

Tutorial: A New Era in Distributed Computing with Blockchains and Databases

C. Mohan

IBM Fellow

IBM Almaden Research Center, San Jose, USA

Distinguished Visiting Professor

Tsinghua University, Beijing, China

LinkedIn/Telegram/Twitter/WeChat: [seemohan](#)

Facebook: [cmohan](#)

Links to Videos, Slides, Bibliography & Twitter Handles @ <http://bit.ly/CMbcDB>

Agenda

Goal: Introduce practically relevant private/permissioned blockchains (BCs) to non-techies AND techies; also, get into details of many private blockchain systems.

- Origin of Blockchains (BCs)
- Related Distributed Systems/Databases Topics
- Evolution: Smart Contracts, Private BCs, ...
- Consortia Approach to Development of Systems
- Applications: Production, PoCs, ...
- Market Scene
- Benchmarks
- Architectural Choices and Relationship to DB Replication
- Technical Details of Representative Systems:
Enterprise Ethereum, Hyperledger Fabric & Composer, R3 Corda, Coco Framework, BigchainDB, Sawtooth, Ripple
- Futuristic Topics

Blockchain (BC)

- Origin in digital currencies (**Bitcoin** - Satoshi Nakamoto, 2008) – anonymity, **open/public/permissionless** environment – **awful performance** (7 TPS, 10 minutes response time) **and widely-varying transaction fees**
- Numerous organizations across the world working on various aspects of it: security, consensus, database, benchmarks, verification, ...
- Banks, regulators, universities, startups, big tech companies, services companies, governments, ... **mostly as part of consortia**
 - 2/2017: First **production** deployment of BC technology by IBM & Northern Trust in Guernsey for administration of **private equity fund** managed by Unigestion – **Hyperledger Fabric 0.6**
 - 4/2017: China's **Tencent** announced **TrustSQL**
 - 7/2017: **Hyperledger Fabric 1.0** Released
 - **Hyperledger Fabric** on IBM Cloud - **IBM Blockchain Platform** (formerly **HSBN**) on highly secure Linux on mainframes (System Z) with security hardware – announced August 2017 – available in Dallas, London, Frankfurt, Tokyo, Toronto, Washington DC, ...
 - 10/2017: **Oracle** announced Blockchain Cloud Service (BCS) - **Fabric 1.0** based
 - 10/2017: China's **Baidu** joined Hyperledger as a **Premium Member** & 1/2018: Announces **BaaS offering**
 - 3/2018: Hyperledger **Caliper** Benchmarking Project initiated
 - 4/2018: **Huawei** announced Blockchain Cloud Service for China & **AWS** announced Blockchain Templates (Fabric/Ethereum)
 - 5/2018: **Enterprise Ethereum Client Specification** Released
 - 7/2018: IBM announces work on **Stablecoin** (pegged to US\$) Stronghold USD
- Grand View Research: Global BC Tech Market **\$7.74B** by 2024
- **My focus: Private/Permissioned** BC Systems!

Blockchain Jobs

BITCOIN

Cryptocurrencies and blockchain are becoming a hot trend in the job market

- On Thursday, CoinDesk, a leading source of cryptocurrency news and organizer of major industry conferences, launched an online "Career Center" with job listings.
- Listings of "blockchain" skills skyrocketed more than 6,000 percent in the first quarter from a year ago, online freelancing database Upwork said in a report Tuesday.
- However, there are many risks. Sometimes a start-up has a good idea, brings people to work on a prototype, but doesn't get funding, said David Gadd, a Canada-based recruiter focused on blockchain talent acquisition. So the company has to close down.

Evelyn Cheng | @chengevelyn

Published 11:45 AM ET Fri, 4 May 2018 | Updated 2:04 PM ET Fri, 4 May 2018



DON'T MISS: Best Places to Work in IT 2018 · Win 7 to Win 10 migration guide · Mingis on Tech · Resources/White Papers · Job Search

≡ **COMPUTERWORLD**
FROM IDG

Home > Careers

NEWS ANALYSIS

Blockchain moves into top spot for hottest job skills

Blockchain development is now the hottest skill in the freelance job market, growing more than 6,000% since this time last year and putting it on pace to be the new "cloud" of the 21st Century, according to a new report.



By Lucas Mearian

Senior Reporter, Computerworld | MAY 1, 2018 6:00 AM PT



HACKERRANK

SPONSORED BY
manifold

Follow



HOME | READER SURVEY | DEVELOPER MARKETPLACE | 🔍



Kirill Shilov

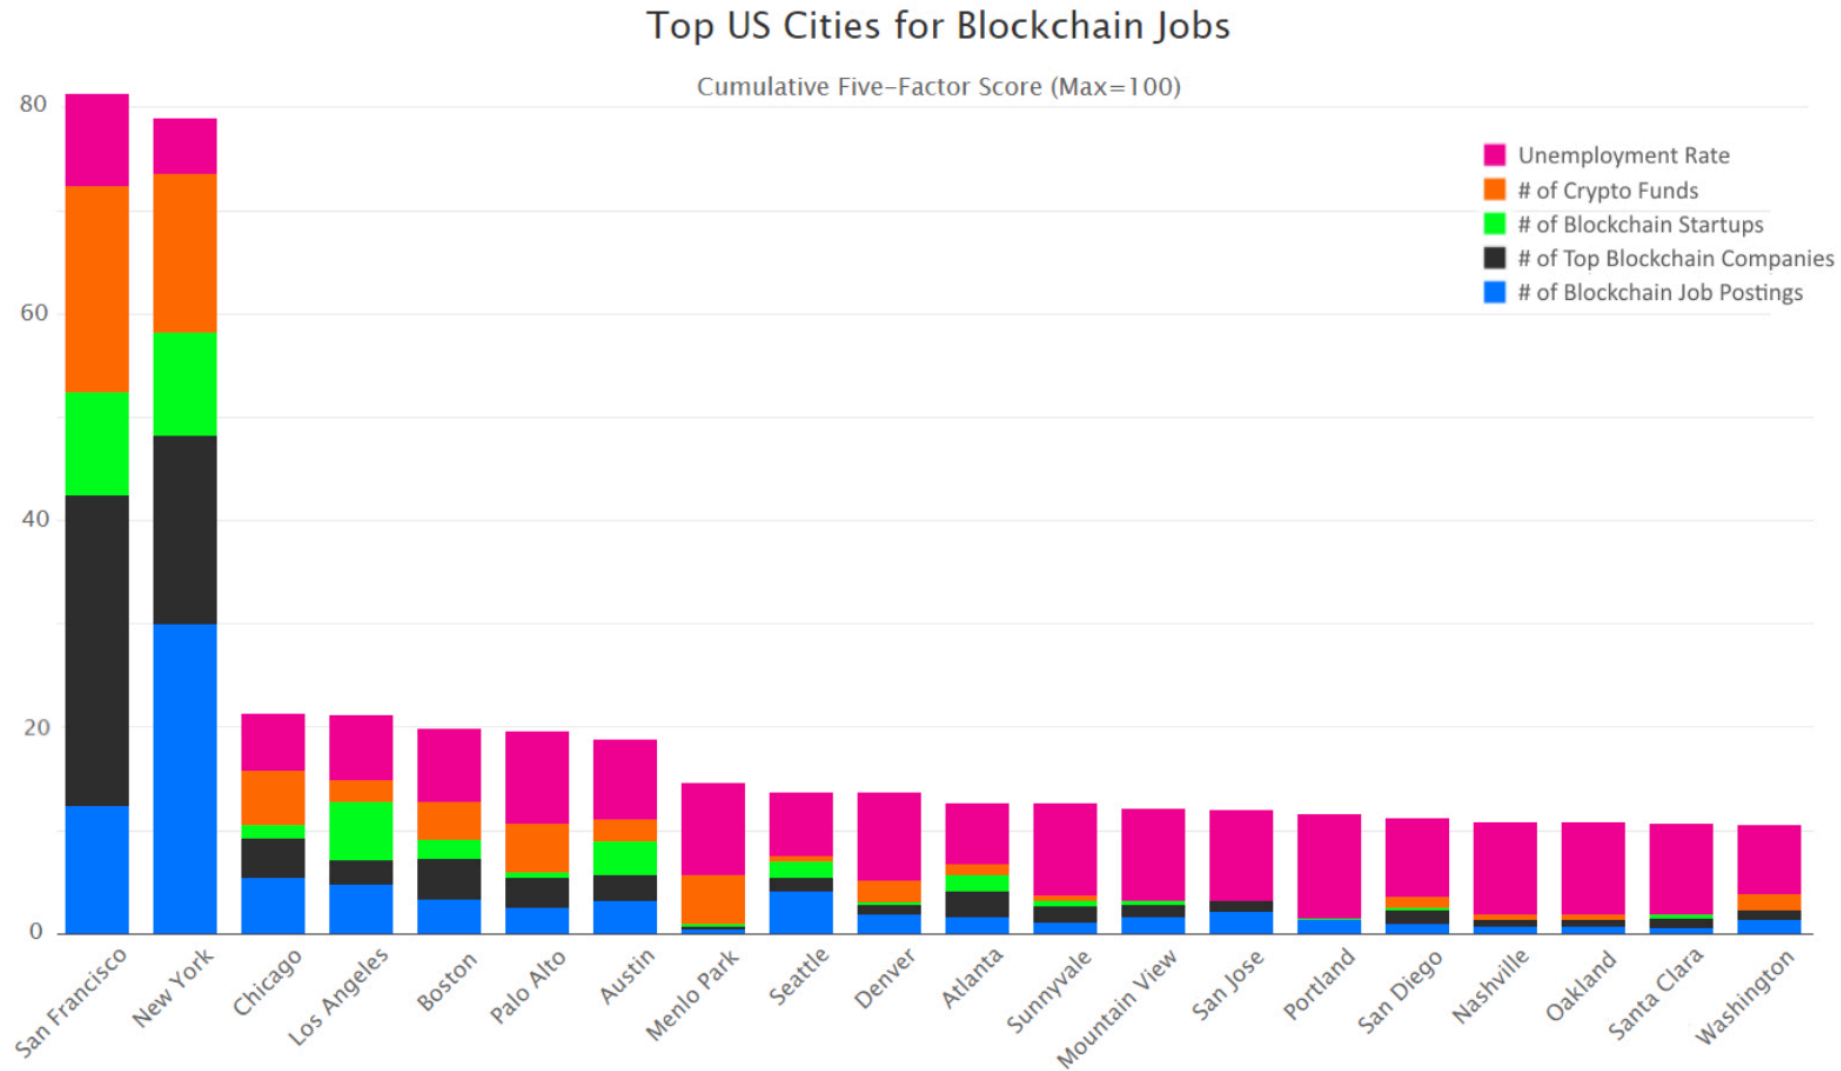
Follow

co-founder howtotoken.com #howtotoken @howtotoken #blockchain #crypto #education

Feb 23 · 10 min read

Blockchain jobs and salaries—2018 report

Blockchain Jobs in USA (Crypto Fund Research 4/2018)



