Vendor Due Diligence & Security Trust

How Cyber Crucible answers enterprise and financial-institution vendor-risk questionnaires — what it holds and does not hold, how it supports regulatory obligations, how AI is governed, service levels, and how its architecture reduces third-party risk. Stated plainly and honestly.

Does Cyber Crucible have a SOC 2 report?

Short answer: No. Cyber Crucible does not currently hold a SOC 2 attestation, and it says so plainly rather than implying otherwise. The reason is architectural: Cyber Crucible is a locally-executing software product, not a cloud custodian of customer data, so the SOC 2 service-organization model — which attests to how a provider stores, processes, and transmits *your* data on its own systems — does not map cleanly to it.

Why most customers have not required one

SOC 2 attests to the controls of a service organization that holds customer data. Cyber Crucible's design removes that premise:

- Threat detection and response run on the endpoint, at the kernel level, on the customer's own device.
- Customer files, credentials, encryption keys, and session tokens are never collected, transmitted, or stored by Cyber Crucible.
- The platform can be deployed fully on-premises or air-gapped, inside the customer's own infrastructure.

Because the highest-risk categories of data are never in Cyber Crucible's custody, a large part of the vendor-risk surface a SOC 2 is meant to address is removed by design rather than merely controlled.

What stands in its place

For reviewers who work from a questionnaire, Cyber Crucible provides a prepared vendor security package that supplies the underlying control evidence directly, and maps its controls to the SOC 2 Trust-Service Criteria and the NIST Cybersecurity Framework so a reviewer can complete a standard questionnaire from it. The package is available under a mutual NDA from **dpo@cybercrucible.com**.

The assurance roadmap is kept under evaluation against customer needs; this page will be updated if that position changes.

Is Cyber Crucible ISO 27001, PCI-DSS, CMMC, or FedRAMP certified?

Short answer: No. Cyber Crucible has not obtained, and does not claim, ISO/IEC 27001, PCI-DSS, HIPAA, CMMC, or FedRAMP certification or authorization. It states this directly, because claiming a certification a vendor does not hold is exactly the kind of inconsistency that undermines trust in a due-diligence review.

What this means in practice

- **PCI-DSS:** Cyber Crucible does not process, store, or transmit cardholder data, so it is out of scope as a card-data processor while still supporting a customer's endpoint-security control objectives.
- **HIPAA:** because no customer content is collected, Cyber Crucible avoids creating a protected-health-information custody relationship; a Business Associate Agreement can be executed where a customer requires one.
- **ISO 27001 / CMMC / FedRAMP:** not held, and not represented as held.

What Cyber Crucible does hold

Independent, objective validation exists where it matters most — see *What independent security validation does Cyber Crucible have?* Framework references elsewhere in this knowledge base are self-assessed mappings that help a reviewer relate Cyber Crucible's controls to a standard they already use; they are not third-party audit opinions.

What independent security validation does Cyber Crucible have?

Short answer: Cyber Crucible's Windows kernel drivers are objectively tested and cryptographically signed under Microsoft's Windows Hardware Compatibility Program (WHCP). That is an external, independent attestation of the quality and tamper-resistance of the product's most privileged component — the code where correctness matters most.

Why WHCP is meaningful

Most of the endpoint market avoided kernel-level development because it is genuinely difficult and expensive. Cyber Crucible operates at the kernel deliberately, and submits its drivers to Microsoft's program so that an outside party — not just Cyber Crucible — validates that the most sensitive code is correct and resistant to tampering.

- The validation is **scoped precisely to the driver**, not represented as a whole-product certification.
- It is an objective, repeatable test rather than a subjective administrative audit.

How it fits the bigger picture

For a reviewer used to looking for a SOC 2 report, WHCP is a concrete, independent data point that a SOC 2 does not provide: direct validation of the kernel component. It complements — rather than replaces — the control evidence in Cyber Crucible's prepared vendor package.

How does Cyber Crucible support GLBA and financial-institution vendor requirements?

