The Virtru Trusted Data Platform:

Accelerating Health Analytics and Collaboration

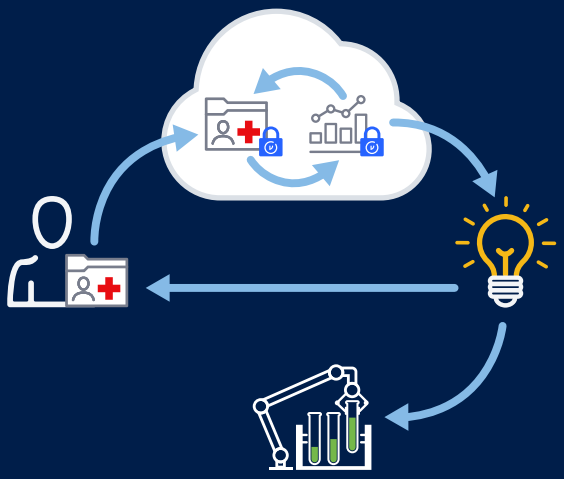

Unmet Expectations for Health Analytics

Enabling data owners to maintain full lifecycle control over health data and securely share it for approved analysis would unlock valuable innovation and deliver faster outcomes in medical research, disaster planning, personalized medicine, diagnostic error reduction, pharmaceutical research, and other critical areas.

In order to generate meaningful insights from their highest-impact analytic models, health institutions must have access to massive amounts of data that can only be gathered by sharing information with other partner organizations. Unfortunately, even though they understand the benefits of broader information exchanges, organizations are usually reluctant to share most of their valuable data with analytics collaborators due to a lack of transparency into, control over, and ownership of their data once it has been transmitted.

Health analytics need a data-centric platform based on cryptographic proof instead of trust – a platform that allows any willing parties to easily collaborate with each other without any party losing control and visibility. In order to protect institution and patient data from misuse, information sharing platforms must employ strong analytic identities and cryptographically-assured originator control over the questions that can be asked using their data. Via the Virtru Trusted Data Platform (TDP), we propose a solution to the healthcare analytics problem that leverages cryptography to mediate data access, usage, and audit through the entire analysis lifecycle.

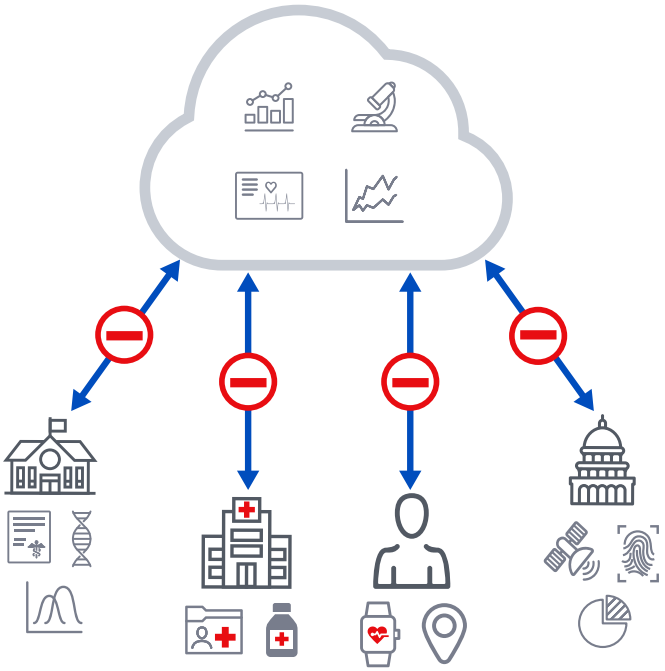

Where Today's Trust Models Fall Short

Many healthcare analytic models try to solve these challenges by establishing multilateral sharing agreements that assign one institution as the trusted party responsible for processing data inputs. These agreements seldom include technical enforcement mechanisms and often include unreliable de-identification methodologies, so they suffer from the same inherent weaknesses of most trust-based models.

Another common solution is to introduce protected enclaves for each output and place access controls at the enclave boundary, so that only parties with access to the enclave can see the actual output. However, this solution siphons off subsets of data from the aggregate, inhibiting the transparency and collaboration that trusted analytics were developed to increase.

Rather than assume the risks and liabilities that come with both of these approaches, most additional parties simply choose not to participate in collaboration or exclude critical data, which means this same information gets left out of the analyses themselves.