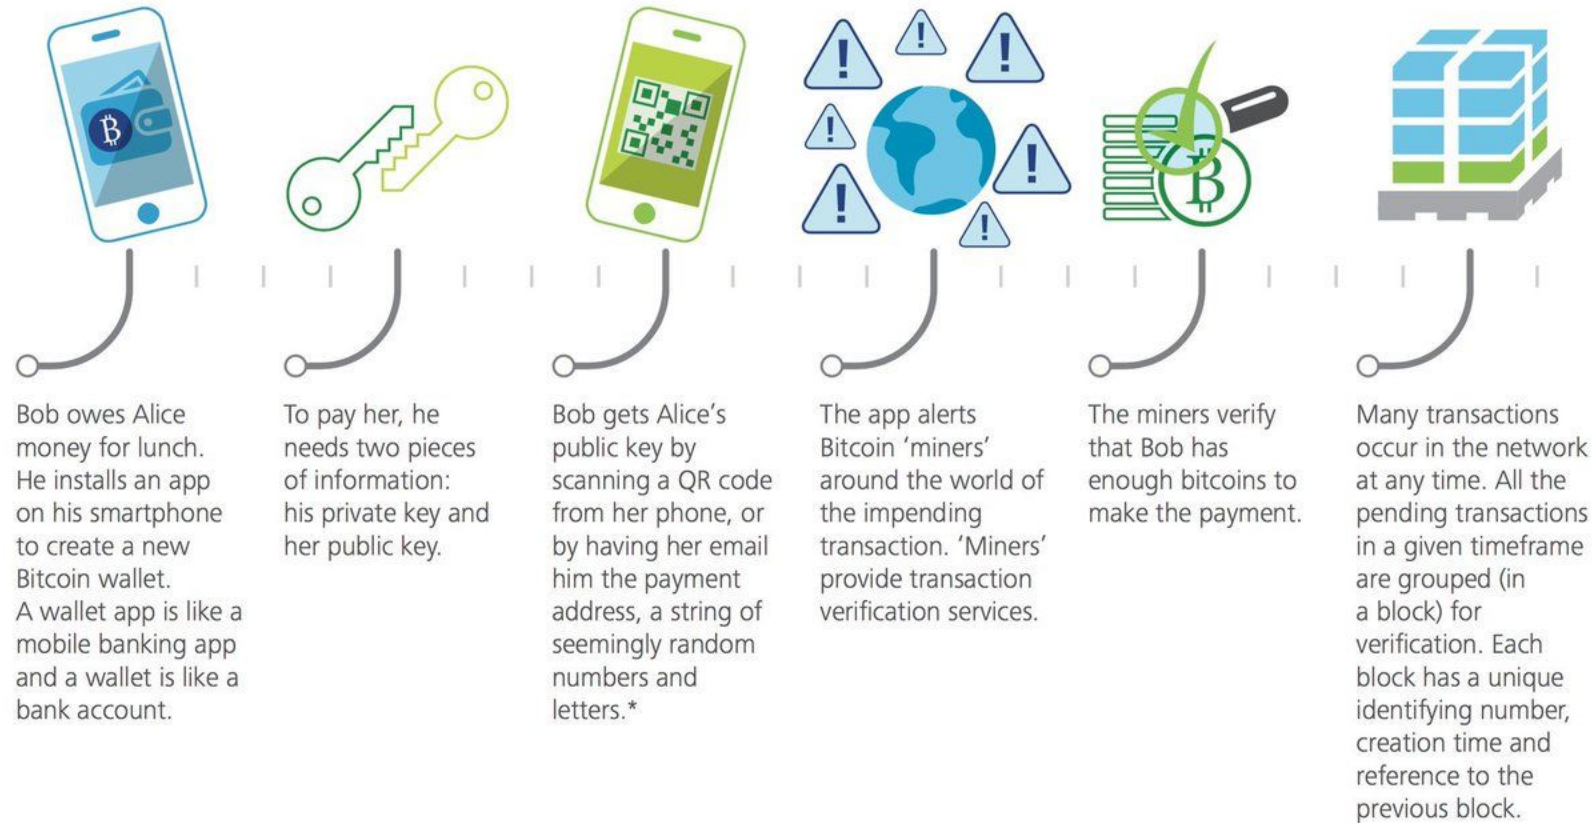
<https://cryptofundresearch.com/top-20-cities-blockchain-jobs-us/>

Horizon 2020 Existing EU Projects on Blockchain

- * D-CENT (social money for democratic societies)/EU-funding ended in May 2016- <https://dcentproject.eu/>
- * DECODE (decentralised management architecture)- <https://www.decodeproject.eu/>
- * MyHealthMyData (blockchain for health and patient-centric system)- <http://www.myhealthmydata.eu/>
- * Bloomend (blockchains for social media)- http://cordis.europa.eu/project/rcn/211092_en.html
- * SUnFISH- <http://www.sunfishproject.eu>
- * Symbiote- <https://www.symbiote-h2020.eu>
- * GHOST- http://cordis.europa.eu/project/rcn/210233_en.html
- * BlockchainKYC (Iceland)- http://cordis.europa.eu/project/rcn/211172_en.html
- * Signaturit (Spain)- http://cordis.europa.eu/project/rcn/205049_en.html
- * Billon (Poland)- http://cordis.europa.eu/project/rcn/212243_en.html
- * BROS- http://cordis.europa.eu/project/rcn/209037_en.html
- * DLInnociate- http://cordis.europa.eu/project/rcn/209748_en.html
- * DEFENDER- http://cordis.europa.eu/project/rcn/210231_en.html
- * TITANIUM- http://cordis.europa.eu/project/rcn/209948_en.html
- * INTERLACE- http://cordis.europa.eu/project/rcn/209089_en.html
- * STOP-IT- http://cordis.europa.eu/project/rcn/210216_en.html
- * CHARIOT- http://cordis.europa.eu/project/rcn/212490_en.html

Bitcoin Blockchain

Figure 1. How the Bitcoin blockchain works



*Anyone who has a public key can send money to a Bitcoin address, but only a signature generated by the private key can release money from it.

Graphic: Deloitte University Press. Source: American Banker²⁰

Bitcoin & Other Cryptocurrencies

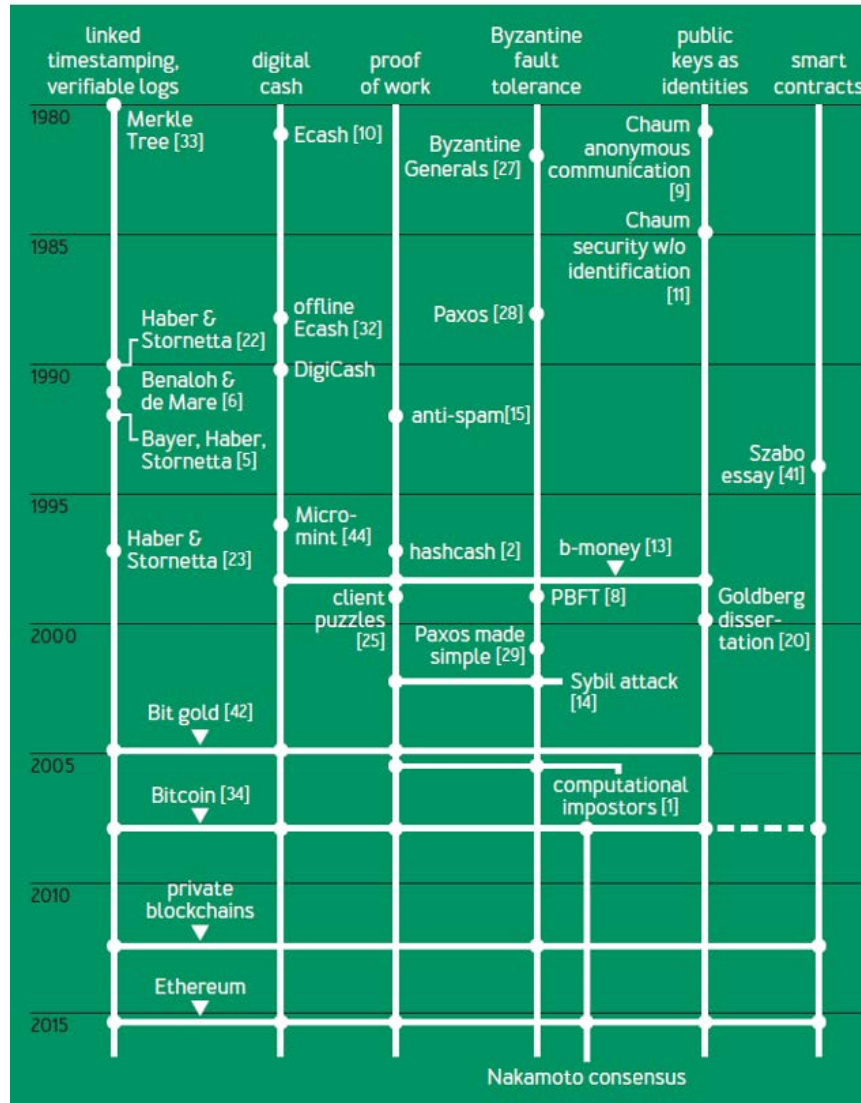
<http://bit.ly/HFpaper>

UTXO Cryptocurrencies

- **Unspent Transaction Output (UTXO)**: Data model introduced by Bitcoin - also used by many other cryptocurrencies and distributed applications (**DApps**)
- UTXO represents each step in the evolution of a data object as a separate atomic state on the ledger
- Such a state is created by a transaction and destroyed/consumed by another unique transaction occurring later
- Every given transaction destroys a number of **input states** and creates one or more **output states**
- A “coin” in Bitcoin is initially created by a **coinbase transaction** that rewards the “miner” of a block. This appears on the ledger as a coin state designating the miner as the owner.
- Any coin can be spent in the sense that the coin is assigned to a new owner by a transaction that atomically destroys the current coin state designating the previous owner and creates another coin state representing the new owner
- Value in the UTXO model is transferred through transactions that refer to several input states that all belong to the entity issuing the transaction
- An entity owns a state because the **public key** of the entity is contained in the state itself
- Every transaction creates one/more output states in the KVS representing the new owners, deletes the input states in the KVS, and ensures that the sum of the values in the input states equals the sum of the output states’ values
- There is also a policy determining how value is created (e.g., coinbase transactions in Bitcoin or specific mint operations in other systems) or destroyed

Bitcoin's Academic Pedigree

FIGURE 1: CHRONOLOGY OF KEY IDEAS FOUND IN BITCOIN




Arvind Narayanan, Jeremy Clark
ACM Queue, August 2017

Cryptoassets & ICOs



Pure-play ICOs are losing steam

Disclosed funding of completed ICOs. January 2017 – March 2018.



CBINSIGHTS

Source: TokenData

ICO MARKET IS LIKELY IN FOR A RUDE AWAKENING

"I have yet to see an ICO that doesn't have a sufficient number of hallmarks of a security."








Jay Clayton
Chairman, SEC
November 8, 2017

CBINSIGHTS

- CFTC considers cryptoassets to be commodity
- FinCEN as money
- SEC as security (Bitcoin/Ether excepted)

Common Blockchain Myths (McKinsey, 6/2018)

Five common blockchain myths create misconceptions about the advantages and limitations of the technology.

