THE CAMPUSCISO® TRUST ANCHOR

AI-augmented, not AI-generated.

Three structural commitments: how we use AI in client work, how we handle institutional data, and how we make the methodology transparent.

We publish this so institutions and partners can examine our approach and decide whether our methods fit their own workflows and expectations.

CONTENTS

What's inside.

Our work is AI-augmented, not AI-generated. This document outlines three structural commitments, allowing institutions and partners to understand the methods we use before, during, and after an engagement.

OUR POSITION

Commitments you can check, not claims you must trust.

> *Our work is AI-augmented, not AI-generated. AI streamlines the analytical groundwork. A senior expert reviews and validates everything before it reaches you.*

In the current artificial intelligence (AI) landscape, institutions are right to be skeptical about how and when AI is used in expert advisory work.

The line between "AI-augmented" and "AI-generated" has been blurred in marketing materials across the sector. Buyers often have to guess what was actually reviewed, where their data goes, and how expert methods are applied during their engagement.

The Trust Anchor is our answer: three structural commitments that define working with us.

01 **AI augmentation in client work**
How we use AI, and the two-checkpoint review discipline that governs every output.

02 **Client data handling**
Where institutional data lives, who can touch it, and what happens when the relationship ends.

03 **Methodology transparency**
The frameworks behind the work, published openly so the method can be read before, during, and after an engagement.

We publish this Trust Anchor so institutions and partners can examine our approach and decide whether our methods fit their own workflows and expectations.

01

AI Augmentation in Client Work.

Every deliverable passes two expert checkpoints before it leaves the practice.

THE KEY IDEA

AI streamlines the analytical groundwork. A senior expert reviews, interprets, and validates every output, twice, before it ever reaches you.

Our position is straightforward. AI used well, with the right human oversight, multiplies what one senior expert can deliver. Used poorly, with no human anchor, it produces output that looks plausible but hasn't been validated. The Trust Anchor is what closes that gap.

What we use AI for

AI and computational processes handle work that doesn't require expert judgment but does require time and consistency:

- Intake normalization
- Cross-reference against benchmarking data
- Pattern recognition across hundreds of assessments
- Initial draft structure
- Consistency checking across deliverables

These are tasks a junior consultant would do at a traditional firm. AI does them faster, more consistently, and at a cost that helps us fit institutional budgets. That's the strategic case for AI in CampusCISO's practice. It's not about replacing expertise. It's about expanding the capacity of senior experts, making our advisory support accessible to institutions that couldn't otherwise afford it.

The expanded affordability is the point. Higher education doesn't need more senior advisory support concentrated at well-resourced research universities; it needs that support accessible across the full range of institution sizes and budgets. AI extends that reach, and the expert review described below keeps the expanded reach from diluting the quality of support.

> *AI workflows are infrastructure. Expert judgment is the product.*

What expert review actually means

The founder reviews, interprets, and validates every CampusCISO deliverable before it's delivered. That review draws on nearly thirty years of higher education technology and security experience: CISO roles at the University of Chicago, the University of Arizona, and the University of Wisconsin-Whitewater; client-facing leadership roles at SunGard Higher Education and FireEye/Mandiant.

The cross-sector arc matters because it's what makes the review substantive. A career spent supporting security programs at hundreds of institutions, both as a campus CISO and as a vendor serving the sector, makes our AI-augmented workflows effective. It takes that depth of experience to identify when AI produced something that sounds right but isn't, and to tell when its output is better than the expert would have written from scratch.

CampusCISO is a deliberately solopreneur practice, and we keep investing in the systems that make that model work: the workflows, templates, prompts, research, and quality checks that direct how AI is used on each task. Think of it as an analyst's workbench built around expert judgment. It's what lets a single senior advisor deliver with the speed and consistency of a larger team. The investment in developing and maintaining this system matters because AI without designed workflows produces generic results.

Inside these structures, expert judgment applies at every step. We enforce two checkpoints every time:

First checkpoint. As the work proceeds, we review and approve each output received from AI analysis before adding it to the deliverable. We reject, refine, or accept every AI recommendation.

Second checkpoint. Before delivering it to the client, we review the assembled deliverable as a whole.

We delegate neither. We automate neither.

What you're actually buying

There are two kinds of AI assessment services, and they are different products. Advisory practices like CampusCISO produce expert-validated analysis informed by AI groundwork. Software platforms produce automated outputs that institutions or their consultants use as raw material.

Both have legitimate uses. The distinction matters for what the client is purchasing. With an advisory practice, the client is paying for expert interpretation, and AI becomes a research tool that makes the expert's time go further. With a platform, the client is paying for software access, with the institution or its consultant providing the interpretation. We operate the first model.

