MIT Connection Science

Towards an Internet of Trusted Data:
A New Framework for Identity and Data Sharing
August 2016

Input to the Commission on Enhancing National Cybersecurity

Alex Pentland
pentland@mit.edu

David Shrier
shrier@mit.edu

Thomas Hardjono
hardjono@mit.edu

Irving Wladawsky-Berger
irvingwb@mit.edu

Massachusetts Institute of Technology
connection.mit.edu

Towards an Internet of Trusted Data:
A New Framework for Identity and Data Sharing

Summary of recommendations from a July 2016 meeting at MIT with senior executives from AT&T, IBM, MasterCard, Qualcomm and U.S. Departments of Treasury and Commerce, as reported by MIT Connection Science leadership.

I.   Executive Summary

As the economy and society move from a world where interactions were physical and based on paper documents, toward a world that is primarily governed by digital data and digital transactions, our existing methods of managing identity and data security are proving inadequate. Large-scale fraud, identity theft and data breaches are becoming common, and a large fraction of the population have only the most limited digital credentials. Even so, our Digital Infrastructure is recognized as a Strategic National Asset which must be resilient to threat. If we can create an Internet of Trusted Data that provides safe, secure access for everyone, then huge societal benefits can be unlocked, including better health, greater financial inclusion, and a population that is more engaged with and better supported by its government.

The future of National Cyber Security should be supported by an Internet of Trusted Data in order to enable both auditable provenance of identity and the credibility of data. in order to enhance economic viability of new technology solutions, policies and best practices. Simultaneously, an Internet of Trusted Data must protect the privacy of people, ensure public safety, economic and national security, and foster public, individual and business partnerships.

In order to accomplish these goals thought leaders in federal, state and local governments should join with academia and carrier-scale private industry to work toward an Internet of Trusted Data.

An Internet of Trusted Data includes:

- **Robust Digital Identity.** Identity, whether personal or organizational, is the key that unlocks all other data and data sharing functions.   Digital Identity includes not only having unique and unforgeable credentials that work everywhere, but also the ability to access all the data linked to your identity and the ability to control the "persona" that you present in different situations. These pseudonym identities, or personas, include the "work you", the "health system you", the "government you" and many other permutations specific to particular aspects of your individual relationship with another party. Each of these pseudonym identities will have different data access associated with them, and be owned and controlled only by the core "biological you". To accomplish this there needs to be a global strategy for Identity and Access Management that genuinely enables trusted, auditable sharing relationships and functions without compromising personal anonymity or security. This is technically straightforward: the

basics were established by the NIST's National Strategy for Trusted Identity in Cyberspace program and now are widely available from, for instance, mobile operators and similar regulated services. The nation now needs to begin requiring such robust digital identity in order to achieve our goals in cybersecurity and universal access.

- **Distributed Internet Trust Authorities.** We have repeatedly seen that centralized system administration is the weakest link in cybersecurity, enabling both insiders and opponents to destroy our system security with a single exploit. The most practical solution to this problem is to have authority distributed among many trusted actors, so that compromise of one or even a few authorities does not destroy the system security consensus. This already standard practice for the highest security systems: no one single actor can launch nuclear missiles, for instance. Now we need to implement this sort of consensus security widely. Examples such as the blockchain that underlies most digital cryptocurrencies show that distributed ledgers can provide world-wide security even in very hostile environments. Today there is a huge amount of investment by private companies to deploy software defined network technology which can transparently expose efficient, convenient versions of this consensus ledger technology, and the U.S. should set policies that take advantage of these new capabilities in collaboration with the private and education sectors, in such a way that digital identities can be originated by individuals and issued with verification from multiple access providers.

- **Distributed safe computation.** Our critical systems will suffer increasing rates of damage and compromise unless we move decisively toward pervasive use of data minimization, more encryption and distributed computation. Current firewall, event sharing, and attack detection approaches are simply not feasible as long-run solutions for cybersecurity, and we need to adopt an inherently more robust approach. The "optimal" technology for such an inherently safe data ecosystem is currently being built and tested, for reference see MIT's Enigma project. Because of the importance of acting quickly, the EU data protection authorities are supporting a simplified, easy-to-deploy version called OPAL (Open Algorithms, which originated at MIT with French support) for pilot testing within certain countries. The concept of OPAL is that instead of copying or sharing data, algorithms are sent to existing databases, executed behind existing firewalls, and only the encrypted results are shared. This minimizes opportunities to attack databases or divert data for unapproved use, but places restrictions on the ability of an ecosystem to collaborate on data when it is in an encrypted state. Note that OPAL may be combined with anonymization identifying elements in order to reduce risk, and in the long run will evolve toward a fully-encrypted, computation friendly model. Approaches such as homomorphic encryption and secure multiparty computation can enable encrypted data to be used in approved, auditable manner by parties that can't decrypt it or read it. In particular, the ability to permissibly ask questions of data in the form of "attributes" will be a key pattern to maintaining digital privacy while enabling innovation ecosystems. The U.S. Government should create a roadmap for progressing

