



HITRUST

From Audit Sprint to
Assurance Engine




Table of Contents



01.

Genesis, Evolution, and HITRUST in the Assurance Landscape

02.

Aligned with Changing Times

03.

AI Assurance: Business Value and How HITRUST Differs from ISO 42001

04.

Treating HITRUST as an Operational Model, Not a One-Time Project

Genesis, Evolution, and HITRUST in the Assurance Landscape

01

Genesis

HITRUST was founded in 2007 by a group of healthcare, technology, and information security leaders to serve as a bridge between what HIPAA required organizations to protect and how to protect it. Every covered entity and business associate interpreted “reasonable and appropriate safeguards” differently, and every customer audit looked different from the last. Over the following two years, HITRUST built and released the HITRUST CSF. Originally the Common Security Framework, now simply the CSF is a single certifiable control set that harmonized existing standards including ISO 27001, NIST 800-53, HIPAA, PCI DSS, and COBIT.

The CSF Assurance Program turned that framework into something an organization could actually be certified against, followed shortly after. The primary aim was harmonization rather than novelty, though HITRUST did add its own layer of prescriptive control specifications, tiered implementation levels, and a maturity based scoring model that none of the source frameworks provided.

Evolution

What started as a healthcare specific framework steadily broadened. Over successive CSF versions, HITRUST expanded its control library, refined its scoring methodology (moving toward maturity-based scoring rather than simple pass/fail), and introduced tiered assessments: e1 (essentials, low assurance, one-year), i1 (implemented, moderate assurance), and r2 (risk-based, the most rigorous, two-year certification with interim assessments). This tiering matters commercially. It let HITRUST move to smaller vendors who couldn't accommodate an r2 engagement.

Evolving from a healthcare specific framework, HITRUST has now been adopted across multiple industries. Over successive CSF versions, HITRUST widened its authoritative source coverage. The framework now incorporates more than 60 frameworks, regulations, and standards, while streamlining requirement counts to reduce redundancy.



Integration with Other Regulatory Factors

Part of HITRUST's design philosophy is that it was never meant to stand alone as one more thing to comply with; it was built to sit on top of, and interoperate with, whatever else an organization already has to address.

On the regulatory side, the CSF's original authoritative sources: HIPAA, PCI DSS, NIST, ISO 27001, meant an organization working toward HITRUST is simultaneously building evidence toward those underlying requirements rather than treating each as a separate project. The overlap is in evidence, not equivalence: a HITRUST certification is not a HIPAA compliance determination and not an ISO 27001 certificate, and it shouldn't be represented as one. As data privacy regulation expanded across states, GDPR for organizations with international exposure, sector-specific rules in financial services, HITRUST extended its mappings to absorb those requirements too, so a single control implementation can serve multiple regulatory obligations at once.

The same logic extends to industry assurance reports. Customers weren't going to stop asking for SOC 2 just because a vendor had HITRUST aligned controls, so the AICPA and HITRUST jointly developed a combined program, SOC 2 + HITRUST, that mapped the Trust Services Criteria to HITRUST CSF requirements and allowed an organization to report on both in a single report. In the combined examination, the CSF's roughly 150 control requirements are mapped to the applicable Trust Services Criteria. However, it is important to note that a SOC 2 + HITRUST certified report can be issued by an entity which is both a HITRUST External Assessor and a CPA firm.



AI Integration

The newest chapter is AI. HITRUST launched a dedicated AI Security Assessment and Certification which layers up to 44 additional, highly prescriptive controls onto an existing e1, i1, or r2 assessment. It isn't a standalone credential. It currently isn't intended for organizations that only use AI tools internally or build workflows on third-party models, and it must be paired with a concurrent or existing HITRUST assessment. It was built to address the model weights, training data, and inference layer specifically, not just the infrastructure sitting around them, and it was designed to align with AI security elements already present in ISO, NIST, and OWASP guidance. HITRUST has also signaled that CSF updates targeted for 2026 will widen coverage to internal AI usage and workflow integrations, which tells you this isn't a side project. It's where the framework is structurally heading.

Aligned with Changing Times

02

HITRUST's durability comes from a decision the organization made early and has stuck to: don't reinvent controls, harmonize them. As new regulatory pressures appear (state privacy laws, evolving HIPAA guidance, cloud-native architectures, third-party risk mandates), HITRUST absorbs the relevant requirements into the CSF rather than asking organizations to run parallel compliance tracks.