1		Myth	Reality
		Blockchain is Bitcoin	<ul style="list-style-type: none"> ● Bitcoin is just one cryptocurrency application of blockchain ● Blockchain technology can be used and configured for many other applications
2		Blockchain is better than traditional databases	<ul style="list-style-type: none"> ● Blockchain's advantages come with significant technical trade-offs that mean traditional databases often still perform better ● Blockchain is particularly valuable in low-trust environments where participants can't trade directly or lack an intermediary
3		Blockchain is immutable or tamper-proof	<ul style="list-style-type: none"> ● Blockchain data structure is append only, so data can't be removed ● Blockchain could be tampered with if >50% of the network-computing power is controlled and all previous transactions are rewritten—which is largely impractical
4		Blockchain is 100% secure	<ul style="list-style-type: none"> ● Blockchain uses immutable data structures, such as protected cryptography ● Overall blockchain system security depends on the adjacent applications—which have been attacked and breached
5		Blockchain is a "truth machine"	<ul style="list-style-type: none"> ● Blockchain can verify all transactions and data entirely contained on and native to blockchain (eg, Bitcoin) ● Blockchain cannot assess whether an external input is accurate or "truthful"—this applies to all off-chain assets and data digitally represented on blockchain

McKinsey&Company

Distributed Systems


- Distributed operating systems
- Distributed virtual memory
- Message passing in distributed computations and distributed checkpoints
- Clock synchronization and event ordering (e.g., Lamport clocks)
- Byzantine agreement and distributed consensus
- Two phase commit optimizations (e.g., Presumed Abort)
- (Transactional) RPCs and distributed file/object systems
- Asynchronous computation via message queues and pub-sub
- Distributed event-based systems
- Client-server, mobile computing and caching, WWW
- Workflow or business process management systems
- Service Oriented Architecture (SOA)
- Public cloud and hybrid cloud
- ...

Data Systems

- Relational DBMSs (e.g., **System R**) and SQL
- Data consistency, degrees of isolation and fault tolerance
- Distributed databases (e.g., **R***) and distributed transactions/queries
- **Synchronous and asynchronous replication with primary copy**
- Update anywhere (multi-master) replication and eventual consistency
- Stored procedures, user-defined types/functions, data provenance, ...
- Data warehousing and parallel DBMSs – OLTP vs OLAP
- Shared Nothing Vs Shared Disks
- Object-oriented databases, XML, schema chaos, data integration, ...
- Web2.0-inspired NoSQL, sharding & massive scaling (e.g., **Spanner**), JSON, ...
- Big Data: Map-Reduce, Hadoop, Spark, ...
- Data privacy, multitenancy and trans-border data flow restrictions
- Multi data centers and disaster recovery
- ...

Problem Being Solved (Export Import Scenario)

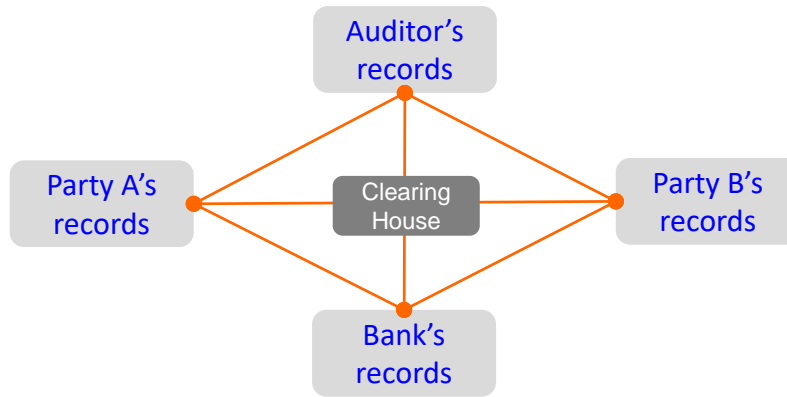
Recording of events is becoming much more complex...



... Inefficient, expensive, vulnerable, lack of transparency


Basic Change to Business Processes

Traditional Way



... Inefficient, expensive, vulnerable

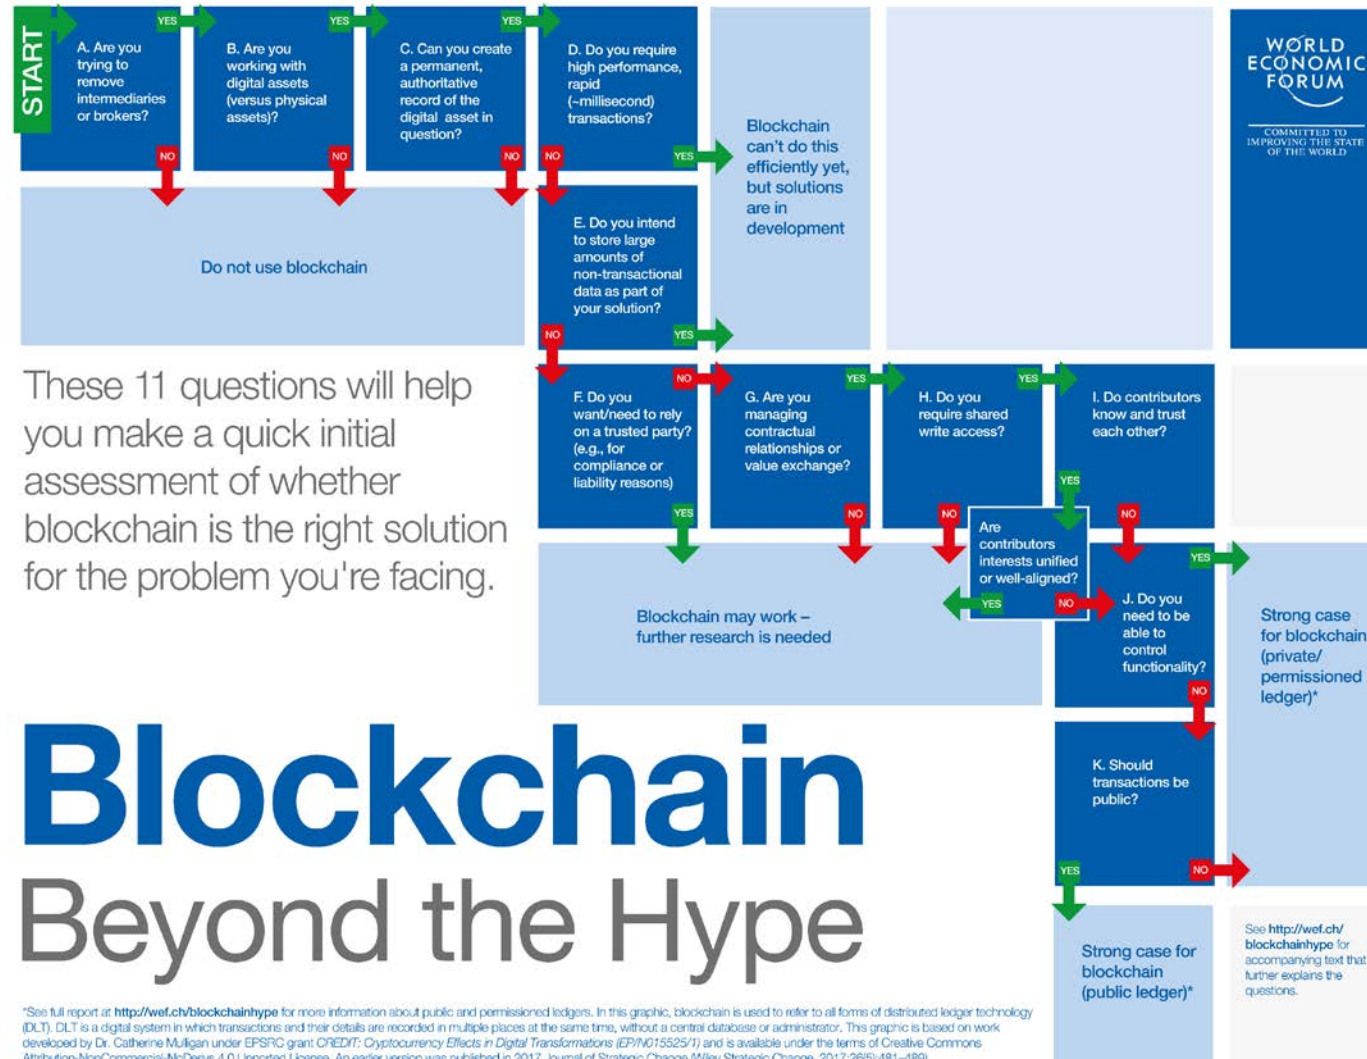
Blockchain Way



... Consensus, provenance, immutability, finality

Which Use Case Needs Blockchain?

World Economic Forum 4/18



Notes: Incorrect Recommendation about use of blockchains for managing physical assets

These 11 questions will help you make a quick initial assessment of whether blockchain is the right solution for the problem you're facing.

*See full report at <http://wef.ch/blockchainhype> for more information about public and permissioned ledgers. In this graphic, blockchain is used to refer to all forms of distributed ledger technology (DLT). DLT is a digital system in which transactions and their details are recorded in multiple places at the same time, without a central database or administrator. This graphic is based on work developed by Dr. Catherine Mulligan under EPSRC grant CREDIT: Cryptocurrency Effects in Digital Transformations (EP/N015525/1) and is available under the terms of Creative Commons Attribution-NonCommercial-NoDerivs 4.0 Unported License. An earlier version was published in 2017 Journal of Strategic Change (Wiley Strategic Change, 2017;26(5):481–489).