Short answer: As a software service provider to a financial institution, Cyber Crucible supports the institution's obligations under the GLBA Safeguards expectations and the interagency information-security guidelines — primarily by never collecting the non-public personal information (NPI) those rules are written to protect. Cyber Crucible is a software vendor, not a chartered financial institution, so it is not itself examined; its controls are built to be examinable as part of the institution's third-party risk program.

Where the architecture does the work

- **No NPI collected.** The product runs on endpoints that may handle NPI without extracting it. There is no consumer financial information for Cyber Crucible to use, share, or re-disclose, which directly supports the GLBA Privacy Rule / Regulation P position.
- **Documented security program.** Access control, encryption, breach response, and change control are documented and evidenceable — the safeguards a bank's program looks for in a service provider.
- **Due-diligence evidence on hand.** A prepared vendor package is structured to serve as the due-diligence and ongoing-monitoring evidence a bank needs under current interagency third-party risk guidance.

The honest boundary

Cyber Crucible supports a customer's compliance obligations; it does not, by itself, make an organization compliant, and this is not legal advice. Specific, evidence-grade detail is provided to a bank's reviewers under NDA. See also *Does Cyber Crucible meet FFIEC and third-party risk management expectations?*

Does Cyber Crucible meet FFIEC and third-party risk management expectations?

Short answer: Cyber Crucible is built to be examinable against the security, operations, and business-continuity expectations that examiners apply to bank technology, and its prepared vendor package is organized to serve as the due-diligence evidence a financial institution needs under the 2023 interagency guidance on third-party relationships.

What a bank's program can rely on

- Control detail mapped to the domains found in standard questionnaires (SIG, CAIQ) and to the SOC 2 Trust-Service Criteria and NIST CSF.
- Independent driver validation (WHCP) as an external quality attestation.
- A documented, tested business-continuity posture, with the important property that endpoint protection continues even during a management-backend outage.
- Insurance evidence, financial information, and a sub-processor list available to reviewers on request.

Why the architecture lowers third-party risk

The single most useful fact for a bank's vendor-risk analyst is that Cyber Crucible never collects customer files, credentials, or keys. The data a reviewer worries about a vendor mishandling is not in Cyber Crucible's custody to mishandle. That is a structural reduction in risk, not a promise to manage it well.

What security and compliance frameworks does Cyber Crucible align to?

Short answer: Cyber Crucible maps its controls to the NIST Cybersecurity Framework, relevant NIST SP 800-53 control families, and the CIS Critical Security Controls, and it maps to the SOC 2 Trust-Service Criteria for reviewer convenience. These are **self-assessed alignments**, offered so a reviewer can relate Cyber Crucible's controls to a standard they know — not certifications or third-party audit opinions.

Alignment, stated honestly

Framework	Status
NIST Cybersecurity Framework (CSF)	Controls self-mapped; not a certification
NIST SP 800-53 control families	Aligned where applicable; not an authorization
CIS Critical Security Controls	Aligned; not certified
SOC 2 Trust-Service Criteria	Mapped for reviewer convenience; no report held
Microsoft WHCP (driver)	Held — an independent, external validation

Why say it this way

A vendor that lists frameworks it has not been audited against invites the reviewer to distrust every other answer. Cyber Crucible states what is a certification (WHCP) and what is a self-assessed mapping (the rest), so a reviewer can weigh each accordingly.

Does Cyber Crucible use AI, and is it a large language model?

Short answer: Cyber Crucible uses proprietary AI — but only during development, not on your endpoint, and it is not a large language model. Its Genetic AI is a design-time discovery tool that identifies the deterministic set of behavioral variables indicating an attack. Those findings become fixed, deterministic kernel-level heuristics. At runtime, no AI, no LLM, and no adaptive model runs on the device.

Why this is a stronger position, not a weaker one

- **No runtime model means no drift.** Endpoint behavior is deterministic, repeatable, and testable — not probabilistic.
- **No live inference surface.** There is no on-device model for an attacker to manipulate or poison.
- **No customer data in any model.** Customer content, credentials, and keys are never used to train, feed, or tune a model, and no customer data leaves the endpoint for AI processing.

For a financial-institution or regulated reviewer, this is usually the AI answer they are hoping to hear: capability without a live, opaque model making decisions on the endpoint.