Few concrete signs of that adaptation:



Cloud native reality

Certification now leans heavily on inherited controls from cloud service providers. They enable infrastructure control inheritance, which cuts assessment scope and hours for vendors built on those platforms, but it also means pre-audit shared responsibility mapping is no longer optional; getting it wrong is one of the more common causes of timeline slippage.



Tiered accessibility

The e1/i1/r2 structure exists specifically because "changing times" meant smaller vendors and startups needed a credible entry point into assurance without absorbing r2 level cost and effort. This kept HITRUST relevant as the buyer base broadened beyond large health systems and payers.



Continuous assurance, not a point-in-time badge

The market has shifted from "Show me a certificate" to "Show me you're still controlled today." HITRUST has responded by pushing toward control maturity scoring, interim assessments between full recertification cycles, and integration with continuous control monitoring tools, moving certification from a periodic snapshot toward something closer to a live posture indicator.

Threat-adaptive controls

Rather than a static control list revised every few years, HITRUST has moved toward a threat-adaptive model, updating requirements in response to the current threat landscape rather than waiting for the next major CSF version.

Cost and procurement dynamics

For US organizations in 2026, certification cost still ranges widely depending on tier and scope, and in regulated healthcare specifically, r2 frequently carries more procurement weight than a standalone SOC 2 report because of its prescriptive healthcare mapping, a good example of HITRUST's positioning strengthening rather than fading as newer frameworks (like ISO 42001) enter the conversation.

The throughline is that HITRUST doesn't try to be first to market with a new compliance concept. It tries to be the framework that absorbs whatever the market has already decided matters, and packages it into something assessable and defensible. Each adaptation is another piece of what turns a one-time sprint into a running engine.



AI Assurance: Business Value and How HITRUST Differs from ISO 42001

03

Why this matters for a company's AI

If your product embeds a model, a generative AI feature, a predictive engine or an agentic workflow, customers and regulators are now asking a question security teams weren't built to answer five years ago: is the model itself secure, not just the infrastructure around it?

Training data poisoning, model extraction, prompt injection, data leakage through inference, and supply-chain risk from third-party foundation models are real attack surfaces that traditional infrastructure-focused audits don't touch.

How HITRUST's AI certification helps

It gives vendors a certifiable, evidence-based answer to third-party risk questionnaires that increasingly include AI-specific sections.



It shortens sales cycles with regulated buyers. One healthcare AI company described the practical effect directly: telling a prospective customer they hold HITRUST certification helps accelerate the sales conversation, and the certification shifted internally from a cost center to something that actively supports go-to-market.



It's modular. It layers onto an assessment a company may already be running (e1/i1/r2), so it doesn't require standing up an entirely separate audit program.



It's prescriptive rather than principle based. How it's different from ISO 42001?



Both are certifiable through third-party audit, so the distinction isn't assurance versus no assurance, it's what gets certified. ISO/IEC 42001 certifies a management system: do you have AI governance, accountability structures, and oversight processes, and do they operate?

HITRUST certifies technical implementation: can you prove, with tested evidence, that the specific security controls protecting your deployed AI system are in place and working? HITRUST frames these as complementary rather than competing. Its own specification calls the certification a complement to, not a replacement for, ISO/IEC 42001:2023.

Treating HITRUST as an Operational Model, Not a One-Time Project

04

The single biggest oversight that organizations make with HITRUST is treating certification as a finish line. It isn't. The certificate is a snapshot of a control environment that has to keep operating effectively for the certification to mean anything, and HITRUST's own direction (interim assessments, continuous monitoring integration, maturity scoring) reflects that the industry has moved past the "get the badge and file it away" approach. This is the point where the audit sprint has to give way to the assurance engine.

The typical HITRUST journey

Readiness / Gap Assessment

Before any formal engagement, understand where current controls stand against the applicable CSF requirements. Skipping this is the most common reason timelines blow up.

Scoping

Determine assessment tier (e1/i1/r2), factor systems, and whether an AI Security add-on applies. Scope drives cost and effort more than almost anything else, so getting this wrong early is expensive to fix later.