from the current situation, through transition technologies such as OPAL, to complete solutions such as MIT Enigma.

- **Universal Access**. The advantages of secure digital infrastructure are diminished without universal access. The U.S. Government can promote universal access by policies that provide for secure, citizen-controlled Personal Data Stores for all citizens in a manner analogous to current physical Post Office Boxes, and promote their use by making government benefits and interactions such as tax transfers and information inquiries conveniently available by mobile devices and web interfaces secured by the citizens' digital identity.   Planning by the U.S. Post Office for such universal Personal Data Stores (Digital Mailboxes) has long been in place, and the secure digital identity infrastructure is already offered by mobile operators and other regulated services.

Investment required.

We recommend that the U.S. Government establish a "Living Lab" to not just test, but actually create a small-scale deployment of this new ecosystem under real-world conditions with all available and necessary technology in order to obtain citizen and stakeholder feedback. A Living Lab would prove concept and build citizen confidence towards large scale deployment and prove viability of technical solutions.   We also recommend that the U.S. Government support "microdegrees" leveraging distance education methods (e.g., MOOC, etc.) in order to upgrade the cybersecurity training of the existing workforce. Such continuing education methods have proven quite cost-effective in changing the technology culture within U.S. companies.

The Living Lab provides a venue where researchers and developers can begin to address challenges around the Internet of Trusted Data, providing them with real-world conditions under which a robust identity system must be deployable with distributed trust authorities, as exemplified by proposals to use blockchain technology. It also permits data sharing to be explored at scale while preserving privacy, where algorithms are sent to existing data repositories based on distributed safe computation. Key to the Living Lab is universal access to the benefits of the convergence of these new solutions.

II.   Introduction

Our economies and societies are going through a historical transition from the industrial age of the past two centuries, - whose models have been mostly based on physical interactions, - to an increasingly digital age based on global, digital interactions. Previous methods for managing identity and data security have proved inadequate in our emerging digital world, and have led to serious cybersecurity breaches.

A number of failure modes emerge at the current transition point.

 While economy and society are becoming digital, identity remains rooted in analog concepts. The consequences of issues such as identity theft include massive fraud, ranging from bank and insurance to tax and even Uber and AirBnb. A parallel issue emerges of equity and fairness: robust digital identities must be available to all individuals.

 Commercial and government organizations have traditionally built silos of IT systems and data stores each are largely incompatible with each other. In order to move to the next level of a digital economy and to attain the speed and efficiency of business moving at the speed of the network, there must be interoperability and sharing to foster public, private and individual collaboration on trusted data. It must be efficient and based on universally agreed protocols while maintaining security and auditability.

At the same time, our increasingly digital world is opening up opportunities for economic inclusion, improved health care, better financial support and populations that are more engaged with and supported by their government.

To help address these threats and opportunities, President Obama issued an Executive Order earlier this year establishing the Commission on Enhancing National Cybersecurity within the Department of Commerce. The Commission was charged with "recommending bold, actionable steps that the government, private sector, and the nation as a whole can take to bolster cybersecurity in today's digital world, and reporting back by the beginning of December."

The Commission invited select members of this working group to participate in a panel on research and development opportunities at a public meeting held in New York City on May 16. In our testimony to the Commission we highlighted six key areas where government, technology companies and academia should work together in order to Increase the speed and quality of the two-way information flows that are essential for developing a data-rich society with a holistic approach to cyber protection:

 Proof of Identity. A new identity management system must be created to replace today's ad hoc systems.

- **Trustworthy, auditable data provenance**. Systems must automatically track every change that is made to data, so it is auditable and completely trustworthy.

- **Secure, privacy-preserving processing**.   We have to enable the entities to engage in transactions and to verify that contracts are being fulfilled but without revealing private or confidential information.