For instance, a hospital may want to contribute to disaster planning, but it may be unwilling to share data about bed capacity with the government, fearing it will be used against them in ways outside the scope of the disaster planning agreement. The risk of misuse of data increases the friction of sharing, which blocks even vital and agreed-upon questions from being accurately answered. This friction causes significant problems for critical emergent research, such as the abandonment of COVID-19 surveillance methods due to privacy concerns.

Introducing the Virtru Trusted Data Platform (TDP)

Health organizations need an analytic platform based on cryptographic proof instead of trust. They need a model that empowers data owners and analytics owners to share insights, while also maintaining awareness and control over the ongoing access and usage of outputs.

The Virtru TDP Leverages Cryptography for Secure Health Analytics

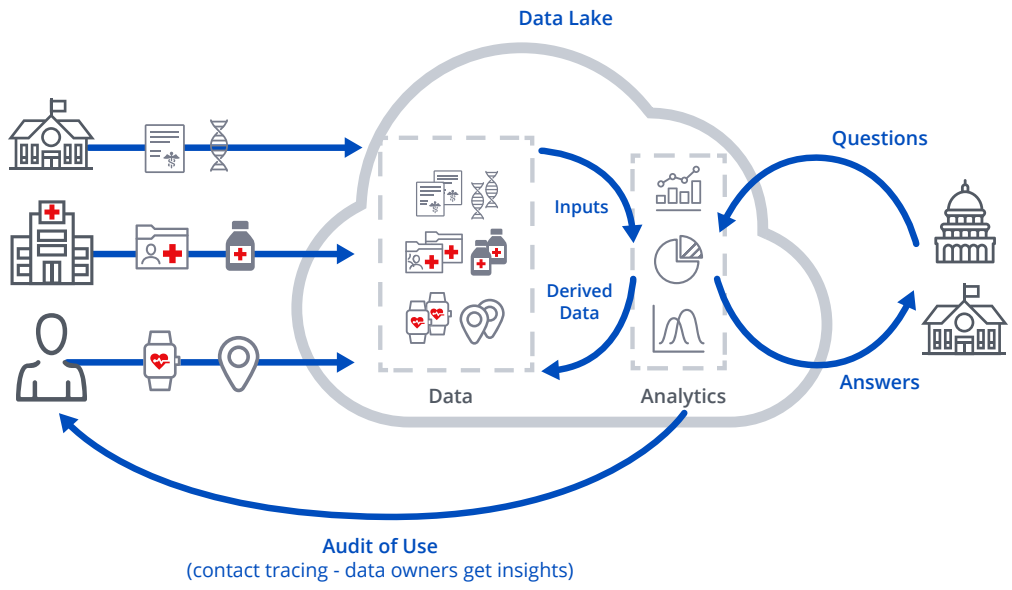

The Virtru TDP delivers on the unmet potential for health analytics by leveraging cryptography to mediate data access, analysis, and audit through the entire shared collaboration lifecycle. By assigning strong analytic identities and cryptographically assuring that original owners always control how analytics can process their data, the Virtru TDP protects health institution and patient data from their biggest risks for misuse.

By ensuring that data remains protected wherever it is transmitted or stored, the TDP renders traditional "data in transit" or "data at rest" protections part of a nice-to-have defense in depth, rather than a fragile imperative. Data owners gain full lifecycle control over their health information while still being able to safely share it for approved analysis, leading to faster, more valuable healthcare outcomes and innovations.

A Closer Look at the Virtru TDP

The Trusted Data Format

The Virtru TDP is powered by the Trusted Data Format (TDF), an open standard, JSON-encoded data format that keeps data protected and under its originator's control no matter where it is created or shared. With the TDF, an organization can apply discrete policies and rules pertaining to Attribute-Based Access Control (ABAC) that travel with the content. By enabling data owners to tie security tags to specific data objects and update these policies over time, the TDF prepares data for ingestion into the TDP.

Local client-side applications, built via the TDF software development kit (SDK), categorize and protect each data object by wrapping it with the TDF, which contains encrypted data that is cryptographically bound to any already applied tags and associated access or handling policies. The local clients then send the encrypted, tagged data to the TDP for processing.