This matters for partners who build on CampusCISO outputs as part of their own client delivery. Whether you present a CampusCISO report within your own engagement or co-deliver alongside CampusCISO, you're presenting expert-reviewed analysis to your client, not raw output from a platform that hasn't been reviewed for accuracy. The Trust Anchor is our structural commitment behind that, and it's what makes CampusCISO outputs durable inputs to your work. Because that expert review is the product and can't be handed off, we work with a deliberately limited number of partners rather than pursuing high-volume work. That keeps our capacity matched to the standard your clients will hold you to.

02

Client Data Handling.

We built data handling into the practice, not around it.

THE KEY IDEA

We limit what you send us in the first place. What you do share lives in your Portal instance, is used only for your engagement, is never used to train AI models, and leaves our active systems when the relationship ends.

We treat responsible handling of client data as a core business obligation, not a management policy bolted on after the fact. Data policies can be revised or leave you wondering whether they're actually followed. Our philosophy grew out of decades spent in higher education leadership roles. We hold ourselves to the standard we spent decades expecting from our own vendors, and then took it further because the sector's concerns around AI, privacy, and data handling have never been higher.

Our vision behind providing cybersecurity benchmarking data also originated from that earlier higher education career. Academic culture asks for evidence-backed justification before making decisions. For years that meant watching teams spend weeks or months assembling cybersecurity comparison data by hand, because no shared reference existed for how the sector actually manages cyber risk. We built Cyber Heat Map® in large part to create a meaningful benchmark for higher education cybersecurity tool adoption. As we've expanded to support additional functions like IT policy development, we've kept the same focus: building the data institutions need to make defensible decisions. We know that reliable benchmarks require trust and open candor from the institutions who work with us. Since we launched, we've made how we collect, curate, and protect that data a core mission at CampusCISO.

Our data commitments serve two purposes at once. They protect each institution's specific findings, and they make the higher education benchmarking dataset trustworthy so institutions can have confidence using and contributing to the benchmark.

> *Protecting each institution and building a dataset the sector can rely on are the same commitment, seen from two angles.*

The benchmarking dataset's proprietary core comes from structured assessments completed during client engagements. Each engagement contributes capability assessments, policy library reviews, and tabletop evaluations as de-identified data. We also analyze publicly available information, watching for sector trends, to supplement this core. The same two-checkpoint discipline applies: we review each contribution at the moment it enters the dataset, and we review aggregate observations drawn from the dataset again before publication.

Where your data lives and who can touch it

We start by limiting what reaches us at all. CampusCISO provides strategic advisory support, so our work runs on structured assessment surveys, policy documents, and the information you choose to share, not the operational details of your environment. We do not access your tools, and we do not collect sensitive configurations, log files, or vulnerability scan output. Under our contract terms, we also do not knowingly store, process, or transmit student education records (FERPA), protected health information (HIPAA), payment card data (PCI), or similar regulated data, and clients agree not to send it without a written amendment first. This intentionally limited data footprint presents a smaller risk by design.

Where it lives. Your institution's working data lives in your dedicated space in the CampusCISO Portal: you upload files there instead of emailing them, we return deliverables there, and your named Portal Users control any sharing beyond that group. The CampusCISO Portal is hosted in the United States.

Who can touch it. The founder does the work, and your Portal Users are the people your institution designates. There is no bench of junior analysts and no one with incidental access, because as a solopreneur practice, there is no staff. Working with a single advisor doesn't create a long dependency on one person: engagements are scoped to complete in days, and memberships run on annual terms rather than multi-year commitments.

Your sessions stay private. Notes and records from your private advisory sessions are part of your institution-specific records and follow the same removal commitments as everything else. Cyber Bridge® group sessions are never recorded; they run under the Chatham House rule so peers can speak candidly.

Independently checkable. We name every provider that touches your data in our subprocessor disclosure and Data Processing Addendum, available on request, and we complete security questionnaires, including the HECVAT.

During an engagement

Used only for your deliverables. We use your institution's data solely to produce the contracted deliverables. Each engagement runs in its own dedicated workspace on a commercial AI platform, separate from other clients' work and from CampusCISO's general AI usage. We don't share knowledge across workspaces.

Never used to train AI. We never use client data to train AI models. We evaluate AI tools continuously as part of the practice, but we use only platforms that meet a fixed standard for client work: we operate every AI platform under commercial terms that bar training on client materials, so even incidental exposure during analysis does not become model memory. We name the platforms we currently use in our subprocessor disclosure, and any change appears there before it touches your data.

