Blockchain and Its Emerging Role in Health IT and Health-related Research

Agency: Department of Health and Human Services

Office: Office of the National Coordinator for Health Information Technology

Debbie Bucci

debbie.bucci@hhs.gov

(202) 690-0213

Technical Point of Contact	**Ketan Paranjape**
Title:	General Manager – Life Sciences
Unclassified Email:	ketan.paranjape@intel.com
Phone Number:	503-712-8643
Company Name:	Intel Corporation – Health and Life Sciences Group
Address	2111 NE 25th Ave, Building JF5
City, State, Zip	Hillsboro, OR. 97124

Collaborators

This whitepaper is written by Intel authors Protik Sandell – Healthcare Software Architect, Health and Life Sciences Group, Mic Bowman – Principal Engineer and Manager – Distributed Ledger Research Group, and Prashant Shah – Director of Engineering, Health and Life Sciences Group. Please forward any comments or questions to Ketan Paranjape – General Manager, Life Sciences, Health and Life Sciences Group.

THIS SUBMISSION IS FOR BLOCKCHAIN AND ITS EMERGING ROLE IN HEALTHCARE RELATED RESEARCH CHALLENGE PURPOSES ONLY. ADDITIONAL DETAILS AVAILABLE UPON REQUEST.

Blockchain and Its Emerging Role in Healthcare Related Research

This whitepaper is a response to a request from the National Coordinator of Health Information Technology, a division of the Health and Human Services Department, soliciting ideas for potential uses for Blockchain Technology in the area of Healthcare.

Contents

Executive Summary........ 2
Distributed Ledger and Blockchain Overview....... 2
Sawtooth Lake Technology.................... 4
Application Layer 5
Ledger Layer 5
Journal Layer............... 5
Communication (Platform) Layer 6
Intel SGX Technology ... 6
Consensus and SGX 7
Proof of Elapsed Time ... 7
Value of Blockchain to Healthcare..................... 8
Proof of Concept - Financial Market 9
Proof of Concept - Fantasy Sports Game 9
Potential Gaps 10
Summary 10
References 10

Executive Summary

Intel recognizes the promise in the new blockchain technology, and has developed the Sawtooth Lake Distributed Ledger Platform (SLDLP), that is targeted at building, deploying, and running distributed ledgers. Intel is testing SLDLP in proof-of-concept (POC) environments in partnership with various external companies to prove the integrity and applicability of the technology. Intel has a Health and Life Sciences division that focuses on advancing personalized medicine, health IT, medical devices, and consumer health. That focus makes Intel well suited to continue to explore the many unanswered questions surrounding applications for blockchain technology and distributed ledgers in the health industry, and to offer a secure end-to-end solution. Intel is interested in doing more collaboration with external healthcare partners and with HHS to address gaps in the use and application of blockchain technology in the healthcare industry.

Distributed Ledger and Blockchain Overview

A distributed ledger can be thought of as a secure database that is implemented among a group of participants without a central authority or administration. Participants in the ledger can submit transactions to add, remove or modify records in the database according to a set of rules that are guaranteed to be enforced by the ledger (for example, the ledger may ensure that you can't spend money that you don't have). Distributed ledgers are powerful enough to use for building a broad class of applications and services like secure, robust cryptocurrencies such as Bitcoin; for providing verifiable ownership of assets ranging from bonds and titles to diamonds and books; and for managing access rights to personal data. And, all of these services can be provided without the requirement that a single organization be trusted with the data.

Underlying the distributed ledger is a technology called “blockchain” that ensures the integrity of the ledger. A blockchain is an immutable series of transactions that is shared by all participants in the ledger. Cryptographic signatures ensure correctness and guarantee “non-repudiation” (that is, that once a transaction is committed to the blockchain, it cannot be un-committed). Distributed consensus algorithms ensure that all participants see the same series of transactions even if bad actors try to compromise the system. The most unique characteristic of a blockchain is that it can provide these capabilities without a central organization to provide authorization or administration.

The blockchain method got its name from the way it builds historical transactions: records of transactions are collected in blocks that are time/date stamped and chained together in chronological order. A new transaction is appended with a timestamp as a new block to the back of the current blockchain. Each valid block in a blockchain contains a reference to the previous valid block, and this creates a chain of blocks that captures the history of a transaction. Blockchain is a machine for creating trust: the blockchain allows a group of users who have no particular confidence in each other to collaborate without having to go through a neutral central authority.