- **Universal access**. Everyone must be able to share in the benefits, and have the protections, of this new trusted data infrastructure

- **Research and development**. We need to dramatically increase the speed and scale of cyber innovation in both the private and public sector by use of "living lab" field trials

- **Workforce development**. Companies face a serious shortage of cyber trained personnel and of management expertise in cybersecurity. We need to increase and maintain the available workforce, which may require greater educational capacity and incentives.

We also emphasized that the essential importance of focusing on complete systems, rather than individual technologies or technology layers, and that they be developed and proven in "living laboratories" with a representative population of users in order to provide feedback about the relevance, efficiency, effectiveness, and ethical dimensions of these new systems.

In discussions with the U.S. Secretary of Commerce, the White House panel on cybersecurity, and the E.U. VP of Single Digital Market, we were encouraged to create a plan for accomplishing these goals that brought together key players from government, industry, and academia.

On July 11 we convened a workshop at MIT to start framing our statement of the problem as well as our recommendations. The workshop participants all agreed that we are at a unique point to move forward, much as was the case with the Internet and World Wide Web in the early 1990. As was the case then, there is a growing consensus on:

1. The problems to be solved:

- The need to bolster cybersecurity in our increasingly digital economy and society
- A requirement for universal, highly secure digital identities covering individuals, private and public institutions and "things" (IoT).
- The need to efficiently access, exchange and share critical data with full security and privacy protection.
- The overall systems have to be "fault tolerant" in the presence of "non-trusted" actors, whether they are competitors or other governments you only want to share limited data with, or "bad actors" with malicious intent.

2. The fact that there is promising evolution of a new set of general technologies and potential solutions to help address these problems, including:

- ☐ Identity, whether personal or organizational, and, moreover, the ability to own and assert identity attributes is a lynchpin concern. There needs to be a kind of "internet of identity" to genuinely enable all other sharing functions.
- ☐ Blockchain networks can provide a single source of auditable truth between organizations and some level of appropriate automation of data processing. However, organizations must decide also on distributed sources of trust for the moderation) of such networks. Identity plays a crucial role in enabling blockchain technology to be adopted broadly.
- ☐ Overlaying inventions such as personal data stores and secure multiparty computation (see for reference implementations MIT OpenPDS and MIT Enigma), we can develop a new digital ecosystem that is secure, trusted, and empowering.

3. The need for the private sector, government and academia to work together to address these critical problems and leverage these promising technologies to enable and incentivize collaboration and innovation.

At the same time, the population needs to be educated in a coherent fashion about the benefits and role that each individual can play in forming this new system.

While the question arises "are we simply replacing known risks with unknown risks?", the shortcomings of existing systems are proving so great that a new approach is needed
In light of these concerns, we have articulated potential solutions around Robust Identity and Trusted Data which enable the auditability and credibility of both while supporting Fair Information Practice Principles.

III.   The Promise of Robust Identity:

Our mission in suggesting a robust identity framework focuses on connecting the individual with the digital identity, while protecting privacy. When we say "robust", we mean both reliable and non-forgeable.

Benefits include better access:
- [ ] to the financial system for the underbanked and unbanked;
- [ ] to the health care system, in a fashion that reduces medical error and improves care;
- [ ] to government services;
- [ ] to other basic services (e.g., making it easier to obtain an apartment or home).

Drivers of need
A number of problems are driving the need for a robust identity:
- [ ] It is fundamental to cybersecurity. Current cybersecurity systems are insufficient to the task, as evidenced by the numerous large-scale data breaches recently experienced by both the private and public sectors globally.
- [ ] We all need identity to access services.
- [ ] The flaws in translating from a physical proof of identity ("I see you in front of me, I know you are you"), to a digital format ("On the internet, no one knows you are a dog"), in a robust and portable fashion.
- [ ] The question of authority: who provides the identity?
- [ ] The need for it to be unique, strong, verifiable, and non-forgeable.

A potential solution
A new paradigm asserts that you are your digital footprint, and you have ownership rights in your data. Just as the Magna Carta established a framework for individual property rights in 1215 A.D., so too a new digital social contract needs to provide for digital ownership rights.

The concept of behavioral biometrics is gaining ground in areas such as financial services and digital authentication. Research conducted at MIT and elsewhere has demonstrated that behavior biometrics are much more difficult to fabricate and deliver 10X+ better security than password-based models[1]. With respect to ownership rights: this "New Deal on Data" was first posed by us in collaboration with the World Economic Forum in 2009 and expanded since[2].