Smart Contracts

Everest Group

Smart contracts: realizing true benefits of blockchain

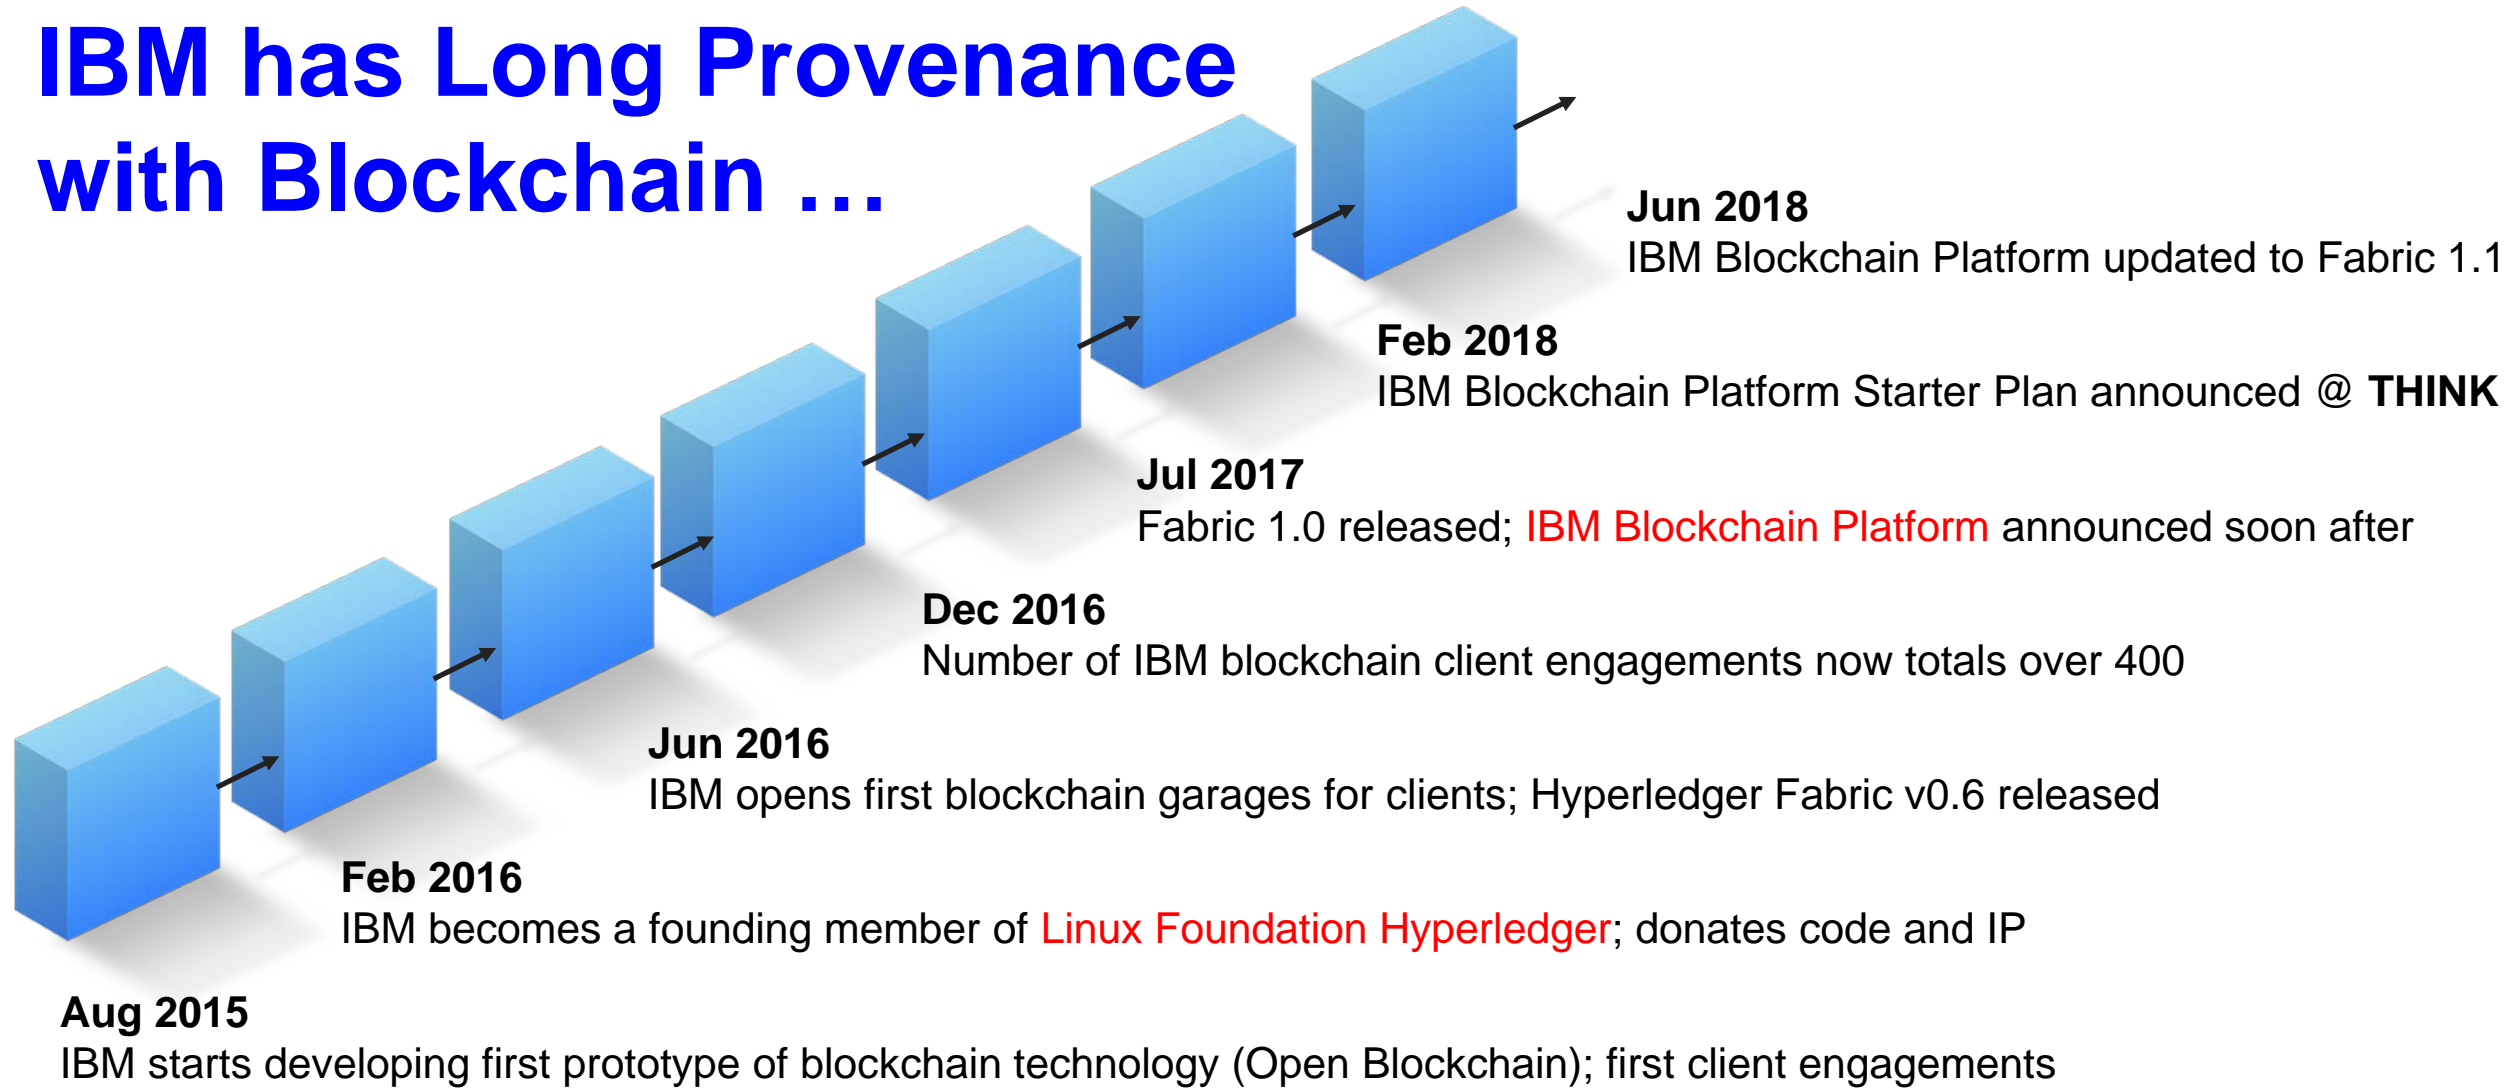
Blockchain is a cryptographic or encoded ledger (database) of transactions in the form of blocks arranged in a chain

Smart contract, a complex set of software codes with components designed to automate execution and settlement, is the application layer that makes much of the benefits of blockchain technology a reality



Everest Group Smart Contracts on Distributed Ledger – Life in the Smart Lane


IBM has Long Provenance with Blockchain ...



BaaS: IBM Blockchain Platform (IBP)

IBM Blockchain Platform is a fully integrated enterprise-ready blockchain platform designed to accelerate the development, governance, and operation of a multi-institution business network

- **Developer tools** that make use of Hyperledger Composer to quickly build your blockchain application
- Hyperledger Fabric provides the ledger; managed through a set of intuitive **operational tools**
- **Governance tools** for democratic management of the business network
- Flexible deployment options, including a highly secure and performant **IBM Cloud** environment



5/2018: IBM Introduces Crypto Anchor Verifier –
special lens added to mobile phone camera

Microscopic details of an object's surface are measured – e.g., optical characteristics such as shape, viscosity, saturation value, spectral values (**AI + optical imaging**)

Platform Value: *Simplicity in the face of overwhelming complexity*

	IBM Blockchain Platform	Community Code Deployment
Inviting members	5 seconds	20 minutes per instance
Installing and instantiating smart contracts	Single click installation	10 minutes per smart contract per peer
Deployment	Specify network parameters and automatically launch ordering service	Not available
Network alterations and additions	Add new members, channels and smart contracts through single clicks, text box or drop down via the UI	CLI driven, and more advanced skills required
Support	Complete support from the HW stack through the blockchain code base included	IBM support options available
Security	Secure container and highest level of security provided	Custom
Migration	Rolling migration and 99.999% availability provided under the covers	Not available

"IBM provides us with the easiest way to develop prototype blockchain applications for our clients. Thank you!"

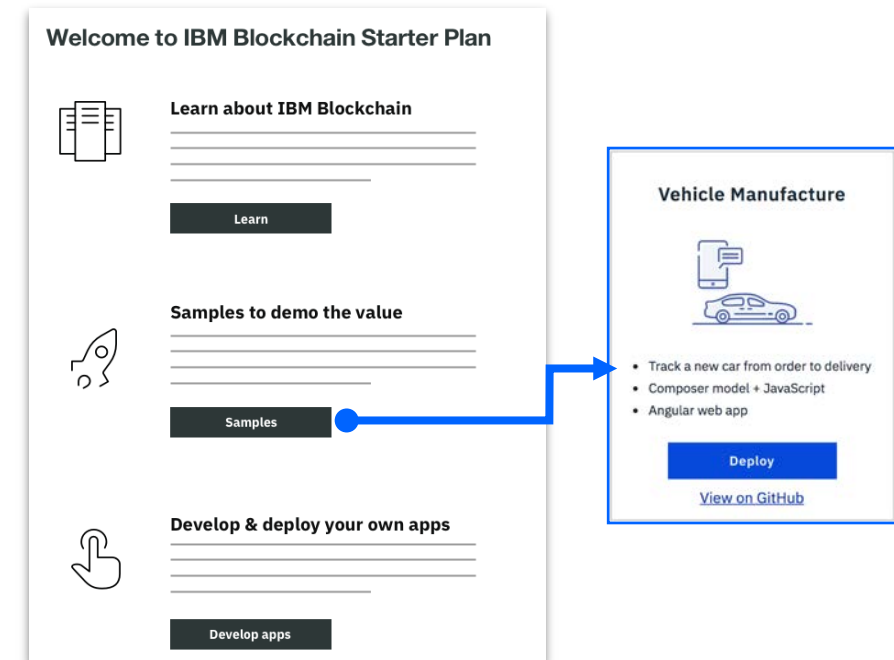
-- Global consulting firm

"IBM has enabled our team to develop our blockchain demo with minimal hassle and gives us a clear path to scale with the tools to manage it"