Distributed ledgers based on blockchain technology present several advantages over current record handling methods that rely on a centralized database: 1) Central intermediaries are removed from the transaction, 2) Connections between counterparts are simplified because the data resides at every participant, and 3) Data is recorded on a tamper-proof and secure block chain. Participants in a blockchain event each have their own copy of the stored data in what can be considered a secure, distributed, shared database. Changes to the data are validated by participants collectively, and updated across the network almost immediately.

There are three common types of distributed ledger in use: public (or open) ledgers, private (or permissioned) ledgers, and consortium ledgers. The main distinction among the types is who has access to the database.

Advantages of blockchain and distributed ledger technology are:

- Convergence on a consensus for changes applied to the data
- Decentralized data and authority
- Non-repudiation of data (digitally signed information)
- Full replication of data (distribution of encrypted storage)
- Distribution of trust (fully auditable)

Disadvantages of blockchain and distributed ledger technology are:

- Inefficient method for data transfer
- No service level agreement (needed for Enterprise applications)
- Many unanswered questions: validation, scalability, transaction complexity, number of participants, viability of privacy

Sawtooth Lake Technology

The Intel Sawtooth Lake Distributed Ledger Platform (SLDLP) is a modular prototyping platform that is targeted at building, deploying, and running distributed ledgers. The Intel SLDLP consists of an SGX-enabled SDK platform, an SGX distributed ledger toolkit (white boxes), and the distributed ledger technology (DLT). Figure 1 shows how these pieces fit together.

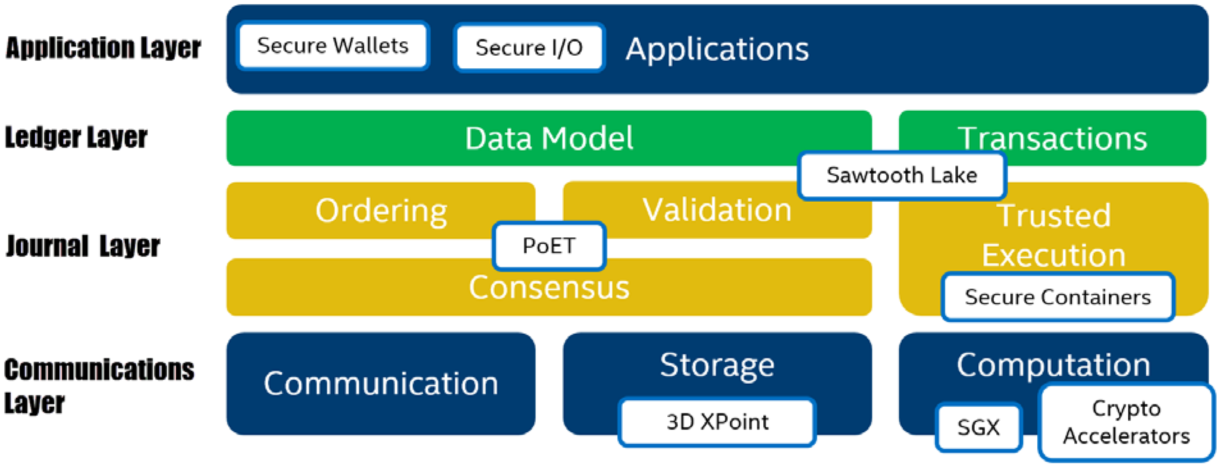

Figure 1. Block Diagram – Intel® Sawtooth Lake Distributed Ledger Platform

The SGX technology (integrated into Skylake) provides an add-on SGX DLT Toolkit (the items shown in the white boxes on Figure 1) to facilitate and streamline operations in eight areas of the distributed ledger environment.

- **Secure Wallets**: Secure management for credentials used to interact with the ledger
- **Secure I/O**: Secure path for information to ensure the integrity of information triggering smart contracts (technology that is in development)
- **PoET**: An algorithm that uses Intel's Software Guard Extensions (SGX) to establish a globally consistent view of the ledger
- **Sawtooth Lake**: Intel's platform for building distributed ledgers
- **Secure Containers**: Secure execution containers (technology that is in development)
- **3D XPoint™ Memory**: Intel's large scale persistent memory
- **SGX**: Intel's Software Guard Extensions, an enclave that protects code and data from exposure and modification
- **Crypto Accelerators**: AES-NI and other accelerators that improve the performance of cryptographic algorithms

The Intel implementation of the Distributed Ledger technology uses a highly modular architecture that consists of four primary layers (an Applications Layer, a Ledger Layer, a Journal Layer, and a

Communications Layer). The modularity of the architecture allows it to be adapted for many different distributed ledger environments, from open consumer markets to closed market services.