Several challenges need to be addressed, including, but not limited to:

- [ ] **"Big Brother":** the fear of a government using panoptic access for dictatorship;

- [ ] **Bad actors**: whether inside an organization or external, accumulating such data creates risk of misuse by bad actors;

- [ ] **Scale and implementation cost**: implementing such a solution globally will have nontrivial scaling and cost functions (and relatedly, who will pay for it and how?);

- **Undocumented residents**: deploying such a system creates potential for institutionalizing a digital divide between rich and poor, and introduces new questions around circumstances such as where municipalities offer undocumented residents a means of identification even if the Federal government hasn't;

- **Equal protection under the law**: how can we protect someone that the system doesn't acknowledge has an existence?

- **Universal vs. silo'd data**: universal data has greater utility, but is generally less secure – siloing can provide a measure of security, but raises issues of interoperability; and

- **Regulatory lag:** there is always a gap between a technology innovation and the ability of policymakers and regulators to implement an appropriate framework around it.

Questions also remain as to how this would be developed. Should it be government led, like the EID or India or Estonia? Should it be furnished by industry, similar to how Internet Domain Name registries are handled? Should it be housed within a nonprofit or academic environment, like the Kerberos Consortium or the World Wide Web Consortium (W34C)? Active dialog with key stakeholders is required to establish the optimal path.

Potential Solution: Core Identities and Persona Identities
At the heart of digital identities is the concept of the *core identity* of an individual, which inalienably belongs to that individual. The core identity serves as the quantum from which emerge other forms of digitally-derived identities (called *personas*), that are practically useful and are legally enforced in digital transactions. An individual must have the freedom to choose to deploy one or more digital personas on the Internet, each used with specific sharing and access permission and tailored to the specific aspect of that individual's life. Each digital personal would carry varying degrees of legal enforceability as relevant to the auditable usage context of that persona.

The individual must be able to use transaction-identities derived from his or her relevant persona, without affecting the privacy of their core identity. This derivation process must also allow the relying party (counterparty) in a transaction to validate the source-authenticity and strength of provenance of the transaction identity, without affecting the privacy of the core identity of the user. New cryptographic techniques – such as zero knowledge proofs – offer a promising direction in providing solutions for privacy-preserving core identities.

As currently configured, existing business models, legal instruments and technical implementations are insufficient to support this type of identity ecosystem. This is because something is missing: an architecture for individual ownership of and primacy over one's own core identity and which entities or relationships have access to attributes of that identity. With such a core identity, it is possible for multiple aliases, accounts and attributes to be authenticated and authorized in a reliable, privacy enhancing and scalable manner. To this end, a viable identity infrastructure provides a way for each person to own their single underlying

core identity and to bind several "personas" to that core identity without the need for other parties to access the core identity or be aware of any other personas. With this approach, government issued identity credentials such as driver licenses, passports, professional licenses, birth certificates, etc.) as well as strong-provenanced sources of attributes about a person (e.g. banks for credit scores, etc) can be leveraged to create a core identity that remains private to the end-user.

A key feature of the new model is that it must allow entities in the ecosystem to (i) verify the "quality" or security of an identity, and (ii) to assess the relative "freedom" or independence of an identity from any given authority (e.g. government, businesses, etc.), and (iii) to assess the source of trust for a digital identity.

We believe a new model for digital identities for future blockchain systems is required, which is summarized in the following progressive steps:

1. **Strong provenanced attributes**: It must be founded on an existing real-world identity which has a high degree of source of trust, where its attributes have a high degree of provenance. This identity maybe issued by an existing identity provider or other trusted third party operating within a legal jurisdiction (e.g. Bank, Government, Service Provider, etc.).

2. **Transitive source of trust**: Create a "Core Identity" based on the existing high quality identity. That is, use a privacy-preserving algorithm that translates the existing real-world identity with strong provenance into a digital core-identity which carries-over the source of trust.

3. **Self-issued derived identities as personas**: Provide users with the freedom (and algorithms/tools) to establish personas and to self-issue anonymous but verifiable transaction-identities, each of which is cryptographically derived from the user's core-identity and each of which carries specific permissioned attributes suitable for the purpose of the transaction-identities. The source of trust from the core-identity must also be carried-over into the derived transaction-identity.

4. **Privacy-preserving verification**: Provide the Relying Parties (counter-party) with a privacy-preserving verification algorithms to validate the source of trust for any given (anonymous) transaction-identity. These verification algorithms must allow a relying party to establish a chain of provenance (from the transaction-identity all the way back to the origin attributes and core-identity), while preserving the privacy of the owner of the identity.