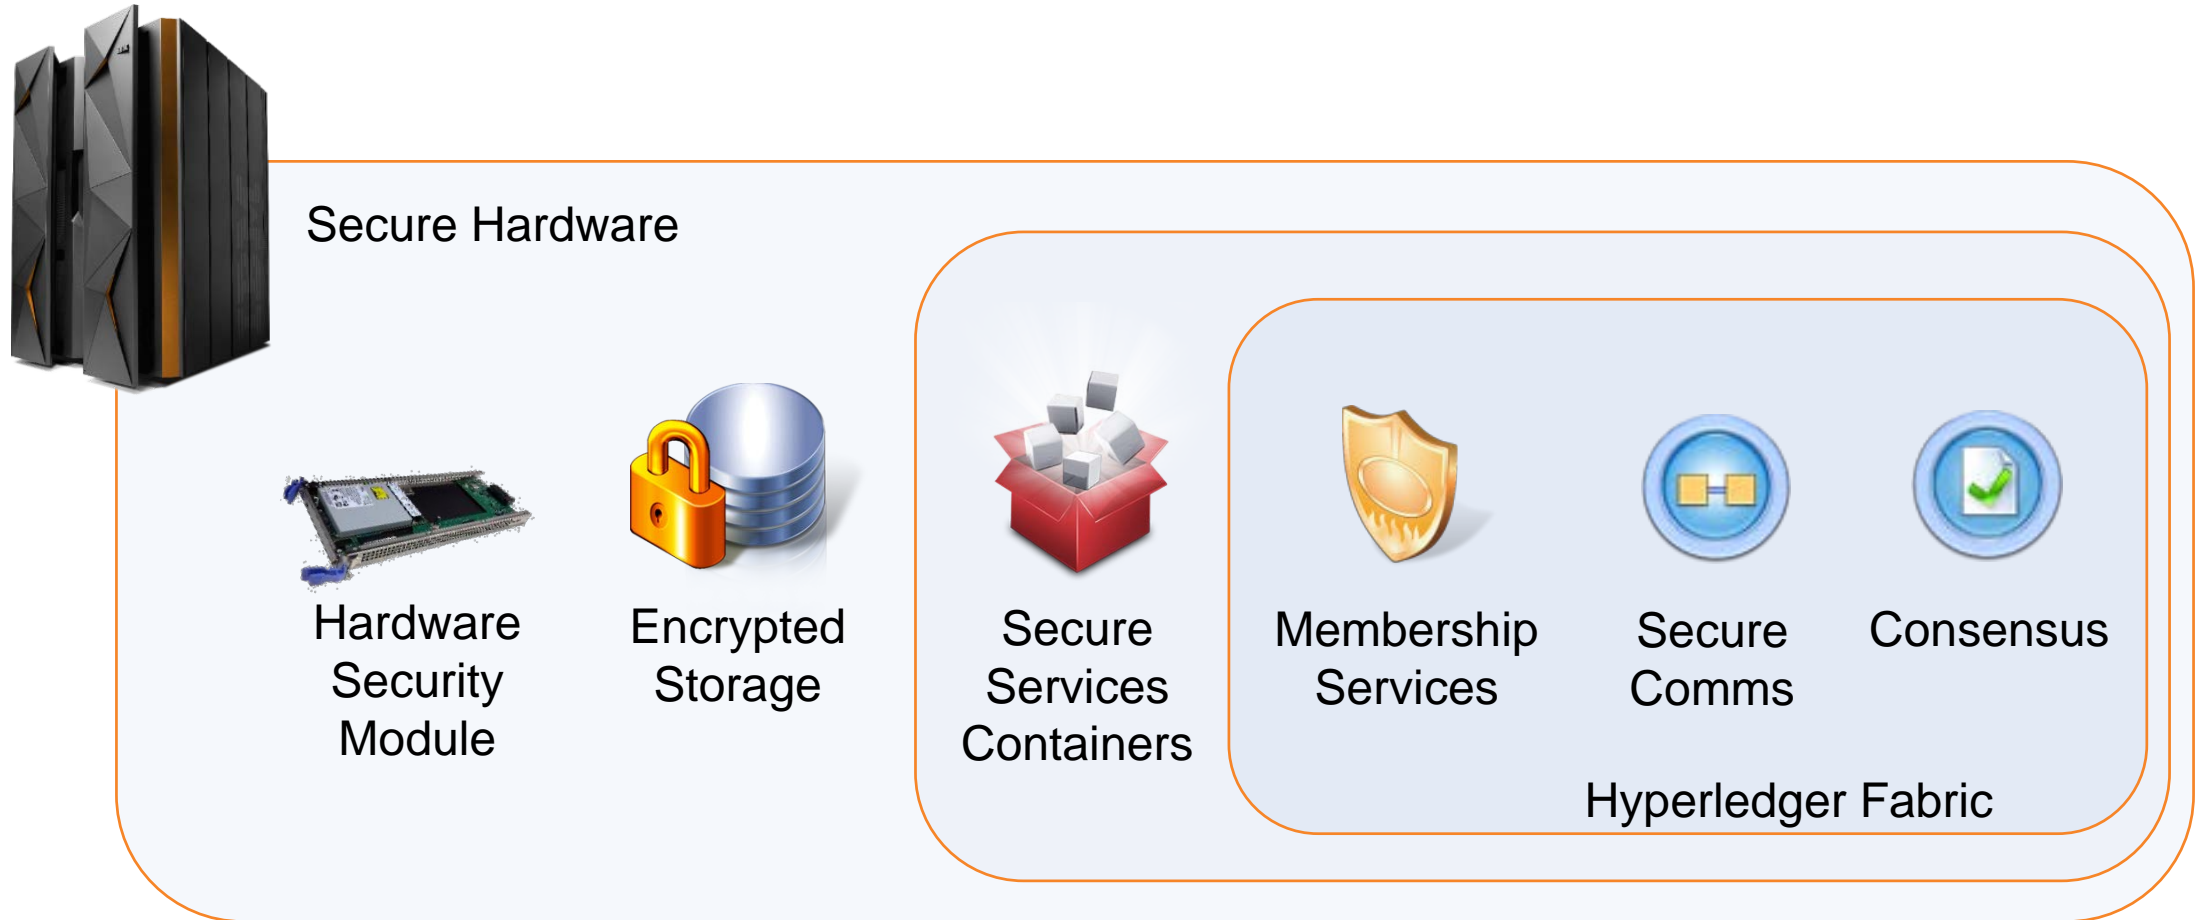
-- Series backed start-up

Starter Plan

- Get started with **IBM Blockchain Platform** with **one-click setup and a fully functional network**
 - Configured for two organizations with one peer each, sample applications and informational tutorials
 - Environment enables iterative development prior to production deployment
 - Same experience as Enterprise
 - Uses SOLO ordering for simplified configuration, development and testing
- Sign up for a **30 day free trial**
 - After that time, there is a monthly charge of \$250 membership fee per month, plus \$125 per peer



IBP: Security at Each Architecture Layer




Blockchain Technical Concepts (Hyperledger Fabric)



Peers are the networked services that maintain ledger state and run smart contracts



Channels are defined subsets of the peer network that share a single ledger




Certificate authorities provide identity services to participants on the network



Smart contracts constitute the transaction logic whose output determines changed asset states