Control Implementation and Evidence Building

Controls generally need to have been operating for a minimum period before fieldwork; HITRUST guidance requires implemented controls to run for at least 90 consecutive days before fieldwork, and newly implemented or remediated policies/procedures need at least 60 days. This is where organizations most often run out of runway because they started evidence collection too late.

Self-scoring and Validated Assessment

The organization first scores itself. Followed by an authorized external assessor who assesses the evidence; they don't implement or remediate controls themselves, only evaluate whether evidence substantiates the requirement. Field work must be completed within 90 days.

HITRUST QA Review

Results go to HITRUST for quality assurance and corrective action plan resolution before anything is issued. This stage typically adds one to two months on top of the assessment, within a journey that commonly runs about twelve months end to end.

Certification and Interim Assessments

r2 runs a two-year cycle with an interim assessment at year one; e1 recertifies annually through a full validated assessment, and i1 offers a rapid recertification option. Either way, maintenance is a cycle, not a single pass/fail moment.

Recertification

The cycle repeats, ideally with far less friction if the organization has kept evidence collection continuous rather than reactive.

Why "operational model" is the right frame

Certification without continuous monitoring degrades into a static artifact, accurate on audit day, increasingly unreliable every day after. Organizations that get the most value integrate control monitoring, configuration drift detection, and automated evidence collection into normal operations, so that by the time the next assessment window opens, most of the evidence already exists. Some mature programs go further and tie HITRUST control maturity scoring into their broader risk models, using improved control effectiveness to justify a lower modeled likelihood of a breach event, which turns the certification into an input for risk quantification rather than a compliance checkbox.

Common pitfalls to avoid



Starting evidence collection too late relative to the 60/90-day operating windows. This alone causes more delays than any technical control gap.



Scoping too broadly or too narrowly. Over-scoping wastes budget; under-scoping creates rework when the assessor pushes back.



Treating cloud inheritance as automatic. Inherited controls from cloud service providers still require the organization to map shared responsibility accurately; assuming the provider "covers it" is a frequent source of last-minute gaps.



Siloing the AI assessment from the core security program. Since the AI add-on is not standalone, it needs to be planned alongside, not after, the base e1/i1/r2 engagement.



Treating recertification as a repeat of year one. If continuous monitoring wasn't built into operations, every renewal cycle recreates the same scramble the first certification did.



Letting certification become a compliance-team-only initiative. The organizations that get commercial value out of HITRUST (faster sales cycles, procurement trust) are the ones where security, product, and sales all understand what the certification actually signals to a buyer.

The organizations that get durable value from HITRUST are the ones that stop asking "When is the audit" and start asking "What does our control environment look like on any given day." That shift, from project to operating discipline, from sprint to engine, is what separates a certificate on the wall from an assurance program that actually protects the business.



ABOUT US

Pierag Consulting was founded in February 2021 by Abhishek Gupta, Thomas Raffa and Pierian Services as a unique business model to serve clients globally by blending domestic proficiency with global expertise. Since then, we have been serving prominent clients across the US, SEA, India and UK in the field of Accounting Advisory, Business Risk, Technology Risk, Assurance, ESG & Sustainability, Deals Advisory, Tax and Digital.

With highly skilled team members, we are fueled by our purpose of 'Inspiring people to do things that inspires them' and our values of 'Excellence, Equity & Empathy'.

Copyright @ 2026, Pierag Consulting (operating under brand name 'Pierag'). All rights reserved. This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax or other professional advice. Please refer to your advisors for specific advice. Pierag expressly disclaims all liabilities in respect to actions taken or not taken based on any or all the contents of this document.

Connect with us



Gaurav Khandelwal

Partner - Risk Advisory Leader

✉ gaurav.khandelwal@pierag.com



Dipesh Khushalani

Technology Risk Advisory Leader

✉ dipesh.khushalani@pierag.com



Karan Trehan

Associate Director- Technology Risk Advisory

✉ karan.trehan@pierag.com



More information

www.pierag.com | info@pierag.com



[pierag/socialmedia](https://www.linkedin.com/company/pierag)



OUR LOCATIONS

Gurugram | Mumbai | Bengaluru | Jaipur | Chandigarh | Washington, D.C. | Melbourne | Singapore