5. **Legal Trust Framework (LTF)**: Establish an identity ecosystem for blockchain based on a LTF for core-identities, personas and anonymous transaction-identities. Such a legal

framework is already in use for identity-federation schemes in the industry today, and may be used as the legal basis for this new model.

A legal trust framework is a certification program that enables a party who accepts a digital identity credential (called the relying party) to trust the identity, security, and privacy policies of the party who issues the credential (called the identity service provider) and vice versa. An LTF applies within a given deployment ecosystem, such as identity-federation or across two partner organizations.

We believe the current LTFs as practiced in the industry can be extended for usage in blockchain systems. New types of entities will be needed specifically for blockchain ecosystems. We denote these as the *Core-Identity Provider* and *Transaction-Identity Providers* which extends the current role of the Identity Provider (IdP).

The Core-Identity Provider takes a user's existing identity which has a high degree of source of trust and converts it using a privacy-preserving function into a private or secret core identity that is maintained as private or secret, and is only supplied to the Transaction-Identity Provider. The later then provides a transaction-identity issuance service to the user, as well as a validation service to the relying parties. The user is free to obtain one or more anonymous transaction-identities from the Transaction-Identity Provider or self-issue a derived transaction-identity, all the while maintaining their privacy. The transaction-identities can be used on the blockchain system with other users (relying parties) or on the Internet. The validation service offered by the Transaction-Identity Provider allows a relying-party to inquire about the status and source-grade of a given anonymous transaction-identity prior to transacting.

In the context of blockchains, the LTF provides the following:

- **Network scalability**: It allows any two parties to transact on a blockchain without prior engagement, thus achieving network scalability.

- **Provenance assessment**: It allows a relying-party (counter-party) to assess the "trustworthiness" (provenance and quality) of an (anonymous) transaction-identity prior to commencing the transaction;

- **Cross-jurisdiction interoperability**: It provides a legal foundation for core-identities and (anonymous) transaction-identities to be recognized in differing legal jurisdictions;

- **New business models**: It incentivizes service-providers (including the Core-Identity Providers and Transaction-identity Providers) to develop new business models around new scalable services and permissible use of attribute data associated with identities;

- **Risk assessment and risk management**: It provides entities in the ecosystem with a means of assessing risk and of legal recourse in unforeseen circumstances (e.g. attacks to the service; identity leaks; identity-data theft, provider negligence, etc.) as specified in the LTF operational contracts.

IV.   Data Sharing

Data is rapidly proliferating from an end-user perspective, without a good solution to manage user data as well as identities efficiently.   We need a new paradigm for data sharing that preserves user privacy, while allowing data to be shared more globally for the benefit of society.

The constituents served by this new system include both enterprises and average consumers. By bringing control back to the consumer, both data and identity we can drive better outcomes. We can create technology that would enable the simplification and securing the digital identities through simpler "form factors."

Questions remain – what exactly does this solution look like? Could it be like Global Entry? Who would manage such a system? In Global Entry's case, it's the TSA.

With respect to R&D, there's a great diversity of work required, and we need to an approach such as outlined in Section V.   Opportunities created from this solution include the creation of digital marketplace for personal data, that would simultaneously provide assurance of your data.

Key Concepts in a Potential Solution
Data sharing in a privacy preserving manner requires a new view on data. There are a number of key concepts and design principles that need to be addressed through an evolutionary proof-of-concept (PoC) implementation. Some of these key concepts are as follows:

(1)   **Moving the algorithm to the data:** The concept here is to perform the algorithm (i.e. query) execution at the location of data (referred to as the data-repository). This implies that raw-data should never leave its repository, and access to it is controlled by the repository/data owner.

(2)   **Open Algorithms**: Algorithms (i.e. queries or scripts) must be openly published, studied and vetted by experts to be "safe" from violating the privacy requirements and other requirements stemming from the context of their use.

(3)   **Permissible Use:** When performing computation on attributes or data associated with identities, respect the explicit and implicit permission, or consent, given for use of the data or identity attributes as part of the transaction.

(4)   **Always return "safe answers" (never raw data)**: When performing computation (e.g. in answering a query), the data-repository must always return "safe answers" and never raw data. This concept seeks to address the issue of data privacy and the potential danger of de-identification (of Personally Identifying Information, or PII) through the correlation of multiple responses.