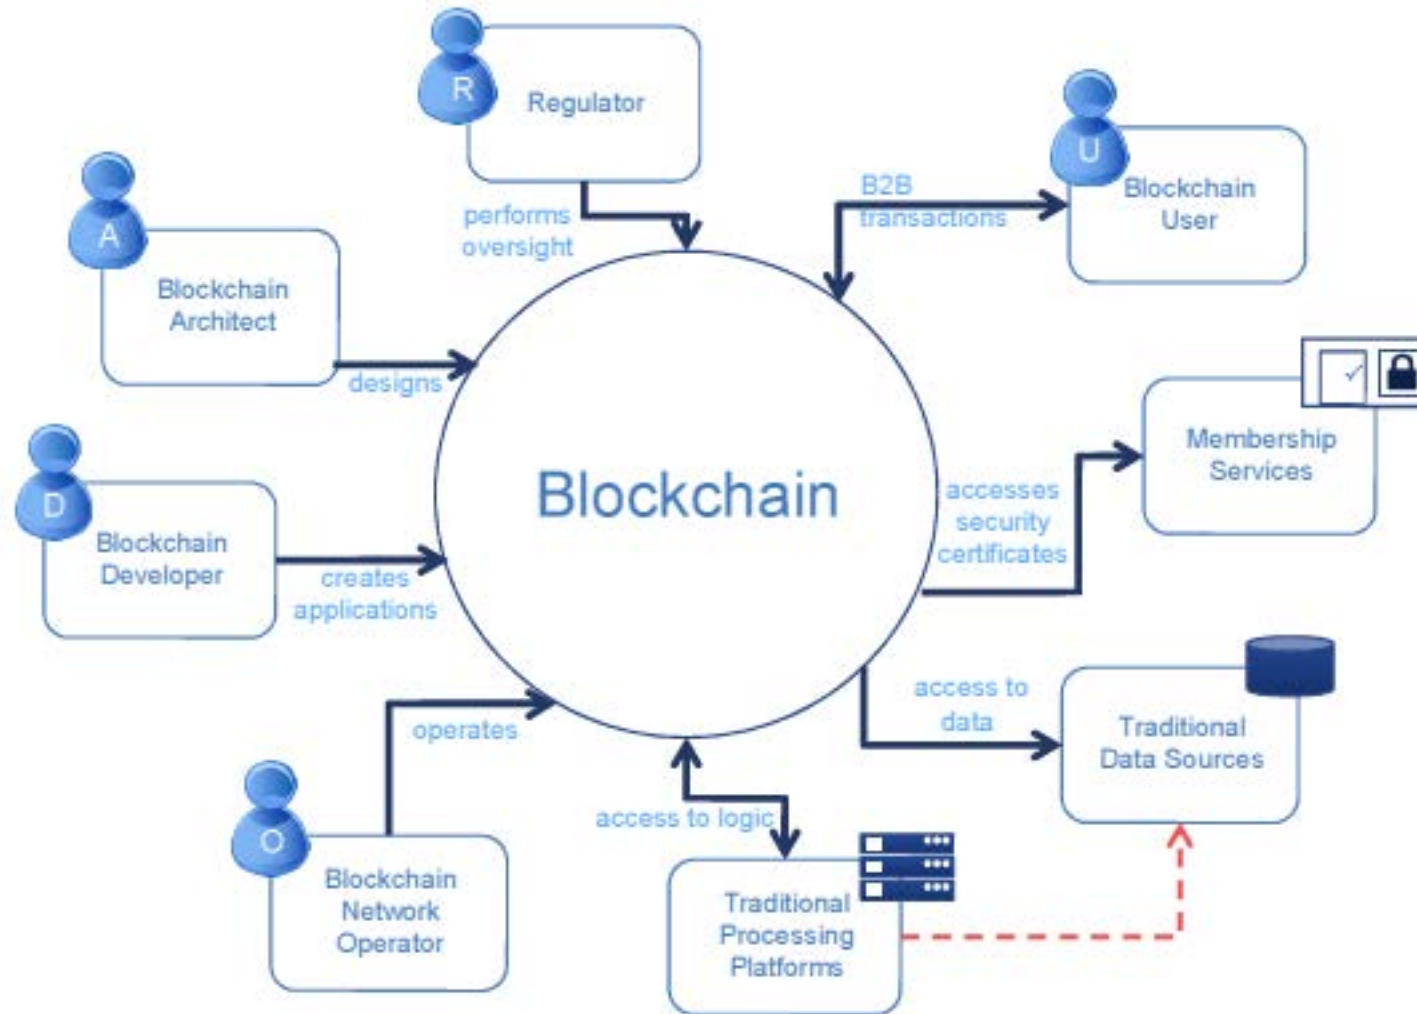


Consensus is the process by which agreement is obtained on the peer network

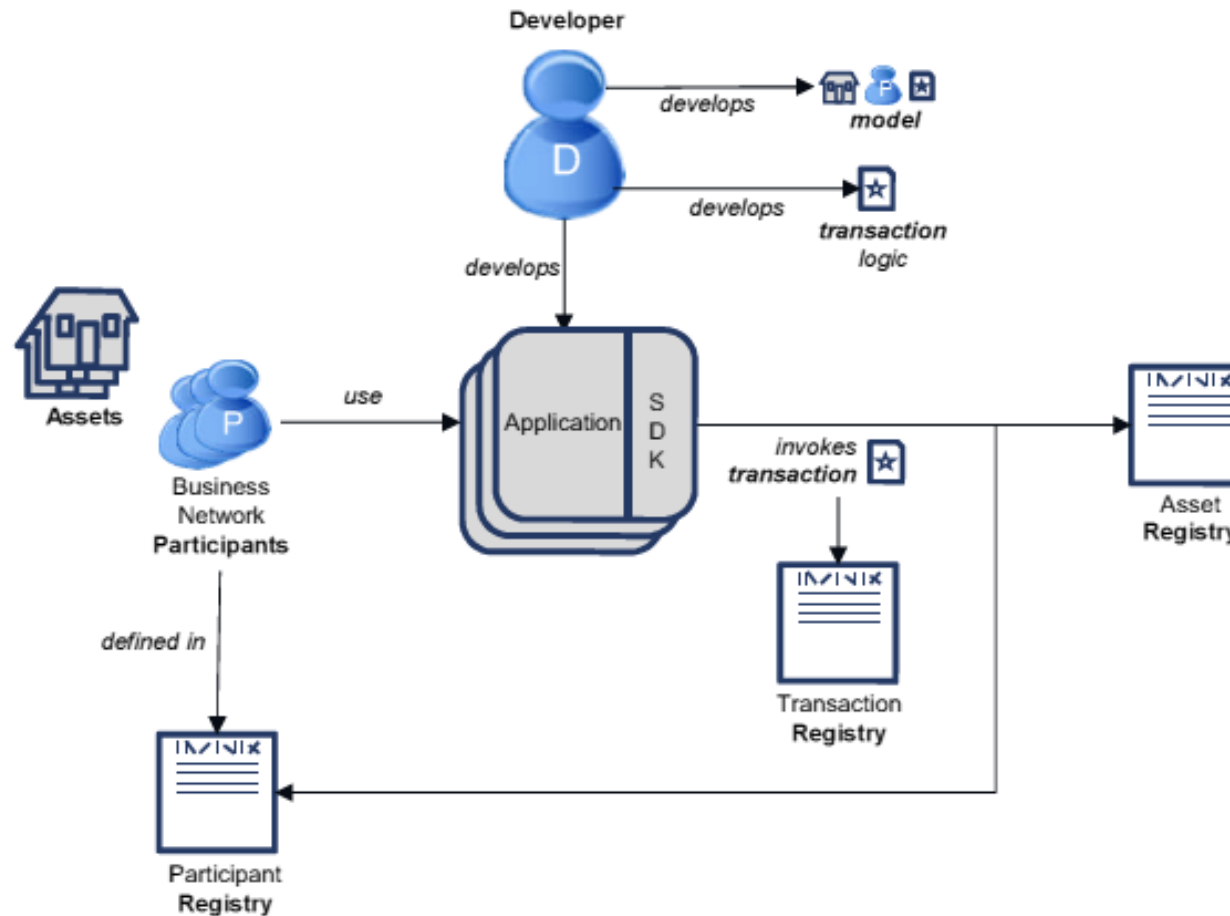


The **Ordering Service** decides transaction sequence and distributes blocks to peers

Actors in a Blockchain Solution




Composer: Workflow of Building a Model



Building Communities in Blockchain Networks

Consortium Based Network




Founders are equal among other participants, may include a joint legal entity among the founders (e.g. – JV)

Examples:



Founder Directed Network




Individual founder in a position to provide strong direction

Examples:



Community Based Network



Driven by industry standards bodies or existing non-blockchain network owners

Examples:



Blockchain Applications

- Track provenance, ownership, relationships & lineage of assets
- Supply Chain – Food Safety ([Walmart](#)), Logistics ([Maersk](#))
- Health Data Exchange ([FDA](#))
- Know Your Customer
- Derivatives Processing
- Trade/Channel Finance ([IGF](#))
- Trade Information Warehouse ([DTCC](#))
- Post-Trade Reconciliation/Settlement
- Private Equity Fund Management ([Unigestion](#))
- Syndicated Loans
- Diamond/Valuables Tracking and Protection – Provenance Management ([Everledger](#))
- Cross-Border Payment, Payments for/by Unbanked Populations
- Low volume stock trading ([JPX](#))

“Blockchain Revolution: How the Technology Behind Bitcoin Is Changing Money, Business, and the World”

Don Tapscott, Alex Tapscott, ISBN 978-1101980132.

IBM Blockchain Engagements


Making blockchain real for business with more than 400 engagements and multiple active networks

Trade Finance	Pre and Post Trade	Complex Risk Coverage
   	   	 
Identity/ Know your customer (KYC)	Unlisted Securities/ Private Equity Funds	Loyalty Program
  	  	
Medicated Health Data Exchange	Fraud/ Compliance Registry	Distributed Energy/ Carbon Credit
		 
Supply Chain	Food Safety	Provenance/ Traceability
 	         	

Opportunities by Industrial Sector (McKinsey, 6/18)

Granular assessments at the use-case level are necessary to determine which blockchain opportunities to pursue.

Blockchain opportunities by industrial sector



Food Trust/Safety

What?

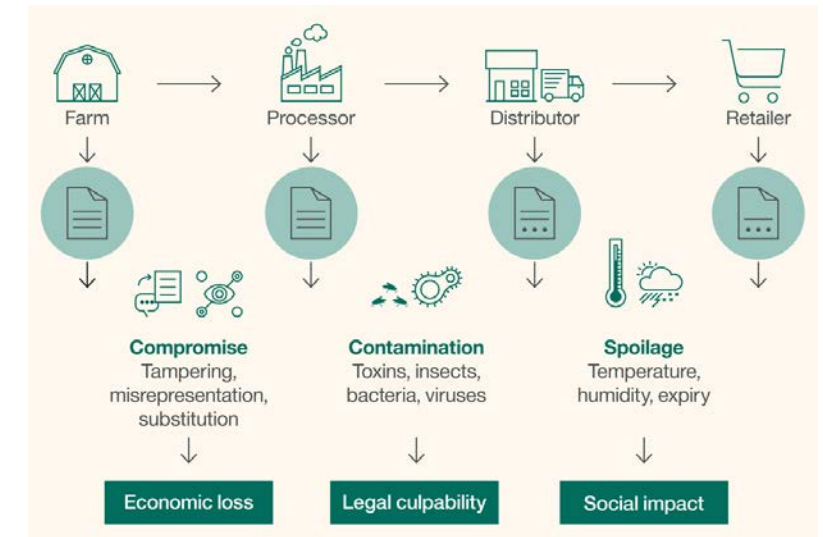
- Provide trusted source of information and traceability to improve transparency and efficiency across food network

How?

- Shared ledger for storing digital compliance documentation, test results and audit certificates network

Benefits

- Reduce impact of food recalls through instant access to end-to-end traceability data to verify history in food network & supply chain
- Help address 1 in 10 people sickened and 400K fatalities worldwide which occur every year from food-born illnesses



Global Trade Digitization (GTD)



What?

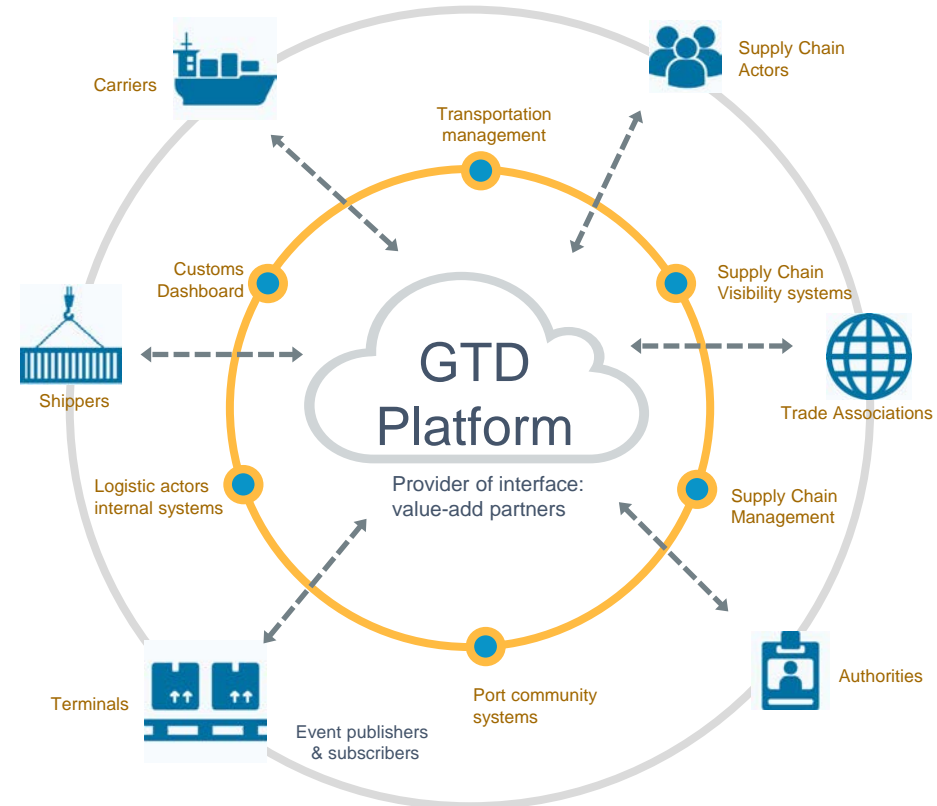
- An open, extensible platform for sharing shipping events, messages, and documents across all the actors and systems in the supply chain ecosystem.

How?

- Providing Shared Visibility and Shared State for Container Shipments

Benefits

- Increase speed and transparency for cross border transactions through real time access to container events.
- Reduced cost and increased efficiency through paperless trade



Dubai Blockchain Strategy

- Aims for Dubai to become the first blockchain-powered city by 2020
- For Dubai government to become paperless by shifting all transactions to Blockchain, and empower Dubai Smart city experience for all
- Based on Three pillars:
 - ✓ **Government Efficiency:** implementing blockchain technology in government services
 - ✓ **Industry Creation:** supporting the creation of a blockchain industry through empowering start ups and businesses
 - ✓ **International Leadership:** leading global thinking on blockchain technology
- the Smart Dubai Office SDO launched Blockchain Challenge in partnership with global accelerator 1776
 - aims to identify the most innovative blockchain ideas from startups around the world and bring them to Dubai
- SDO launched a city-wide effort to implement blockchain in city services
- Partnerships with IBM as a Blockchain Lead Strategic Partner, and Consensys as Blockchain City Advisor.