Figure 1 shows how the layers interact in a blockchain environment: Application layer, Communications layer, Journal Layer, and Ledger (or Transaction Family) Layer.

Application Layer

At the Application layer, secure transactions are initiated, key management is applied, and external data adapters may be used. Some examples of vertical markets that use distributed ledger technology today are the Financial Services (FSI) Industry, the Internet of Things (IoT), Supply Chain Management, and Currency Markets (Bitcoin). We believe that the basic technology and concepts inherent in the distributed ledger technology can also be applied to the healthcare industry.

A recent security breach at one of the Bitcoin exchanges highlighted the importance of security at the Application layer. Insecure management of private keys in the exchange allowed hackers to steal several million dollars worth of Bitcoins. Intel's distributed ledger solution provides capabilities to improve the security of key management, extending the trusted computing base from the blockchain all the way to the edge of the network.

Ledger Layer

The Ledger (or Transaction Family) layer provides the Data Model with well-defined constraints on correctness. In this layer, transactions use smart contracts to modify data. Ledger layer defines and enforces semantics for transaction families. The data model and transaction language in the SLDLP are implemented as a transaction family. We expect ledgers to build custom transaction families that reflect the unique requirements of the ledger/industry. However, the SLDLP product provides three transaction families that are sufficient for building, testing, and deploying a marketplace for digital assets:

- EndPointRegistry: A transaction family to register ledger services
- IntegerKey: A transaction family to test deployed ledgers
- MarketPlace: A transaction family to buy, sell, and trade digital assets

These three transaction families provide an "out-of-the-box" ledger that can be used to quickly implement a fully functional marketplace for digital assets in many environments.

Journal Layer

The Journal layer is a globally consistent, immutable ordered log of transactions. The layer provides secure sandboxed execution. The Journal layer drives consensus on an ordering of transaction identifiers across all nodes in a validator network. The distinction between *transactions* and *transaction identifiers* provides the basis for separating *consensus* (the role of the Journal layer) from *transaction semantics* (the role of the transaction families in the Ledger layer). When a client submits a new transaction, the corresponding transaction family in the Ledger layer first ensures that the transaction is semantically correct. Once that is established, the identifier for the transaction is passed to the consensus protocol in the Journal layer; this process separates consensus from correctness.

Communication (Platform) Layer

At the Communication layer, the physical platform consists of an SDK with a trusted computing base, a cryptographic accelerator, and long-term persistent storage. This layer implements a gossip protocol for communicating among participants in a validator network. The layer provides basic facilities for sending messages directly to other validators and broadcasting messages to the entire validator network. Messages are typed, signed by the originator, and encoded in CBOR for transmission. Other Sawtooth Lake components can register type-specific handlers for messages.

The Communication layer provides two related abstractions that enable customization for different deployment environments:

- **Topology –** A protocol for establishing connections between peers in the gossip network can be customized. The SLDLP currently implements three protocols, one for random connections (that is most resilient to malicious manipulation), one for scale-free topologies (that supports efficient broadcasts), and one based on a distributed hash table (that supports routing to "interest groups").
- **Routing –** A default method for forwarding messages through the network is a simple broadcast to all peers. However, this may be overridden as is the case with a distributed hash table implementation.

Intel SGX Technology

Intel® Software Guard Extensions (Intel® SGX) is an Intel technology for application developers who are seeking to protect select code and data from disclosure or modification. Intel SGX makes such protections possible through the use of enclaves, which are protected areas of execution. Application code can be put into an enclave by special instructions and software can be made available to developers via the Intel® SGX SDK. The Intel SGX SDK is a collection of APIs, libraries, documentation, sample source code, and tools that allows software developers to create and debug Intel SGX enabled applications in C/C++.

Intel designed the SGX technology to satisfy eight objectives:

1. Allow application developers to protect sensitive data from unauthorized access or modification by rogue software running at higher privilege levels.
2. Enable applications to preserve the confidentiality and integrity of sensitive code and data without disrupting the ability of legitimate system software to schedule and manage the use of platform resources.
3. Enable consumers of computing devices to retain control of their platforms and the freedom to install and uninstall applications and services as they choose.
4. Enable the platform to measure an application's trusted code and produce a signed attestation (rooted in the processor) that includes this measurement and other certifications that the code has been correctly initialized in a trustable environment.
5. Enable the development of trusted applications using familiar tools and processes.

