulting, and air-gapping strategies against Commvault® Cleanroom™ drills.
- ☐ If penetration testing is not already run against backup infrastructure, begin now.

Why this matters: Frontier AI–enabled reconnaissance can map backup topology before a payload executes. If recovery infrastructure is reachable from production, it can be compromised before it is needed. Protection from the current threat is only half the job. The other half is maintaining clean recovery options for the vulnerability you haven't patched yet, because that exploit is already happening.

3 PRIORITIZE SYSTEMS THE BUSINESS CAN'T LIVE WITHOUT

DEFINE YOUR MINIMUM VIABLE COMPANY (MVC) – AND BUILD RECOVERY SEQUENCING AROUND IT.

- ☐ Define your MVC: the smallest set of systems required to keep the business running.
- ☐ Map recovery dependencies (e.g., identity platforms, billing systems, operational databases, cloud services) and document the sequence.
- ☐ Extend dependency mapping to AI operational infrastructure: data pipelines, model repositories, vector databases, agentic workflows.
- ☐ Build recovery sequencing around that definition before an incident forces the conversation.

Why this matters: Most organizations spend the first 24–48 hours of an incident deciding what matters rather than restoring it. As AI becomes embedded in operations, recovery complexity grows. Sequencing cannot be improvised under pressure. It needs to be mapped, tested, and validated in advance.

4 TEST CONTINUOUSLY

AUTOMATE RESILIENCE AND TEST CONTINUOUSLY – NOT ON A CALENDAR SCHEDULE.

- ☐ Automate threat scanning, clean recovery point identification, dependency-aware restoration, and recovery orchestration.
- ☐ Run cleanroom drills against all MVC systems on a regular cadence – not once.
- ☐ Publish drill results internally against a documented baseline: average end-to-end recovery time including patching and application verification.
- ☐ Validate AI system recovery specifically: training datasets, vector databases, model weights, and inference pipelines.
- ☐ Build a regular cadence for executive review of drill results.

Why this matters: Most recovery plans are tested once, under controlled conditions, against an environment that has since changed. Continuous testing finds what annual reviews miss: the dependency that shifted last month, the recovery sequence that breaks under a new workload, the identity service that has three times as many restoration steps as the runbook assumes. Those gaps don't announce themselves. They surface during an incident.

To learn more, visit **commvault.com/cyber-resilience**