Features of the Virtru TDP

1. Trusted Data Format
2. Policy Tagging & Configuration
3. Key Authorization
4. Access Control
5. Analytic Container Integrity
6. Output Audit & Protection

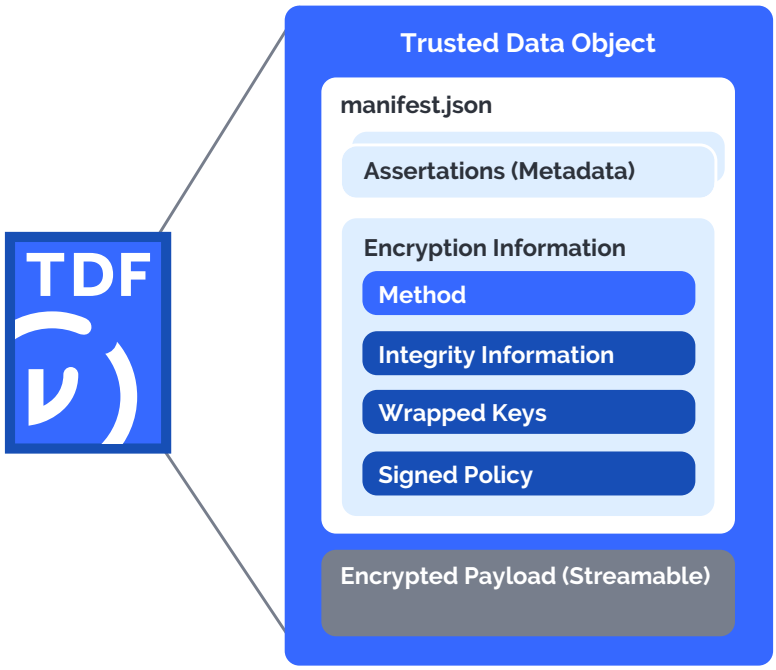

Policy Tagging & Configuration

Simple Attribute-Based Access Control Example

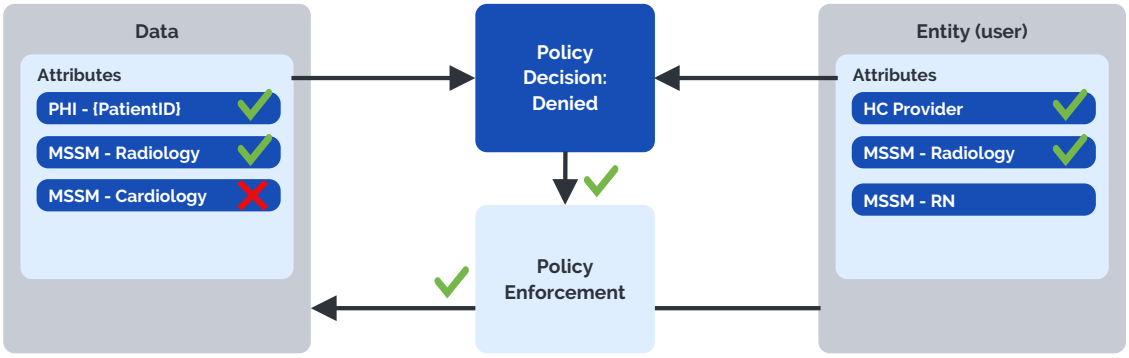

Once protected data is ingested into the TDP, data owners can manage user and analytic access to both their data and its derived outputs – because the TDP leverages a cryptographically-enforced ABAC model.

In this ABAC model, access decisions are made based on data attributes, policy logic, and user attributes. Before they can gain access to a dataset, data tag owners must assign all applicable data entitlements to their users and analytics. The TDP ensures that access to datasets remains cryptographically enforced by encrypting all data and enforcing access policies on their associated key servers.

Whether managing unstructured patient notes, telemetry from a sensor, or genetic code, tagging data and tying policies to those tags is a powerful way to ensure persistent, secure, multi-party control of data at scale.

Key Authorization and Access Control

In order for an analytic to process any encrypted data, the data owner must first grant authorization to the object encryption keys, which are stored in servers separate from the data itself. Data owners can independently control and audit the use of their data by managing the data tags and associated policies used by the key servers, even after the data is mixed with other data, be it in an S3 storage bucket or another type of container.

Data owners may choose to use on premise, cloud, or third party hosted key servers. In order for analytics to request keys, the TDP gives each process a strong digital identity that can be used to connect to the data owner's key server.

After data is encrypted and co-located, data owners can approve analytics to access their data. At the time of approval, analytics must define how the output or derived data will be tagged. In some cases, all of the tags of all the input data may be required to persist on the output. In other cases, data may be summarized or obfuscated enough as to not require all of the raw data tags.

To ensure integrity of analytic logic, registered analytics are hashed prior to approval and cannot be modified without being re-approved. In order for a piece of new analytics software to be granted access to the data and the TDP, it must be able to prove its identity and confirm it has been securely configured.