6. Allow the performance of trusted applications to scale with the capabilities of the underlying application processor.
7. Enable software vendors to deliver trusted applications and updates at their cadence, using the distribution channels of their choice.
8. Enable applications to define secure regions of code and data that maintain confidentiality even when an attacker has physical control of the platform and conducts direct attacks on memory.

Consensus and SGX

The Intel SGX technology provides a means for gaining consensus that is key to distributed ledger environments. Transaction families require a consensus protocol, a feature that is show in the Journal layer in Figure 1. The consensus protocol in SLDLP can be modified or replaced easily based on the specific requirements of the ledger implementation. The default consensus protocol uses the Trusted Execution Environment (TEE) feature in the Intel's Software Guard Extensions (SGX) technology. TEE is a highly resilient consensus protocol that does not carry the burden of extremely high computational requirements (as in the Bitcoin version called "proof of work"). The Intel method ("proof of processor" algorithm) scales to thousands of participants and can be run efficiently on any Intel processor that supports SGX[1] technology.

SLDLP is designed to take advantage of the features of Software Guard Extensions (SGX) that are built into the Intel microprocessors (starting with Skylake) to streamline the implementation of distributed ledgers. SLDLP abstracts the core concept of consensus, isolates it from transaction semantics, and provides two consensus protocols with different performance trade-offs. The first protocol, called PoET for "proof of elapsed time", is a lottery design that builds on trusted execution environments provided by Intel's Software Guard Extensions (SGX) to address the needs of large populations of participants. The second, Quorum Voting, is an adaptation of the Ripple and Stellar consensus protocols and serves to address the needs of applications that require immediate transaction finality.

Proof of Elapsed Time

The SLDLP uses a Nakamoto-style consensus algorithm called Proof of Elapsed Time (PoET). As a proof algorithm, PoET uses a lottery for leader election that is based on a guaranteed wait-time provided through a trusted execution environment. For the purpose of achieving distributed consensus efficiently, this lottery function provides several characteristics:

- **Fairness**: The function should distribute leader election across the broadest possible population of participants.
- **Investment**: The cost of controlling the leader election process should be proportional to the value gained from it.
- **Verification**: It should be relatively simple for all participants to verify that the leader was legitimately selected.

[1] SLDLP also provides an insecure emulation of the consensus algorithm that can be deployed more broadly for testing. The emulated version is not recommended for any production use.

Value of Blockchain to Healthcare

Healthcare is undergoing a transformation. These changes are driven by number of federal government initiatives such as Health Information Technology and Economic and Clinical Health (HITECH) Act, President Obama's Precision Medicine Initiative (PMI) Cohort Program and Cancer Moonshot project. These initiatives require the capture of health information using Electronic Health Records (EHR) and make the information available for research and care improvement. Providers are actively using information from the EHR along with genetic data to obtain a holistic view of the patient. Sophisticated use of health information technology is contributing to improving health outcomes, improved health care quality, and lower health care costs – the three overarching aims that the U.S. is striving to achieve (also known as the "Triple aims"). The continued success of using healthcare data in spearheading transformative changes is being inhibited by major hurdles.

Presently, there are significant security and privacy apprehensions (among both patients and providers) that impede sharing of health records. These concerns are fueled by recent high profile security breaches in patient health care records. These include theft of information for 78.8 million patients at Anthem and ransomware attacks at Medstar. The *Sixth Annual Benchmark Study on Privacy & Security of Healthcare Data* (May 2016) revealed that for the sixth year in a row, data breaches in healthcare are consistently high in terms of volume, frequency, impact, and cost. The primary concern of patients is that they have little (or no) control over their information after it has been provided to a payer, provider, or healthcare exchange. Patients want greater insight into how their data is used, who has access to it, and when it is being modified.

Blockchain applications offer great opportunities to address the privacy and security concerns of patients. Mitigating these concerns will facilitate the sharing of secure data and accelerate the PMI and EHR interoperability initiatives from the Federal Government. The key benefit of using blockchain is that it empowers the patient to control access to their health records. Patients can authorize a new doctor to review their health record and obtain a second opinion or grant viewership rights to a guardian, a doctor, a pharmacy, an insurance company, as needed via their private key. Individual patients can contribute part or all of their information to a genetic database (such as PMI) and still retain control over who has access privileges to their information. Blockchain provides the structure for health data such that it can be analyzed, but remain private. Taking advantage of the pseudonymous (i.e., coded to a digital address rather than to a patient name) nature of blockchain technology and its privacy, personal health records could be linked through the blockchain [1].

