

VENDOR PROFILE

# TripleBlind: A Lateral Approach to Privacy-Enhanced Data Sharing

**TripleBlind Is Shifting the Paradigm in Privacy-Enhancing  
Technology for Research Data Sharing**



**Steve Wilson**  
VP and Principal Analyst

# TABLE OF CONTENTS

Executive Summary..... 3

Privacy Technology Matures ..... 4

The State of the State in Privacy-Enhancing Technologies ..... 6

The TripleBlind Solution ..... 8

Key Differentiators..... 8

Endnotes ..... 11

About Constellation Research..... 13



## EXECUTIVE SUMMARY

TripleBlind is a leader in the relatively new and fast-evolving category of privacy-enhancing computation (PEC). A great many options are emerging for unlocking the inherent value of data in decision-making while maintaining privacy and control over data locality and visibility. Some use new mathematical techniques such as homomorphic encryption, statistical perturbation, or pseudonymization to enable third-party analytics without exposing personal details (although these methods are complex and fragile), whereas others copy data to an intermediate clean room for hosted analytics (a comforting and transparent approach but one that challenges data localization policies).



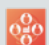
In contrast, with the TripleBlind APIs and user interface, customers have access to secure multiparty computation (SMPC) and other advanced privacy tools. TripleBlind claims that its SMPC includes proprietary improvements over other SMPC options—including TripleBlind’s own computationally superior Blind Learning (a patented solution for distributed, privacy-first, regulatory-compliant machine learning at scale), Advanced Encryption Standard (AES) inference, distributed inference, and distributed regression techniques—that make it suitable for different use cases. All data remains localized within the customers’ own networks, within which they choose the analytics to run. Customers retain full control over their own security posture and their own standards of care for their valuable data assets. TripleBlind’s code has undergone rigorous independent evaluation by the MITRE Corporation and has proven its high technology readiness.

In the busy data protection market, with many new algorithms in play against regulatory tensions, Constellation believes that strong R&D and rigorous validation are at a premium. One mistake in execution or the math, and customer privacy can be irreparably damaged. Yet the potential rewards for safely unlocking data are enormous. TripleBlind’s award-winning research-driven solutions may help shift the paradigm in a highly risk-averse environment by enabling organizations to share analytic outcomes instead of sharing data.

### TripleBlind

- **Headquarters:** Kansas City, MO
- **Founded:** 2019
- **Type:** Privately held
- **2020 Revenue:** Unknown
- **No. of Employees:** 35
- **Website:** <https://tripleblind.ai/>
- **Twitter:** @TripleBlindAI

### Business Themes

-  Marketing Transformation
-  Data to Decisions
-  Digital Safety & Privacy

## PRIVACY TECHNOLOGY MATURES

Data privacy—or simply *data protection* in many jurisdictions such as the European Union—generally has been more about policy and rules than about technology. In Constellation’s view, privacy is fundamentally about restraint; it is less about what organizations do with personal data than it is about what they choose not to do. Looking at privacy this way puts the stress on business practices rather than technology.

There has long been a category of tools known as *privacy-enhancing technologies* (PETs), with a focus on encryption to hide personal data. Yet encryption is not sufficient to ensure privacy, nor, paradoxically, is it even necessary. This is because privacy is not about secrecy but, rather, about *control*. Data privacy rules apply to *identifiable* (personal) information; the rules that restrain the collection, use, and disclosure of personal information apply to a class of data that is not secret.

In the past 10 years or so, the focus in data privacy has shifted to questions that are more regulatory in nature, such as geographical localization (hence the conspicuous rush to build cloud data centers all over the world) and reasonable grounds for data processing (as opposed to old-fashioned blanket consent). The United States over this time frame has started to adopt more internationally contemporary principles-based privacy laws, superseding the prescriptive notice-and-consent approach of past laws such as the Children’s Online Privacy Protection Act (COPPA, 2000) and the Health Insurance Portability and Accountability Act (HIPAA, 1996). And leading regulators have stressed precautionary approaches. For instance, under Europe’s infamous but widely misunderstood General Data Protection Regulation (GDPR), pseudonymized data should still be treated as personal data.<sup>1</sup> Constellation advises its clients to not rely on encryption alone but, rather, to deploy layered privacy protections to control the flow of any potentially reidentifiable encrypted files.

Therefore, any effective privacy solutions must deploy technologies in concert with processes and rules. Good privacy-by-design (PbD) practices should be grounded in a privacy impact assessment (PIA—aka data protection impact assessment, or DPIA) to gauge the material risks to privacy in the business context and help guide an optimum and well-understood balance of technologies and processes.