Dubai launches Blockchain strategy to become paperless by 2020

Hamdan unveils ambitious plan to save 25 million work hours annually through paperless transactions

Published: 20 OCTOBER 11, 2016

GULF NEWS



Source: Saeed Al Dhaheri, Etisalat Academy 3/2017

IBM-FDA Partnership

What?


- Create and promote secure, efficient and scalable exchange of health data using blockchains

How?

- Create an electronic ledger of where and how data is transferred and exchanged
- Initial trial focus on oncology data

Benefits

- Creating an audit trail through the ledger, healthcare professionals will be able to:
 - hold information leakers accountable
 - maintain transparency in what data is going where
 - secure weak spots in sharing process



IBM bets on the blockchain to keep your medical data safe

Big Blue believes the secure transfer of medical information can be achieved through technology associated with Bitcoin.



IBM has announced a new partnership with the US Food and Drug Administration (FDA) in a study designed to determine whether blockchain technology can be used to keep medical data transfers safe from theft or exploit.

Blockchain Companies/Consortia & Banks

New kid on the block

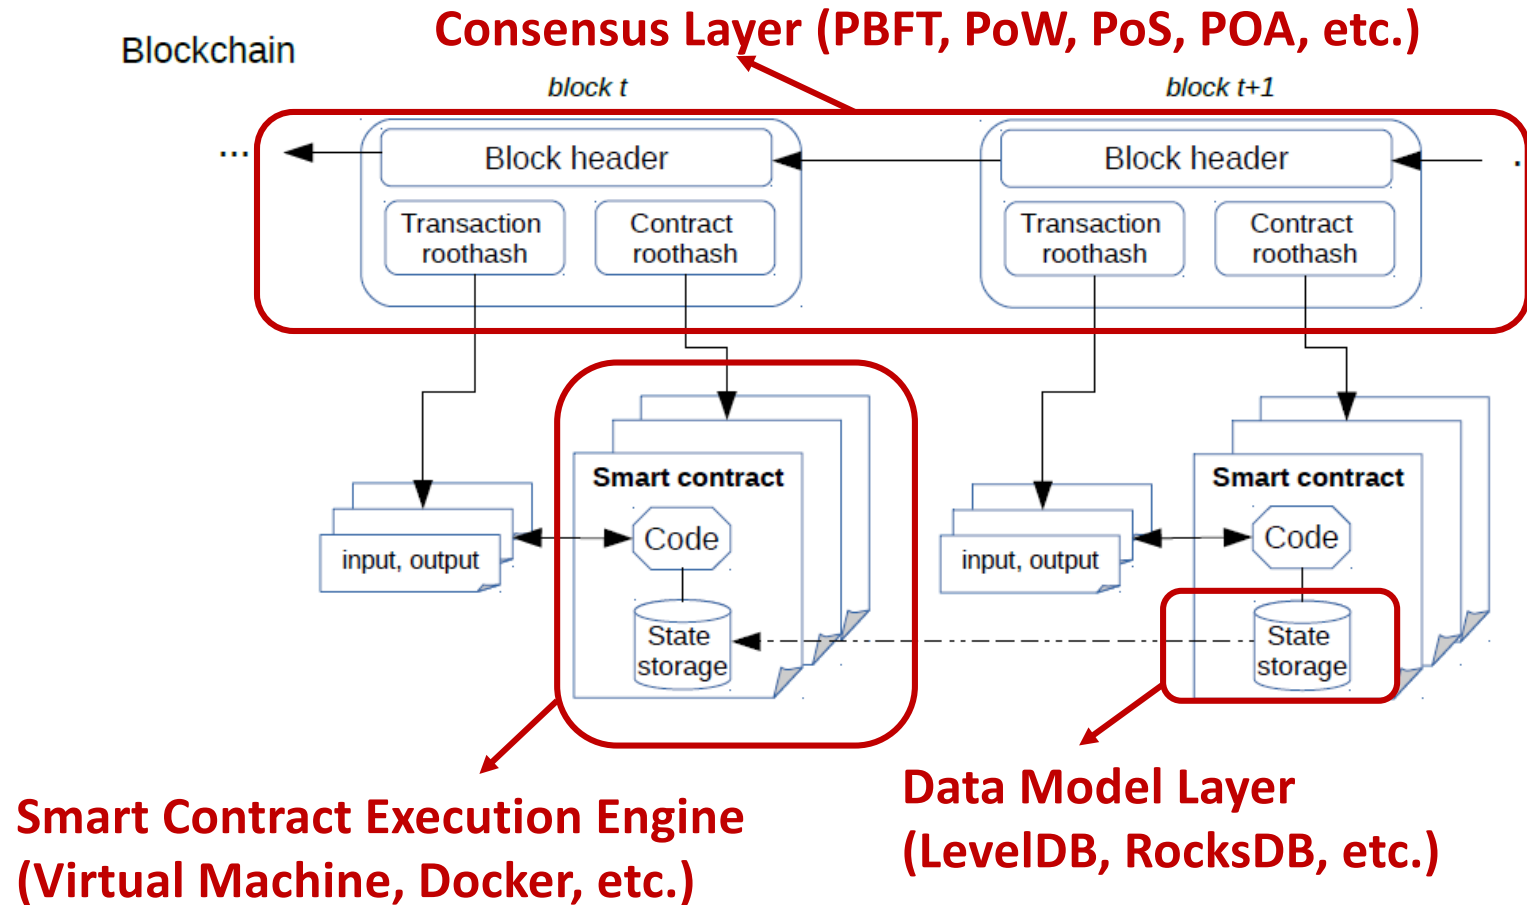
Enterprise Ethereum Alliance is the latest addition to a list of companies and consortiums focused on blockchain where banks serve as investors or partners

	No. of partners
Utility Settlement Coin	5
Global Payments Steering Group	6
Chain	9
Digital Asset Holdings	15
Enterprise Ethereum Alliance	30
R3	81
Ripple	90+
Hyperledger	122

Source: Staff research

Source: American Banker 2-2017

BC Software Stack



Source: Anh Dinh, et al., SIGMOD 2017

Blockchain Architecture/Feature Choices

- Cryptocurrencies Vs Generalized Assets
- Permissionless/Public Vs Permissioned/Private
- Byzantine Vs Non-Byzantine fault model
- Consensus approach: PoW, PoA, PoET, PBFT, ...
- SQL Vs NoSQL data stores
- Transactional stores Vs Non-transactional stores
- Versioned/Unversioned state database
- On-Chain Vs Off-Chain data
- Parallelism exploitation during different phases of transaction execution
- Pluggable features: consensus protocol, state DB, smart contract language, ...

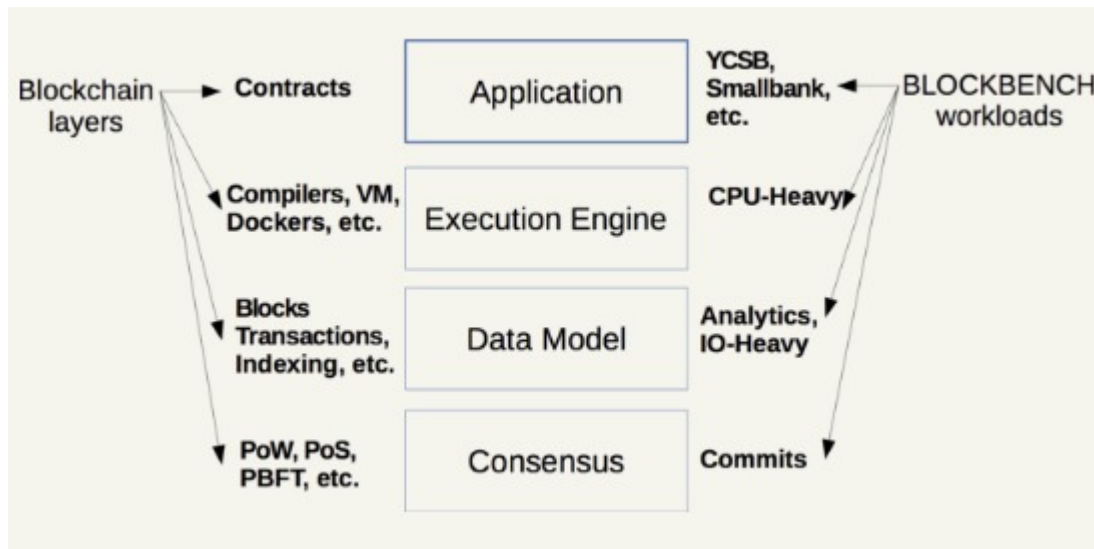
Good Survey Paper: [Untangling Blockchain: A Data Processing View of Blockchain Systems](#), A. Dinh et al.

Database Replication

- Primary **log replay** at replica – homogeneous systems with full DB replicas, typically done for disaster recovery (DR) backup
- Log **capture generates DML statements from what is logged** and **apply** executes those statements (e.g., IBM Q Replication)
 - Can handle non-determinism and partial replicas
 - Requires dependency analysis to leverage parallelism at apply time
 - <https://www.ibm.com/developerworks/data/roadmaps/qrepl-roadmap.html>
- Capture DML statements **as issued by application** and re-execute them at replica (e.g., H-Store/VoltDB)
 - Cannot handle non-determinism
 - Typically, serial execution of transactions

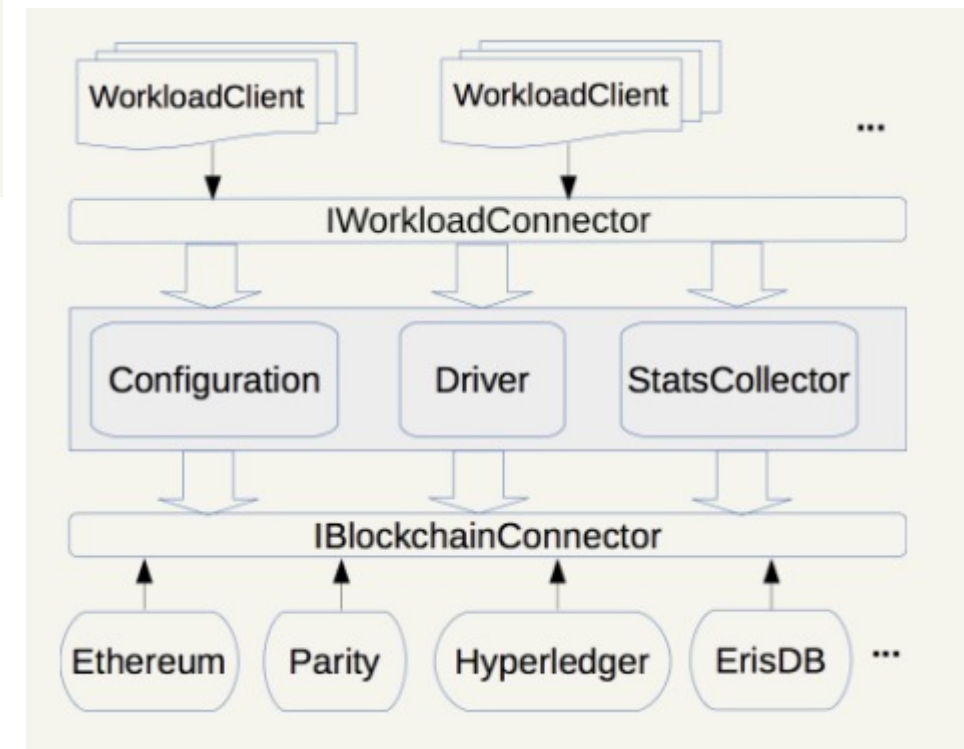
Upfront (fairly random, unoptimized) ordering of transactions in blockchain systems – leads to all sorts of issues!