(5)   **Data always in encrypted state**: Data should be encrypted at all times, namely at-rest, in-transit and during computation. Data should not need to be decrypted prior to computation and then re-encrypted afterwards. Advanced cryptographic techniques are now emerging that allows limited forms of computations to be performed on the encrypted data. There two broad scenarios we seek to address:

- **Computation by individual repository**: Here, computation over encrypted data is performed by a single repository, which may employ a physically distributed set of nodes (e.g. P2P collaboration network) to collaboratively store parts of the data for increased resiliency against attacks. Cryptographic techniques such as secret-sharing provides for interesting possibilities in addressing these requirements.

- **Collaborative computation by multiple parties (multiparty computation)**: Some form of queries may require answers to be computed by multiple participants (e.g. repositories) in a collaborative/quorum method whilst maintain the privacy of data stored at each repository. Each participant may see the final group-computed value, but they must not see each other's raw data. Cryptographic techniques such as Multi-Party Computation (MPC) offer a path forward in solving these scenarios.

(6)   **Networked Collaboration Environments and Blockchains for audit and accountability**: Peer-to-Peer (P2P) networks – such as those underlying the Hyperledger system – offers an attractive solution for data resiliency and scalability specially when combined with cryptographic techniques such as secret-sharing and MPC. The consensus-based ledger mechanism underlying these blockchain systems offers a way to perform logging, audit and accounting of queries executed against data in distributed repositories.

(7)   **Social and economic incentives**: For privacy-preserving data sharing to scale and be adopted by a wide range of stakeholders, social and economic incentives must be provided, not only for persons or organizations holding "edge data" but also for infrastructure providers. These infrastructure providers are entities who deliver P2P network scalability, as well as efficient edge computing services that yield real-time edge analytics and visibility into the state of data sharing.

The overall goal of any proposed solution should be increased data protection and privacy, together with scalability, performance and interoperability. In the following, we describe different evolutionary phases of a solution, where each phases focuses on one or more of the above key concepts.

Solution - Phase I "Deployed Blockchain"

The goal here is to explore the use of small number of independent data repositories together with P2P nodes and blockchain technology such as Hyperledger:

- **API-driven query/response controls**: The API defined at the data-repository provides a "hard-wired" query capability. The Querier performs the query by sending a message to the API at the repository, which in turn results to the Querier. The API itself defines the type of query accepted, and the granularity of answers being returned.

- **Aggregate answers only**: The API will return "safe answers" in the form of aggregate/statistical answers only. This approach conforms to the key principle of never allowing raw-data to leave the repository.

- **Multiple repositories**: Multiple data-repositories is envisioned, where each repository may store only one kind of data (e.g. GPS location) that are constrained through its APIs. To the Querier, the nodes on the Networked Collaboration Environment allows Querier to get access to large number data-repositories – represented as nodes of the P2P network (i.e. nodes on the blockchain).

- **Metadata for discovery**: Sharing of data presumes the existence of data is known. A number of nodes on the Networked Collaboration Environment may take the role of "metadata directories", where they can return information regarding the location of nodes with desired data.

- **All access request/response logged**: Direct capture of logs into a blockchain is inefficient and does not scale well. Instead, approaches such as those in Blue Horizon offer as more promising solution. Each party (the Querier and data repository) tracks the API calls, and then hashes the logs. Each party must generate the same hash for the calls, as well as the same hash for data sent/received. These hashes can be written to the blockchain and can be audited by each party only (for privacy preservation).

- **Authorization and Identity**: Authorization tokens will be used in conjunction with a basic blockchain identity, with the focus primarily on authorizations to access a data-repository through the published APIs. The member services capability of systems such as Hyperledger can begin to be explored in this phase.

Solution - Phase II "Bring Algorithms to the Data"

Building on the previous phase, the goal in this phase is to introduce the use of algorithms or scripts that are vetted to be safe for a given classification of data. These vetted algorithms can be signed and published (e.g. at nodes of the Networked Collaboration Environment). Invocation of one of these published algorithms by a Querier will require all the conditions stated be fulfilled (e.g. identity of the Querier is verified, target data repositories, authorization tokens, etc.). The approach build on access APIs from the previous phase, with tighter

restrictions expressed through vetting of algorithms. The Querier's choice of algorithm can recorded on the blockchain, with the returned results being logged and recorded on the blockchain. Furthermore, a published algorithm can be expressed as a smart contract residing in one or more of nodes on the Networked Collaboration Environment.