## TripleBlind Takes a Lateral Approach

TripleBlind has taken a contemporary lateral approach to privacy, developing a software-only product featuring SMPC and other advanced privacy tools. TripleBlind claims that its SMPC includes proprietary improvements over other SMPC techniques, suited to different use cases.

One such technique is TripleBlind's Blind Learning, an alternative to federated machine learning. TripleBlind asserts that its Blind Learning ensures that counterparties never expose their partially trained models to each other, has a lower computational load, and trains faster than comparable solutions. For other use cases, TripleBlind deploys other PEC methodologies, such as AES inference, distributed inference, and distributed regression.

TripleBlind is deployed within the customer's network to undertake localized data analytics. The raw data never leaves the customer's site, and it is one-way-encrypted within the analytics space for SMPC or whichever other technique is employed. Algorithmic details of any contributor are also kept confidential. No hardware or network modifications are involved.

TripleBlind offers the following benefits:

- No preprocessing, perturbation, or degradation of data; no restrictions on the applicability of analytics, so customers can run any analysis they like
- Data remaining local to the customer, without the need for cross-border flows
- No reliance on assumptions of deidentification or on the quality of sometimes-novel algorithms
- Elegant informatic architecture, more easily proven to be mathematically correct than most cryptographic algorithms (especially novel homomorphic encryption)
- Customers remaining subject to the legal jurisdiction of their choosing and keeping full control of the security model

To appreciate the properties and power of TripleBlind's approach, it's helpful to review the world of privacy-enhancing technologies.

# THE STATE OF THE STATE IN PRIVACY-ENHANCING TECHNOLOGIES

The prevailing data privacy regulatory model is concerned with *identifiable data*—namely, any data that can reasonably be associated with an identifiable natural person, now or in the future, on its own or in combination with other data. The prevailing regulatory term of art internationally is *personal data* or *personal information*. Classical data privacy rules (traceable to the Organization for Economic Co-operation and Development’s foundational privacy work of 1980<sup>2</sup>) do not apply to truly anonymous data, so data scientists, privacy champions, and regulators have all become fixated on identifiability. To maximize privacy, the classical privacy protection is to minimize identifiability.

So the focus of PETs as a class has been on encryption and/or deidentification. These tools make data less usable (indeed, that is really the whole point of encryption), but there are alternatives.

## Deidentification

Deidentification, pseudonymization, statistical data linkage keys, and the like create new indexes for individual data records, seeking to join all records about the same person to a new index or pseudonym without identifying the person. The trouble—especially these days when so much extra data is available—is that additional data can be found and folded into a pseudonymized dataset to make the records recognizable again.

## Homomorphic Encryption

Regular encryption scrambles data to make it unrecognizable, which hides the data from adversaries but also puts it beyond the reach of analytics. However, newer homomorphic encryption does not scramble data to the extent that it cannot be processed. Fully homomorphic encryption (FHE) is a cutting-edge technique that, in theory, preserves all mathematical properties of a dataset. Researchers, for example, can still run all the statistics they like on a set of health records. A fundamental problem is that, logically, if there is some morphology left in the encrypted data, that also leaves structures and clues for attackers to exploit for illicit reidentification. These algorithms are so new that their resistance to reidentification attacks is not yet properly understood.

There are also practical problems of performance degradation: Any homomorphic encrypted data brings performance overhead, compared with processing unencrypted data, when it is being processed. This is still an active area of academic research. On the one hand, the overheads might continue to be improved to become manageable, but on the other hand, this math is new, and thus the security promises of FHE have not yet been demonstrated to the levels of confidence that established encryption algorithms have earned.

## Differential Privacy

Differential privacy hides identifying details by perturbing data—injecting noise—in a way that purportedly preserves statistical significance so that mathematical analyses are still meaningful. Nevertheless, some damage is necessarily done in the process, and the extent to which analytic results suffer is difficult to gauge. Debates continue in the data science community about the efficacy of differential privacy for protecting privacy, because it requires such a light touch.

## Data Clean Room

Deidentification, FHE, and differential privacy all trade off the quality of data or its utility for analytics against privacy, hence tending to limit the research possibilities. And so it follows that research data sharing and open data programs in general are widely regarded as necessitating some sort of compromise. Policymakers are forced to argue for a public-good outcome at the cost of personal privacy. The arguments are inherently controversial and almost always loaded.

One alternative is the data clean room, which secretes data in third-party or neutral-party quarantine areas for analysis in much the same way that a mergers-and-acquisitions deal room is often set up at a legal or accounting practice to host sensitive documents. A key advantage is simplicity: No complex algorithms and no technically complicated accuracy trade-offs are involved. Remember that complexity is the enemy of security and privacy too. Data clean rooms are an easily understood proposition. On the other hand, even if well managed by careful and reputable operators, they do involve moving data between locations and jurisdictions, which is frowned upon in most contemporary data privacy regimes.

