Compliance & Risk

How autonomous, on-device prevention affects regulatory obligations, breach disclosure, data sovereignty, cyber insurance, and board-level risk reporting.

Does Cyber Crucible collect encryption keys?

The shortest answer is, “No.”

There was a time where Cyber Crucible software was not stopping ransomware encryption pre-encryption. Instead was collecting possible encryption keys or settings, then using a variety of techniques to figure out the proper decryption package (keys included) to decrypt after a ransomware attack.

From the standpoint of what we collected, a good analogy might be a video camera that would take pictures of everything that might be a key to a locker. Then, a server would work out the various math to determine the proper key and settings for the proper file. Computer scientists would call this machine learning.

Before we moved on from the old “collect anything that might be a key” software model, we were organizing thousands of keys and encryption settings per machine automatically in attacks.

The cryptographic analysis software module has now evolved in an additional behavioral analytic module which is used as one data source to determine if software is trying to improperly access identity data or “data data” to steal or encrypt.

If an attack is prevented, do we still have to report a breach?

Short answer: Generally no. When an attack is stopped before execution and no data is accessed, altered, or exfiltrated, there is typically no breach to disclose. Disclosure obligations are usually triggered by unauthorized access to protected data — not by an attempt that failed.

A real example

In the financial services deployment where Cyber Crucible intercepted nearly 10,000 malicious processes, the outcome for the CISO included zero regulatory disclosure. Because the attacks were stopped pre-execution and no data was touched, there were no required notifications to customers, vendors, or government regulators — and no headlines, which preserved customer trust.

An important caveat

Disclosure requirements vary by jurisdiction, industry, and contract, and the specific facts of an incident matter. Cyber Crucible is not a law firm and this is not legal advice. The reliable principle is that prevention narrows your disclosure exposure dramatically compared with detection-after-the-fact — but your counsel and compliance team should make the determination for any given event.

How does Cyber Crucible support HIPAA compliance in healthcare?

Short answer: All analysis happens locally at the kernel level, so protected health information never leaves the endpoint for security processing. That avoids the HIPAA exposure created by security tools that upload sensitive files or keys to third-party clouds, and it stops attacks without taking clinical systems offline.

The problem with vendor-hosted security

Many vendor-hosted SOC and MDR arrangements require uploading sensitive keys, files, or telemetry to third-party cloud databases. For a covered entity, that creates real compliance risk and expands the number of parties handling protected data.

Local analysis, local data

Cyber Crucible analyzes on the endpoint. Sensitive assets never leave the secure boundary of the device — which supports both data sovereignty and the regulatory position that goes with it.

Continuity matters clinically

Containment that requires isolating a machine or forcing a reboot is not acceptable when patient care depends on it. Selective suspension of only the malicious process means the attack is neutralized in under 200 milliseconds while medical operations, connected devices, and clinical systems keep running.

How does Cyber Crucible support FERPA and student data privacy?

Short answer: FortressAI checks data access on the device, so student records and family data can't reach a public AI tool by accident. Cyber Crucible stops ransomware autonomously without requiring a 24/7 security operations center — which most districts and institutions cannot fund.

The two education problems

Schools and universities are prime ransomware targets while rarely having budget for round-the-clock SOC staffing. At the same time, classroom AI adoption has moved much faster than governance for it.

Addressing both

- **Ransomware without a SOC:** autonomous, kernel-level prevention stops attacks in under 200 milliseconds with no analysts to hire and no 24/7 monitoring contract to fund.
- **AI use aligned with FERPA:** FortressAI evaluates every prompt on the device, so protected educational records and personal family data don't end up in a public AI system — including as training data.

That combination protects students and staff without adding significant IT overhead.

What does data sovereignty mean for Cyber Crucible deployments?

Short answer: Your data stays where you put it. Analysis and enforcement happen on the endpoint, and Cyber Crucible can be deployed fully on-premises — running in physically secured server racks rather than a public cloud, with no third-party platforms, no external access, and no data sharing.

Why centralized security data is a liability

Security architectures that centralize sensitive material — encryption keys, session tokens, identity data — create a single point of failure and a high-value target for criminal and state-sponsored attackers alike.

There's a second risk that receives less attention: a government could compel access to centrally held data through a classified subpoena or national security letter, without the customer's knowledge or consent. That's a profound data-sovereignty exposure with no transparency.

The on-premises option

For organizations requiring absolute control, Cyber Crucible offers full on-premises installation, including air-gapped deployment and multi-tenant architecture suitable for managing multiple departments or agencies from a single secure instance.

What compliance risk does encryption key escrow create, and why doesn't Cyber Crucible do it?

Short answer: Storing encryption keys centrally creates a single point of failure, a high-value target, and exposure to compelled disclosure. Cyber Crucible's prevention approach doesn't depend on capturing or escrowing keys, so there is no central key store to secure, subpoena, or breach.

> For the direct product answer, see the existing article **"Does Cyber Crucible collect encryption keys?"** This page covers the compliance and data-sovereignty reasoning behind that answer.

Why this is deliberate

Cyber Crucible originally developed and patented key-capture technology for ransomware defense. By 2021, ransomware strains had evolved to defeat it — using custom encryption libraries, unique per-execution implementations, and streaming ciphers immune to key capture. We presented that cryptanalysis publicly at BSides Pittsburgh in 2021.

The risk we chose not to create

Beyond no longer working, key capture requires centralized storage of encryption keys. That introduces a single point of failure, a high-value target, and the possibility of compelled disclosure without customer knowledge.

Preventing the attack before encryption begins removes the need for keys entirely — which is both a stronger defense and a smaller compliance surface.

How does autonomous prevention affect cyber insurance and board reporting?

Short answer: Prevention changes what you report. Instead of documenting incidents, downtime, and remediation costs, you report attacks that were stopped with no business interruption, no disclosure, and no ransom paid.

What the board actually cares about

In the financial services deployment referenced throughout this knowledge base, the CISO was able to report four outcomes:

- **Zero business takedown** — attacks suspended in memory; operations and legitimate transactions continued.
- **Zero regulatory disclosure** — nothing was accessed, so nothing required notification.
- **Zero PR damage** — no headlines, customer trust preserved.
- **Strong ROI** — a fraction of the cost of the legacy stack did the work the legacy stack missed.

On ransom payments

Roughly 90% of ransomware victims paid the ransom in 2023. Cyber Crucible customers paid nothing, because the attacks did not reach the encryption stage.

For insurers

Underwriters increasingly ask what controls prevent an incident rather than what tools detect one. An autonomous, kernel-level control that operates without human response time is a substantive answer — though coverage terms are always set by your carrier.