- **Query control using vetted algorithms/scripts**: Queries are expressed as executable "algorithms" or scripts. The Querier sends the algorithm to the relevant end-points located at the data-repositories.

- **Flexible queries**: In this phase the queries have greater flexibility than the API-driven approach in Phase I. Subset SQL or Python may be considered as the query/scripts language.

- **Vetting of safe algorithms**: Algorithms are first vetted by experts for their safety and impact to privacy and to the correctness of returned results. Only aggregate/statistical queries are permitted. Copies of all approved/vetted queries are signed and then stored at a number of nodes on the Networked Collaboration Environment (participating in blockchain) for public verification.

- **Repositories evaluate and execute algorithms**: Each repository must dedicate computational power ("compute engine") to execute/evaluate a received query (in the form of an algorithm) against the available local data. Raw-data itself is never returned.

- **Richer data repositories**: Each repository is assumed to store richer sets of data of varying types.

- **Authorization & Privacy-preserving Identities**: The identity of the Querier and the repositories need to be preserved from "leakages" of information through methods such as correlations and others that disclose private information. Basic transaction identities can be deployed, as part of managing identity, privacy and confidentiality on the network. Some blockchain systems (e.g. Hyperledger) provide a suitable framework for "member services" (i.e. user's core identity and transaction identities).

- **Blockchain used for logs and monitoring**: A blockchain is used to log all access request/responses for audit and post-event traffic analysis.

Solution - Phase III "Basic MPC"
This phase introduces more sophisticated cryptographic techniques that support privacy-preserving distributed computations over data. Once such technique is *secure multi-party computation* (S-MPC). Here an SMPC cryptographic algorithm is used for computation over data that is stored in plaintext at the repository. A select number of nodes on the blockchain have SMPC capability, and can participate in Secure MPC computation instances. Focus is on the performance aspects of a small number of MPC-nodes, including computational

performance and network bandwidth measurements. An MPC-node may be implemented as an "overlay" over nodes in a blockchain system. A rudimentary proof-of-MPC-completion maybe recorded on the underlying blockchain.

- **Raw data remains locally at the repositories**: The data remains private, and none of the nodes see each other's raw data. Furthermore, each data-item is located at its "home" data repository (i.e. not at a P2P set of nodes).

- **Data at rest in plaintext (not encrypted)**: Data-at-rest a given repository is in plaintext, and not hidden using secret-sharing encryption. This allows focus to be directed to MPC algorithms and their performance.

- **Simple MPC configuration of known nodes**: A small number of repositories (e.g. 3 nodes) will be used to create a group of parties involved in the MPC computation. Each node in an MPC instance will employ a secure channel (e.g. TLS1.2) to ensure integrity protection from attacks, and each know the others identity (e.g. via X509 certificate).

- **Predefined simple queries**: Only simple queries will be addressed, possibly as pre-defined (template) queries. Complex queries (e.g. inner/outer joins of tables) will be left for future work.

- **Metadata service**: Types of data and available simple operations are "advertised" at a special server, to which Queries can locate relevant repositories.

- **Rudimentary Proof-of-MPC-Completion**: MPC-nodes need to record the completion of their MPC-computation on the underlying blockchain, with a matching verification/validation mechanism. The proof is to be determined, but may consists of each MPC-node listing its steps and message flows (with other MPC-nodes), and recording these on the blockchain for later replay/verification. This aspect is relevant for providing economic incentives nodes to participate within a given MPC computation.

Solution - Phase IV "Full MPC with Secret Sharing"
This phase employs a combination of two sophisticated cryptographic techniques, namely the Secure MPC technique (from the previous phase) together with data encrypted into pieces (or "shares"). A minimal "threshold" number of shares are required to re-construct the original data item. A select number of nodes on the blockchain store "shares" of a data item but never the complete set of shares, providing resiliency against attacks seeking to recover the data item. These are called "Shares-node". The MPC-nodes are now also responsible for collecting the relevant shares of each data item from the underlying Shares-nodes in the blockchain. Each Share-node may hold shares corresponding to different data-items, which in turn belong to different owners. The Querier sends the algorithm/query to a coordinating-node that represents the Querier to the MPC-nodes.

- **No centralized data-repository**: Data-repositories no longer hold any complete data, but only location of other nodes (shares-node) in the decentralized Networked Collaboration Environment that hold data (in the form of encrypted data-shares). As the "owner" of a data-item, the data-repository must perform shares management (i.e. shares creation, distribution and re-locations).