Deleted when the work is done. When the deliverable is finished, we transfer the final work into your space in the CampusCISO Portal and delete the dedicated workspace and any local working files used in preparing the analysis. Deleting the workspace removes your institution's data from our active systems, and our AI providers permanently delete it under their commercial terms, subject to limited retention required by law or to enforce their acceptable-use policies.

Aggregated only in cohorts. De-identified observations contribute to the benchmarking dataset, and we publish benchmark statistics only when a cohort reaches at least 30 institutions, with most cohorts at 100 or more, so no one can identify an individual institution from aggregate findings.

When the client relationship ends

A 90-day window to retrieve everything. When a client relationship ends, we keep your portal space available for up to 90 days so your team has time to download everything. We'll also delete it sooner if you ask.

Then it leaves our active systems. After that, we remove institution-specific data from our active systems. This includes assessment scores, policy inventories, year-over-year trending history, and diagnostic outputs. De-identified contributions to the higher education benchmarking dataset remain as part of the aggregate, without institutional identifiers. Former clients who return for new engagements start with a fresh baseline.

If your institution has specific AI or data governance requirements

Discuss them with us before the engagement, not after. CampusCISO's pricing depends on streamlined workflows that have been refined for consistency and quality over many engagements. Most institutional requirements are compatible with how we already work. Some require accommodation, which we can often arrange. Some are incompatible with the streamlined model and would materially change the engagement.

IN PRACTICAL TERMS

COMPATIBLE

Standard AI governance frameworks
(for example, NIST AI RMF principles, institutional AI use policies, vendor AI disclosure requirements) are typically compatible with CampusCISO's existing workflow. Bring them into the initial conversation, and we can confirm alignment.

OFTEN WORKABLE

Specific requests about platforms, providers, or controls
(for example, requirements to avoid certain AI vendors, use specific data residency arrangements, or apply enhanced logging) can sometimes be accommodated. The accommodation may affect engagement timeline or cost, depending on what the request changes about our standard process. We'll be transparent about both.

NOT A FIT

Requirements that would replace the streamlined workflow
(for example, requirements to perform analysis without AI assistance, or to use a fundamentally different platform stack) generally don't fit CampusCISO's commercial model. If your institution needs accommodations that our streamlined model can't support, we'll tell you up front, before any commitments are made, so you can find a firm whose support model fits your needs.

The principle behind all three: discuss it before, not after. A conversation at the proposal stage is much cheaper than a conflict mid-engagement and gives both sides a clear answer before commitments are made.

03

Methodology Transparency.

We publish the frameworks behind the work, rather than hold them as proprietary knowledge.

THE KEY IDEA

We publish our frameworks openly, so you can read the exact methods before working with us, check your deliverable against them, and use the methods whether you hire us or work on your own.

We decided to publish our frameworks rather than hold them as proprietary knowledge. We want higher education to succeed with this work, whether you're a client, a consulting firm building on our outputs, or a team that just needs a reference model you can cite and use internally. The transformation started with our 2026 service refresh: some frameworks are published now, others are working through their updates, and every new service going forward will publish its framework as a matter of course.

What methodology transparency means

We publish our frameworks under Creative Commons Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0). You're free to read, cite, and redistribute them with attribution. You can apply them in your own work and build tools, workflows, and services around them, including commercial work. What the license reserves is the right to publish a modified version of the framework.

Where each framework stands today. Each engagement already applies its documented methodology; publishing it makes that method available to everyone, not just our clients.

FRAMEWORK	STATUS	LICENSE
CampusCISO IT Policy Framework	Published now	CC BY-ND 4.0
Cyber Heat Map Framework	Next, with its 2026 update	CC BY-ND 4.0
CampusCISO TTX Coaching Framework	After that, with its 2027 update	CC BY-ND 4.0

Transparency is structural to the practice, not a marketing position. Once a framework is published:

- **Prospective clients** can read it in full before the proposal conversation begins.
- **Current clients** can compare the deliverable they receive to the published methodology and verify the work was done the way CampusCISO said it would be.
- **Partners** who build on CampusCISO outputs can explain to their own client exactly what method underlies the analysis they're delivering.

None of these depend on CampusCISO taking action. You can use the published methodology for reference and planning whether or not you engage with a paid CampusCISO service.

The benchmarking dataset is separate from the frameworks. We publish aggregate observations drawn from the dataset in articles and other resources. The dataset itself remains proprietary to CampusCISO.

WHY WE PUBLISH THE METHODOLOGY OPENLY

A Note From the Founder

In addition to providing transparency, helping colleges and universities succeed is personal to me. I've seen firsthand how higher education cybersecurity teams often spend too much time on audits, compliance checklists, and reactive GRC cycles while spending too little time on the continuous improvement strategies that actually build resilience. My whole career has been about simplifying seemingly complex systems: we set out to create methodologies that use the 80/20 principle to help teams remove the busywork and focus on the handful of tasks that actually move the needle. Nearly three decades supporting hundreds of institutions convinced me this is a sector-wide problem, not an institution-by-institution one.