Related

See *How is AI governed in Cyber Crucible's development lifecycle?* for how the design-time work is controlled.

How is AI governed in Cyber Crucible's development lifecycle?

Short answer: Because Cyber Crucible's Genetic AI is used only at design time, it is governed under the secure development lifecycle (SDLC) rather than as a live production system. The discovery work happens in a controlled internal environment, its output is validated and translated into deterministic kernel heuristics, and the resulting drivers are independently signed under Microsoft WHCP before release.

The governed lifecycle

Phase	Governance
Design	Security requirements set; AI discovery is part of the design phase, under the SDLC
Discovery	Genetic AI identifies deterministic attack-indicating variables using internal research datasets only
Derivation	Findings become fixed, deterministic kernel heuristics — no live model is deployed
Validation	Deterministic outcomes are repeatable and testable; models and drivers validated through QA
Release	Windows drivers WHCP-signed; changes follow formal change management

Data and integrity

No customer content, credentials, or keys are used to train, feed, or tune a model. Models are proprietary and built in-house, with no third-party model dependency. Model integrity is maintained through version control, SDLC/QA validation, and WHCP signing; confidentiality through a controlled development environment.

Where is Cyber Crucible's data processed — is it stored in the cloud?

Short answer: Threat analysis and response happen locally on the endpoint, not in the cloud. Cyber Crucible does not collect customer files, credentials, or keys, and it does not store customer content with a third-party public-cloud provider. The management and reporting backend runs on Cyber Crucible-operated infrastructure in the United States, and a fully on-premises or air-gapped deployment is available for organizations that require zero external data flow.

Why local processing matters

- **The protective decision never leaves the device.** Detect–Decide–Respond runs in the kernel on the endpoint, typically in under 200 milliseconds, with no cloud round trip in the critical path.
- **Backend-independent protection.** A management-backend outage does not reduce endpoint protection.
- **Deployment choice.** Organizations with strict data-residency or sovereignty mandates can run Cyber Crucible entirely within their own secured infrastructure.

What a reviewer can rely on

Because customer content is never collected, the question "where does our data go?" has a simple answer for the highest-risk categories: it does not go anywhere. Detailed infrastructure and sub-processor information is provided to reviewers under NDA.

What is Cyber Crucible's service availability and support SLA?

Short answer: Cyber Crucible's published Service Level Agreement commits to **99.9% monthly service availability** (excluding scheduled maintenance and causes beyond its reasonable control), with a defined downtime-credit remedy. Support response times are tiered by severity, measured from ticket entry through the support portal.

Support response targets

Severity	Description	Response
SEV1 — Critical	Service down or critically affecting production; no workaround yet	≤ 2 hours
SEV2 — Major	Service impaired in a way that limits protection or response	≤ 4 hours
SEV3 — Minor	Service impaired without affecting protection or response	≤ 12 hours
SEV4 — Minor	Non-critical issue, information request, or enhancement	≤ 48 hours

Where to find it

The full Service Level Agreement is public at cybercrucible.com/service-level-agreement, and support tickets are entered at support.cybercrucible.com. A committed breach or incident notification timeframe can be set contractually for regulated customers.

A note on availability

Because endpoint protection runs locally, the availability figure applies to the management and reporting backend. Protection on the device itself continues even if the backend is temporarily unavailable.

How does Cyber Crucible handle data retention and deletion?

Short answer: Cyber Crucible minimizes what it holds and retains data only as long as necessary to provide the service. It never collects customer content, credentials, or keys in the first place, so the highest-risk categories have nothing to retain. Behavioral telemetry and derived security metadata are aged off on defined schedules, and at the end of a contract all associated data is made available for deletion with confirmation.

The principle

- **Minimize first.** The most sensitive data is never collected, so there is nothing in those categories to retain or destroy.
- **Retain by purpose.** Behavioral telemetry and security metadata are kept only as long as they serve detection, investigation, and reporting, then aged off.
- **Delete on exit.** At end of contract, associated data is made available for deletion and Cyber Crucible provides confirmation.

What is documented

A documented destruction process covers both printed and electronic information. The specific retention periods by data type are provided to reviewers under NDA, and can be tightened by contract.