- **Shares distribution, re-collection and management**: Each data-repository implements a "standardized" shares location-management function that support the creation of shares-coordinates (of all the relevant shares for each data item). The shares-coordinates are then given to the designated MPC-node involved in a given query instance.

- **MPC only nodes**: A new category of nodes will be introduced whose purpose is to participate in and complete an MPC computation instance. This distinction allows for support of an "outsourcing" model whereby the data repository (who now holds no actual data) delegates the MPC computation instance to a given MPC-node.

- **Shares-based primitive operations**: Simple operations (additions and multiplications) over the encrypted-shares.

- **Query-to-primitive translation**: Simple translator from subset SQL into the relevant MPC primitives (additions and operations).

- **Authorization for invoking MPC-nodes**: The Querier must provide authorization evidence that it authorized is to request the set of MPC-nodes to collect corresponding shares and to perform MPC computation on these shares.

V.   Investment Required

R&D investment is needed to implement the solutions suggested herein. In evaluating an R&D plan, what can we leverage that already exists out there, and what would we need to build.

We recommend that the U.S. Government establish a "Living Lab" to not just test, but actually deploy, this new paradigm. A Living Lab would prove concept and build confidence towards a national and international rollout. MIT has employed the Living Lab model in communities as diverse as Hamburg Germany, the country of Senegal, and Cambridge MA.

We would envision a test starting with 10,000 people, then 100,000+, then 1+ million (roughly logarithmic proof of concept scaling). It could begin with a neighborhood or small community, then a town or larger neighborhood, then a midsize city. It could be deployed to a specific department within the Federal Government. Or, it could be structured on an interest group or common problem area ("affinity-driven"), such as veterans or federal workers.

Competitions such as the Department of Transportation Smart City Challenge, with a $50 million prize for the winning proposal, illustrate a viable model for generating a diverse set of perspectives on the problem and potential solutions.

Criteria for Site Selection
1) Town-scale
2) Proximity to a substantial federal facility
3) End points: consumers, health system, financial system
4) Make interoperable with things that are not new?
5) Advisable: strong local university partner
6) Coordinating "owner" locally

Geographies that meet these characteristics include Boulder CO, Austin TX, Rochester NY, Hartford CT, and Boston MA. Affinity-driven examples include the Veterans Administration or the USPS.

A federally administered competition could stimulate innovation much in the same way that the ARPANet and NSFnet led to the development of the commercial internet.

Background

The Massachusetts Institute of Technology Connection Science initiative (MIT Connection Science) hosted a working group session on 11 July 2016 in Cambridge, MA to answer the call posed by the White House Commission on Cybersecurity, and separately discussions with the U.S. Secretary of Commerce Penny Pritzker and the European Union's Vice President Ansip, in charge of of the Single Digital Market. This document reflects a distillation of the discussion from the July 11 working group.

Participating were:
- Alex Pentland (Professor, MIT)
- Irving Wladawsky-Berger (MIT Connection Science Fellow)
- David Shrier (Managing Director, MIT Connection Science)
- Jerry Cuomo (IBM Fellow and Vice President Blockchain Technologies)
- Steve Davis (Senior Consultant, Payments Innovation, MasterCard)
- Michael Frank (Program Director, Blockchain Technologies, IBM)
- Thomas Hardjono (Chief Technology Officer, MIT Connection Science)
- Guerney Hunt (Research Staff Member, IBM)
- Cam Kerry (Visiting Scholar, MIT; formerly General Counsel and Acting Secretary, U.S. Dept. of Commerce)
- Mark O'Riley (Office of the General Counsel, Government and Regulatory Affair-Technology Policy, IBM)
- Chris Parsons (Vice President Big Data Strategy and Business Development, AT&T)
- Gari Singh (Distinguished Engineer & Blockchain CTO, IBM)
- Anne Shere Wallwork (Senior Counselor for Strategic Policy, Office of Terrorist Financing and Financial Crimes, U.S. Department of the Treasury)
- Rod Walton (VP Qualcomm)
- Irida Xheneti (Entrepreneur in Residence, MIT Connection Science)

---

[1] Yampolskiy R et. al. (2008) "Behavioral Biometrics: a survey and classification" Int. J. Biometrics, Vol. 1, No. 1, 2008

[2] World Economic Forum (2009) "The New Deal on Data") http://hd.media.mit.edu/wef_globalit.pdf