# THE TRIPLEBLIND SOLUTION

TripleBlind researches and develops a suite of novel PEC techniques and has brought several well-thought-through solutions to market. This profile focuses on TripleBlind's Blind Data and Algorithm API, a patented solution that employs SMPC and other data privacy methodologies and includes related capabilities such as digital rights management, asset registering, audit trails, and containerized access points.

With TripleBlind's solution, raw data is never moved from the customer to another jurisdiction, nor is it exposed outside the data owner's own network. The data owners retain full control of their security posture and data protection standards. As a result of SMPC and the other data privacy techniques discussed, the TripleBlind solution reduces risk, cost, and oversight effort without reducing data utility or the depth and breadth of the analytics that are possible. In theory, the full analytic value of data can be unlocked, because with no preprocessing needed and no change of custodianship, no limits are placed on the analytics that can be run. The data owners control their own research agenda. Furthermore, TripleBlind never sees any of the customer's data. The solution is not a managed service of TripleBlind's (so, strictly speaking, the company lies outside the data-protection-as-a-service category) but instead is a self-contained software-only solution deployed within the customer's IT environment (either behind the customer's firewall or within a hosted environment of the customer's choosing).

## KEY DIFFERENTIATORS

### Elegant Privacy-Enhancing Computation

With its Blind Data and Algorithm API, TripleBlind makes an elegant and easily verified privacy promise. The architecture is simple, and there is no complicated cryptology to be proven in order to have confidence in the system and no interference with any of the raw data, unlike in the case of homomorphic encryption or differential privacy.



## Independent Evaluation

TripleBlind works continuously with independent advisers and peer reviewers, including the MITRE Corporation, Accenture, the University of New Brunswick in Canada, and the law firm Polsinelli, to validate its work.

MITRE Engenuity is a wholly owned nonprofit subsidiary of MITRE Corporation, founded to convene impartial, noncompetitive collaborations of experts, organizations, and investors in innovation for the public good. Among other things, MITRE Engenuity provides independent evaluation of cybertechnologies. A prospective customer of TripleBlind commissioned MITRE Engenuity to validate TripleBlind's capabilities and claims.

A convenience sample of data privacy technologies was pulled from MITRE Engenuity's recent routine research, and several synthetic use cases were derived to approximate the representative types of studies seen in the fields of observational research and clinical trials between 2020 and 2021. From the synthetic use cases, MITRE Engenuity built a list of features it judged necessary for a research network to effectively address each use case. It then compared the capabilities identified in its review with the requirements identified from the synthetic use cases. Separate MITRE Engenuity staff members not involved in the initial identification of features reviewed the requirements before the entire team examined how the TripleBlind solution matches those requirements.

Of the MITRE Engenuity-derived requirements, four are mandatory, five are desirable, and two are business-related metrics that gauge technology readiness and the cost of deployment. The details of the requirements and the tested product compliance are confidential, but Constellation has confirmed that TripleBlind met all four mandatory requirements, met three of the desired requirements, and partially met two of the desired requirements. TripleBlind also scored well on the business metrics, being judged to be fully operational, as well as low costs in terms of operations and maintenance.

The TripleBlind code has undergone rigorous independent evaluation as well, proving both the fundamental mathematics and its high technology readiness. And the company has impressive reference implementations with prestigious institutions such as the Mayo Clinic.

## Research and Development

Major intellectual horsepower is needed to drive and guide a business such as this. The TripleBlind board is rich with senior researchers, healthcare researchers, academics, and deep experts in artificial intelligence (AI) and data science.

The TripleBlind R&D team has amassed an important body of published academic papers and patent filings. As further evidence of the high quality of the company's research, TripleBlind Chief Technology Officer Dr. Craig Gentry was recently jointly awarded the prestigious 2022 Gödel Prize for his original work on FHE.<sup>3</sup>

## Conclusion

Good-quality data is becoming as important to society as clean drinking water. A whole range of commensurate digital safety measures is emerging, and new categories of data protection based on evolved concepts of privacy and data access are forming. Society at large now expects to share the benefits of data as a communal resource without actually sharing the data.

Privacy-enhancing computation is still taking shape as a category within the broad field of privacy-enhancing technologies and data protection as a service. TripleBlind is poised to be a leader and could shift expectations of privacy in data analytics. Maximum privacy does not mean minimum identification if the analytics can be safely carried out locally. And data owners do not need to share, relocate, or lose control of their precious assets if the analytics can come to them.