Does Cyber Crucible have a Data Privacy Officer?

Short answer: Yes. Cyber Crucible has engaged a contracted Data Privacy Officer with cross-jurisdictional expertise spanning the EU/UK GDPR, California CCPA/CPRA, HIPAA, PCI-DSS, and global data-protection regimes including the Saudi Arabian PDPL. The DPO is reachable through the monitored inbox **dpo@cybercrucible.com**.

What the DPO covers

- The privacy program and data-processing agreements.
- Data-protection impact and transfer-risk assessments, including Standard Contractual Clauses for cross-border transfers.
- Breach-notification determinations and support for a customer's own notification obligations.
- A knowledgeable point of contact for privacy and compliance questions during procurement.

Why a contracted DPO with global expertise

Several non-U.S. regimes — the Saudi PDPL among them — can be more restrictive than current U.S. law. Engaging a DPO fluent across GDPR, CCPA/CPRA, HIPAA, and global PDPLs lets Cyber Crucible support customers to the stricter standard, which generally satisfies the U.S. baseline as a subset.

Is Cyber Crucible subject to US export controls?

Short answer: Yes — Cyber Crucible's software is subject to the U.S. Department of Commerce Export Administration Regulations (EAR). It is **not** ITAR-controlled and is not administered by the Department of State. No special license is required to provide the software to customers, and export classification and screening are handled as part of Cyber Crucible's compliance program.

What this means for a buyer

- **EAR, not ITAR.** The product falls under Commerce Department jurisdiction, the framework that applies to most commercial encryption-bearing software — not the defense-article regime.
- **Patented, and reviewed.** The technology has been patented internationally following U.S. Patent and Trademark Office review that included Department of Defense review.
- **No licensing barrier.** Providing the software to a customer does not require a special export license.

Export classification detail is available to reviewers on request.

How does Cyber Crucible reduce third-party and vendor risk for regulated buyers?

Short answer: By removing the risk rather than only managing it. The three facts that answer the hardest vendor-risk questions are: Cyber Crucible collects no customer content, credentials, or keys; protection runs locally and survives any backend outage; and its AI lives only in development, never on your endpoint. Together these shrink the surface a due-diligence review is designed to probe.

The three structural risk reducers

- **Zero content.** No customer files, credentials, or keys are ever collected. That single fact answers the PCI, confidential-data, and "what if you're breached" lines of questioning — the data is not there to lose.
- **Backend-independent protection.** Detect–Decide–Respond runs locally, with an effective recovery-time objective of zero on the endpoint. A total backend outage does not reduce protection.
- **Deterministic AI.** AI is used in development only; deterministic heuristics run at runtime. No live model, no drift, testable outcomes.

Independent reinforcement

Where independent attestation carries weight, Cyber Crucible points to what is real: Microsoft WHCP driver signing, and — for the managed backend — the third-party attestations of the infrastructure providers it relies on. With no SOC 2 of its own, these carry the load honestly rather than being oversold.

How do I get Cyber Crucible's security due-diligence package?

Short answer: Request it from **dpo@cybercrucible.com**. Under a mutual NDA, Cyber Crucible provides a prepared vendor security package that answers the domains of the standard questionnaires (SIG, CAIQ) and stands in for a SOC 2 report where a reviewer would ordinarily look for one.

What the package and its companions include

- A prepared vendor security and trust package: architecture, data handling, encryption, access management, secure development, incident response, business continuity, regulatory alignment, and a control-framework mapping.
- One-page program summaries for compliance, audit, and AI/SDLC governance.
- A certificate of insurance, W-9, and financial information, available on request.
- Standard Contractual Clauses, a Transfer Risk Assessment, and a Data Processing Agreement, available on request.

Public documents you can read now

Document	Location
Service Level Agreement	cybercrucible.com/service-level-agreement
Mutual NDA	cybercrucible.com/mutualNDA
MSA & EULA	cybercrucible.com/msa-and-eula
Privacy Policy	cybercrucible.com/privacy-policy

Everything beyond the public agreements is provided to reviewers under NDA, subject to reasonable confidentiality measures.