Key Authorization and Access Control

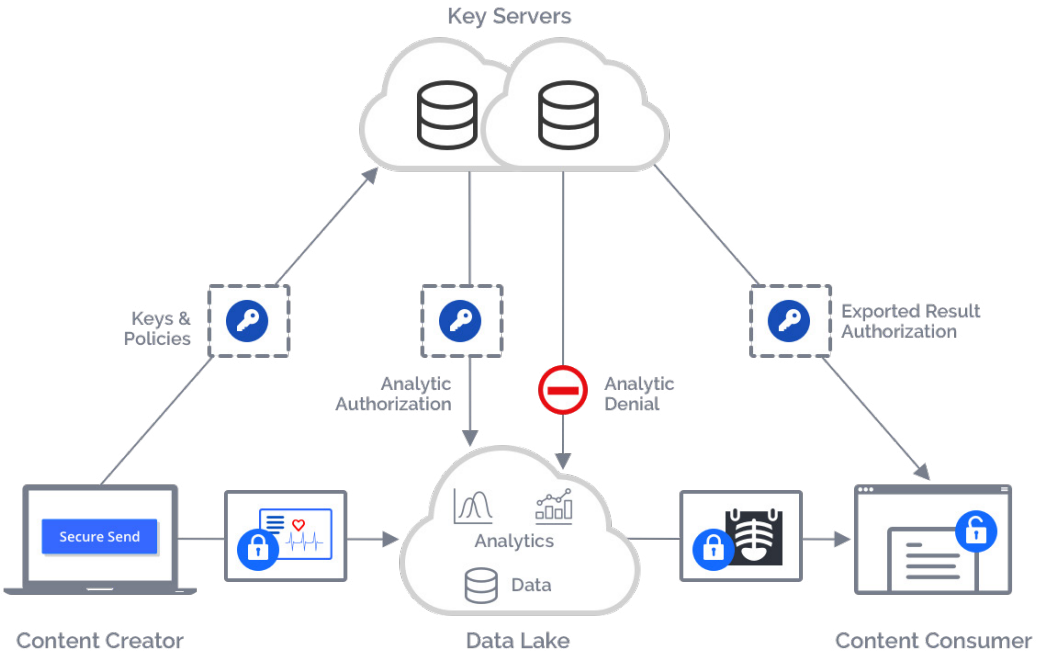

Analytic Container Integrity

After an analytic has been assigned a strong identity and been approved to run over one or more datasets, it also needs a secure location in which to run. The Virtru TDP Analytic meets this requirement by offering secure machine images where analytic owners can run their containerized software. The image provides cryptographic integrity measurements of the software image, as well as tamper-resistant configurations that are required by policy. All of these parameters are digitally signed and part of the identity used to request data access.

Attribute authorities can require that TDP container identities be issued attributes, like the analytics, before they can be approved to process data. In addition, tamper protections can ensure that clear text is never allowed outside of the secure analytic container unless authorized, and that container configurations cannot be modified without forcing re-authentication via a trusted measurement and attestation process.

Entities can require that approved analytics are running in a secure TDP container in order for key authorization to be granted to a collaborating party.

Output Audit & Protection

The TDP can enforce object-level encryption of analytic outputs or other derived data by either manually approving output tags or creating dynamic tags based on a combination of input data and analytic attributes. Only entities that have been granted authorization by these tags will be able to access the platform's data outputs.

For outputs that require a tag from the originator, the originator has the ability to audit all subsequent access to the derived data, change the access policy associated with their data tag, or revoke access entirely.

Extending the Data Collaboration Lifecycle

In addition to giving output owners full control over access and sharing permissions, the TDP's policy manager also enables data owners to securely chain new inputs as follow-on values to be added to existing outputs. As a result, analytic workflows remain fluid and open to additional collaboration beyond a specific use case or stage of analysis.

By providing highly controlled interfaces for the encryption and decryption of analytic processes, there are no mathematical limitations placed on the kinds of questions that can be asked of and answered by the data on which they run.

If you'd like to learn more about using Virtru's Trusted Data Platform to integrate healthcare analytics into your infrastructure, please contact us at platform@virtru.com.

About Virtru

At Virtru, we empower organizations to unlock the power of data while maintaining control, everywhere it's stored and shared. Our portfolio of solutions and tools—built on our open data protection platform—governs data throughout its lifecycle. More than 5,000 organizations trust Virtru for data security and privacy protection. For more information, visit virtru.com or follow us on Twitter at @virtruprivacy.