## ENDNOTES

---

<sup>1</sup> “General Data Protection Regulation Recital 26,” Intersoft Consulting, April 27, 2016. <https://gdpr-info.eu>

---

<sup>2</sup> “Recommendation of the Council Concerning Guidelines Governing the Protection of Privacy and Transborder Flows of Personal Data,” Organization for Economic Co-operation and Development (OECD), 1980; amended 2013. <https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0188>

---

<sup>3</sup> 2022 Gödel Prize announcement, European Association for Theoretical Computer Science and the Association for Computing Machinery, June 1, 2022. <https://sigact.org/prizes/g%C3%B6del/citation2022.html>

## ANALYST BIO

# Steve Wilson

Vice President and Principal Analyst

Steve Wilson is a vice president and principal analyst at Constellation Research and leads the firm's work in digital safety and privacy. A 20-year veteran in cybersecurity, Wilson is one of the world's most original thinkers in digital identity.

Wilson is a researcher, innovator, and R&D leader with 30 years of experience in information technology. Since 1995 he has been dedicated to digital identity and privacy, responsible for numerous breakthroughs in smart technologies, identity management, privacy-enhancing technologies, and national identity frameworks. Wilson has been awarded nine cybersecurity patents and is currently undertaking a Ph.D. on the evolution of identity ecosystems.

Wilson advises chief information security officers, chief privacy officers, strategists, and ICT architects seeking to optimize data protection in complex digital systems. He provides Privacy Impact Assessments, builds robust security strategies, and helps architect identity for big data, Internet of Things, and cloud rollouts. His coverage areas include digital safety and privacy, data to decisions, and consumerization of IT.

---

 [@Steve\\_Lockstep](https://twitter.com/Steve_Lockstep)  [constellationr.com/users/steve-wilson](https://constellationr.com/users/steve-wilson)  [linkedin.com/in/lockstep](https://linkedin.com/in/lockstep)

# ABOUT CONSTELLATION RESEARCH

Constellation Research is an award-winning, Silicon Valley—based research and advisory firm that helps organizations navigate the challenges of digital disruption through business model transformation and the judicious application of disruptive technologies. Unlike the legacy analyst firms, Constellation Research is disrupting how research is accessed, what topics are covered, and how clients can partner with a research firm to achieve success. Over 350 clients have joined from an ecosystem of buyers, partners, solution providers, C-suite, boards of directors, and vendor clients. Our mission is to identify, validate, and share insights with our clients.

## Organizational Highlights

- Named Institute of Industry Analyst Relations (IIAR) New Analyst Firm of the Year in 2011 and #1 Independent Analyst Firm for 2014 and 2015.
- Experienced research team with an average of 25 years of practitioner, management, and industry experience.
- Organizers of the Constellation Connected Enterprise—an innovation summit and best practices knowledge-sharing retreat for business leaders.
- Founders of Constellation Executive Network, a membership organization for digital leaders seeking to learn from market leaders and fast followers.

---

 [www.ConstellationR.com](http://www.ConstellationR.com)

 [@ConstellationR](https://twitter.com/ConstellationR)

 [info@ConstellationR.com](mailto:info@ConstellationR.com)

 [sales@ConstellationR.com](mailto:sales@ConstellationR.com)

---

Unauthorized reproduction or distribution in whole or in part in any form, including photocopying, faxing, image scanning, emailing, digitization, or making available for electronic downloading is prohibited without written permission from Constellation Research Inc. Prior to photocopying, scanning, and digitizing items for internal or personal use, please contact Constellation Research Inc. All trade names, trademarks, or registered trademarks are trade names, trademarks, or registered trademarks of their respective owners.

Information contained in this publication has been compiled from sources believed to be reliable, but the accuracy of this information is not guaranteed. Constellation Research Inc. disclaims all warranties and conditions with regard to the content, express or implied, including warranties of merchantability and fitness for a particular purpose, nor assumes any legal liability for the accuracy, completeness, or usefulness of any information contained herein. Any reference to a commercial product, process, or service does not imply or constitute an endorsement of the same by Constellation Research Inc.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold or distributed with the understanding that Constellation Research Inc. is not engaged in rendering legal, accounting, or other professional services. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. Constellation Research Inc. assumes no liability for how this information is used or applied nor makes any express warranties on outcomes. (Modified from the Declaration of Principles jointly adopted by the American Bar Association and a committee of publishers and associations.)

Your trust is important to us, and as such, we believe in being open and transparent about our financial relationships. With our clients' permission, we publish their names on our website.

San Francisco Bay Area | Boston | Colorado Springs | Denver | Ft. Lauderdale | New York Metro  
Northern Virginia | Portland | Pune | San Diego | Sydney | Washington, D.C.