In today's environment, patients interact with a staggering number of health care providers as they live their lives. They may have medical history from a pediatrician, a university physician, a dentist, an employer health plan provider, or a medical specialist. As they live their lives, patients leave data scattered across many areas within the greater healthcare system [2]. The result is a frayed trail of health records that are hard to collect and difficult to piece together because providers want to retain primary ownership of medical records [3]. Blockchain provides a novel way to securely create a virtual lifelong longitudinal health record by storing encrypted access-links to individual records from disparate health systems into a distributed ledger application, and make the links accessible to authorized users.

Efforts such as PMI aim to follow a cohort of patients through their health journeys across multiple years. Blockchain can help this effort by creating a tamper-proof *distributed event ledger*, or a "breadcrumb" of that patient's journey, with new entries that get added with every health care encounter or event. The ledger acts as a unique fingerprint that can be used to establish the identity of the patient. An authorized researcher from PMI could use the breadcrumb to subsequently query the source system to retrieve the relevant patient record.

Additionally, blockchain provides mechanisms to:

- Ensure that the contents of the records have not been changed (as hash generated would be unique to the document contents)
- Provide an audit trail by adding timestamp to the records
- Restrict access to authorized users
- Enable attribution that encourages individual participation in large-scale projects

Proof of Concept - Financial Market

Forty banks in the R3 CEV blockchain consortium ran a test of blockchain technology in a pilot program that was launched in February of 2016. The test included products from five blockchain vendors and three cloud providers. Technology groups at Bank of America Corp., Deutsche Bank, Morgan Stanley, Royal Bank of Scotland Group PLC and 36 other banks collaborated to build the ledgers (using base technology from rivals Chain Inc., Eris Industries Ltd., Ethereum, Intel Corp. and International Business Machines Corp. Cloud infrastructure came from Amazon.com Inc., IBM and Microsoft Corp) and test the technology. The effort focused on:

- How different combinations of technology handle simulated transactions in commercial paper
- How to execute smart contracts, or coded business rules set up to conduct transactions without human intervention
- How CIOs can determine what criteria to use to evaluate blockchain technologies from competing vendors
- How many institutions can run many ledgers in parallel in a rigorous, scientific way

The R3 CEV Consortium has announced plans to conduct similar tests with the participation of government regulators and tests for integrating blockchains with banks' legacy transaction systems.

Proof of Concept - Fantasy Sports Game

As an internal test for proof-of-concept, Intel Corporation tested the blockchain technology by using Sawtooth Lake as a digital asset exchange platform for a fantasy sports game. In the test, Intel had over 400 participants who entered over 10k transactions over the course of three weeks. In the fantasy game, users traded shares of March Madness NCAA basketball teams and were awarded with digital currency that is based on how well those teams did in the tournament. The effort was successful in testing the blockchain technology.

Potential Gaps

Use of blockchain technology in healthcare is at a promising stage in development. Blockchain-based applications in the healthcare industry are yet to be demonstrated as a viable platform for exchanging and reviewing information. Blockchain technology does hold promise of being a widely adopted mechanism for resolving issues that have been long-term concerns of the industry and inhibited research.

For blockchain technology to succeed, current healthcare applications and processes would need to be extensively modified. Care delivery processes would need to be redesigned so that they are patient-centric. Patients would gain more control over who has access to their health records, and the healthcare industry would have to reduce their control.

Summary

The healthcare industry values many of the basic underlying tenants inherent in blockchain technology, like trusted execution, non-repudiation of data, auditable trails and records for transactions, full replications of data in a secure environment, consensus on data changes, and decentralization of authority/data. At the same time, there are many areas of blockchain that are relatively untested in a healthcare environment, like the need for a service level agreement, viability of privacy, scalability of a system to handle large numbers of participants, control and restrictions around access to patient data, and issues around patient record ownership. Recent security breaches at Bitcoin are not failures in blockchain technology, but rather failures to securely manage the private keys on the exchange. We feel that blockchain technology holds high promise of being a widely adopted mechanism in the healthcare industry for resolving issues that have long concerned the industry. And further, we believe that the Intel solutions could extend the security of transactions beyond just the blockchain/distributed ledger environment to provide an end-to-end solution that includes security for the client key management.

References

[1] - Blockchain: Blueprint for a New Economy - Melanie Swan

[2]"Public Standards and Patients' Control: how to keep electronic medical records accessible but private". BMJ. Vol. 322. (2001): Num. 7281. 283-287.

[3]"Who Owns Medial Records: 50 state comparison". Milken Institute School of Public Health. George Washington University, (2015)