The dataset is our core business, so we can't give it away. The methodology is different: we've decided to share our methodologies rather than lock them away for paying clients, so any institution can benefit from the work that's gone into refining these models. Our paid advisory engagements fund continued development of the frameworks and benchmarks that the whole sector benefits from.

Our free resources help institutions move from reactive compliance to strategic continuous improvement, year over year, without having to build their plan from scratch.

— Chris Schreiber, Founder, CampusCISO®

04

Our Commitments.

The operating commitments behind the Trust Anchor, and why specifics, not rhetoric, are the test.

What We Will and Will Not Do

These are the operating commitments behind the Trust Anchor. They apply to direct engagements, partner-channel engagements, and the benchmarking research that supports both:

ON HOW WE USE AI

- **We will use AI to assist with analytical groundwork.** Intake normalization, benchmarking cross-reference, draft structure, pattern recognition across the dataset.
- **We will review, interpret, and validate every AI-produced output during the work and every client deliverable before delivery.** No exceptions. No AI-only outputs reach a client.
- **We will not present AI outputs as expert analysis.** What we deliver is expert analysis informed by AI groundwork, not the reverse.

ON HOW WE HANDLE YOUR DATA

- **We will delete each engagement's dedicated workspace when the work is complete.** Each engagement runs in its own dedicated workspace on a commercial AI platform, separate from other clients' work and from CampusCISO's general AI usage. We transfer the finished work into the client's space in the CampusCISO Portal and delete the dedicated workspace and any local working files when the deliverable is complete. When a client engagement ends, we delete institution-specific data from our active systems, and our AI providers permanently delete it under their commercial terms, subject to limited retention required by law or to enforce their acceptable-use policies.
- **We will review and score data before it enters the benchmark database.** This applies to de-identified contributions from completed engagements and to the public data analysis we compile while watching for sector trends. No automated data ingestion occurs without expert qualification.
- **We will pool benchmark data into segments of at least 30 institutions before publishing any benchmark statistic.** Most benchmark cohorts include 100 or more institutions. This pooling threshold is a structural privacy commitment that makes reverse identification impractical even in higher education's small, distinctive cohort, where stripping institutional names alone would not reach the same protection.
- **We will not use client data to train AI models.** Not now, not later, not as a condition of any future pricing tier.

ON HOW WE SHARE OUR METHODOLOGY

- **We will publish the methodology behind our engagements under Creative Commons Attribution-NoDerivatives 4.0 (CC BY-ND 4.0).** The frameworks are available to read, cite, and redistribute with attribution. Any institution can apply them whether or not they engage CampusCISO.
- **We will not gatekeep the methodology behind paywalls or non-commercial use restrictions.** A free community edition of every CampusCISO framework will be available to everyone. A prospective client can read the exact method before deciding to engage, can compare deliverables against the published methodology during the work, and can use the published methodology after the engagement ends.

Why This Is Structural, Not Cosmetic

Higher education has good reason to be skeptical of AI in advisory work. Vendors are marketing AI capabilities aggressively, and the substance behind those capabilities is rarely available for scrutiny. Institutions evaluating advisors are often left guessing on the three questions that matter most: how the work is actually done, what happens to their data, and what method underlies the analysis.

Our position is that the difference between a structural commitment and a cosmetic claim is shown in specifics, not rhetoric. Here is where the difference actually shows up:

- **AI use.** Who reviews each deliverable, what they bring to the review, where in the workflow the review happens, and whether the review can be skipped without anyone noticing.
- **Client data handling.** Whether the protections are built into how the practice operates or exist only as a policy that can be revised, and whether the vendor names its storage, its providers, and its data lifecycle, or hides them behind "secure systems." We name ours.
- **Methodology.** A test anyone can run: the published frameworks let a client compare what was promised to what was delivered.

We've named all three so they can be examined and held to.

That's our commitment, and it is sustainable because we built the practice on it rather than retrofitting to it.

CampusCISO®

IN CLOSING

Bring your questions.

We would rather answer them in the proposal conversation than have them surface during a procurement review or mid-project.

Explore the methodology

Our published frameworks show exactly how we work, so you know what to expect before engaging. Free community editions are available to everyone.

campusciso.com/guides

Start a conversation

Discuss any institutional AI or data governance requirements before the engagement, not after. A conversation at the proposal stage is much cheaper than a conflict mid-engagement.

campusciso.com/free-consultation