Benchmark Framework: BLOCKBENCH (NUS)



- OLTP & OLAP load measured
- Metrics: throughput, latency, scalability, fault tolerance, security
- Consensus methods: Ethereum (PoW), Fabric (PBFT), Parity (PoA)
- Old version of Fabric (pre-V1)
- Fabric performs better
- Fabric scales well up to 16 nodes

Source: Anh Dinh, et al., SIGMOD 2017



Hyperledger Caliper

- Allows users to measure performance of a specific blockchain implementation with a set of predefined use cases
- Will produce reports containing a number of performance indicators, such as TPS (Transactions Per Second), transaction latency, resource utilization, ...
- Intent is for Caliper results to be used by other Hyperledger projects as they build out their frameworks, and as a reference in supporting the choice of a blockchain implementation suitable for a user's specific needs
- Initial contributors: Developers from Huawei, Hyperchain, Oracle, Bitwise, Soramitsu, IBM and Budapest University of Technology and Economics
- <https://www.hyperledger.org/projects/caliper>

Ethereum, Fabric, Corda Comparison

Table 1
Comparison of Ethereum, Hyperledger Fabric and Corda

Characteristic	Ethereum	Hyperledger Fabric	R3 Corda
Description of platform	– Generic blockchain platform	– Modular blockchain platform	– Specialized distributed ledger platform for financial industry
Governance	– Ethereum developers	– Linux Foundation	– R3
Mode of operation	– Permissionless, public or private*	– Permissioned, private	– Permissioned, private
Consensus	– Mining based on proof-of-work (PoW) – Ledger level	– Broad understanding of consensus that allows multiple approaches – Transaction level	– Specific understanding of consensus (i.e., notary nodes) – Transaction level
Smart contracts	– Smart contract code (e.g., Solidity)	– Smart contract code (e.g., Go, Java)	– Smart contract code (e.g., Kotlin, Java) – Smart legal contract (legal prose)
Currency	– Ether – Tokens via smart contract	– None – Currency and tokens via chaincode	– None

FSBC Working Paper, 6/2017

Ethereum

- Public blockchain system like Bitcoin
 - Extends it with Smart Contracts
 - Uses PoW for consensus
 - Own machine lang & VM
 - *gas* charging!
- Most apps relate to its currency Ether
- *Enterprise Ethereum Alliance (EEA)*: JPMorgan Chase, Microsoft, Intel, Accenture, Banco Santander, BNY Mellon, ConsenSys, Credit Suisse, ING, Thomson Reuters, UBS, Wipro
 - EEA will add confidentiality (Quorum), scalability (pluggable consensus) and permissioning to Ethereum
 - Focus on specification, **EntEth** 1.0 with Python reference client, benchmarking, compliance testing and tools
 - Develop standards for Ethereum: best practices, security, privacy, scalability, interoperability
- **Quorum** from JPMorgan; Support for PBFT added in 7/2017 by AMIS

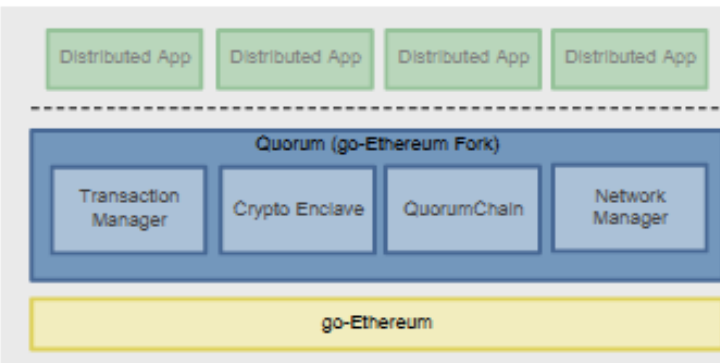
Quorum Overview J.P. Morgan

Highlights

- **Built on Ethereum**
 - First mover advantage. In production since July, 2015.
 - 50,000+ unit tests, Security Audits, Bounty Program
 - Largest Ecosystem of Developers, Tools, DApp's
 - Public Ethereum blockchain protects over \$1B+ Ether¹
- **Simple Privacy Design**
 - Supports both private and public transactions and smart contracts
- **Single Blockchain Architecture**
 - All public and private smart contracts and state derived from a single, common, complete blockchain of transactions validated by every node in the network
 - Private smart contract state validated by parties to contract only
 - Best of both worlds... every node validating the list of transactions while only exposing details of private transactions and contracts to relevant parties
- **High Performance**
 - Able to process **dozens to hundreds of transactions per second**, depending on system configuration; enough to support institutional volumes

¹As of 22-Sep-2016

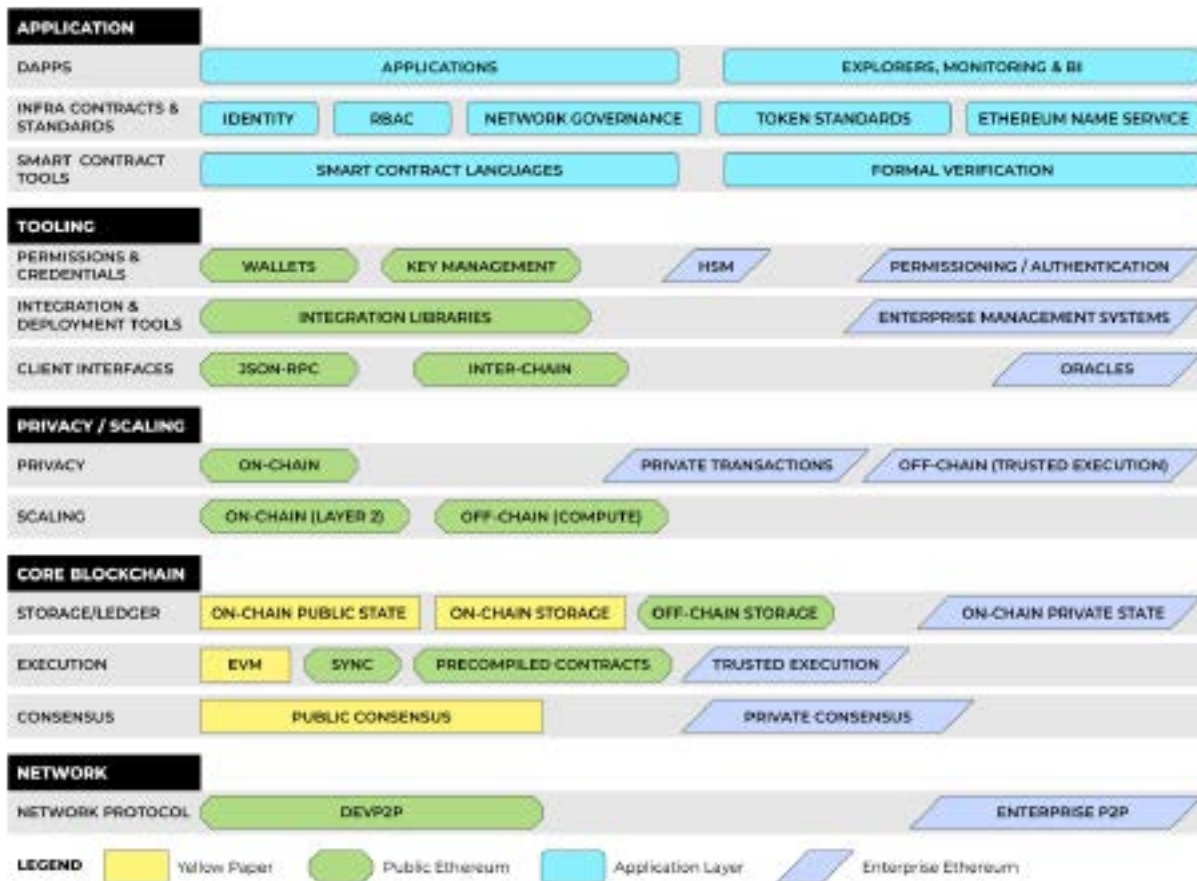
Architecture



Components

- **Transaction Manager** – allows access to encrypted transaction data for private transactions, manages local data store and communication with other Transaction Managers
- **Crypto Enclave** – responsible for private key management and encryption and decryption of private transaction data
- **QuorumChain** – voting-based, BFT-hardened consensus mechanism that utilises core Ethereum features to verify and propagate votes through the network
- **Network Manager** – controls access to the network, enabling a permissioned network to be created

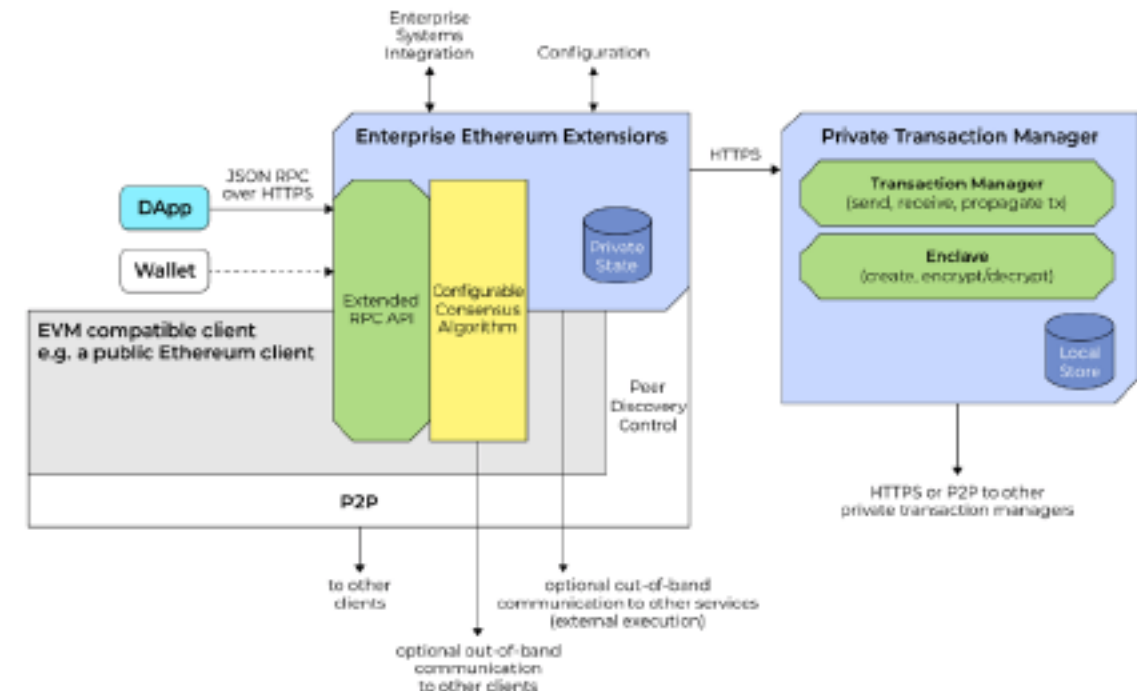
EEA Client Specification & Sample Architecture (5/2018)



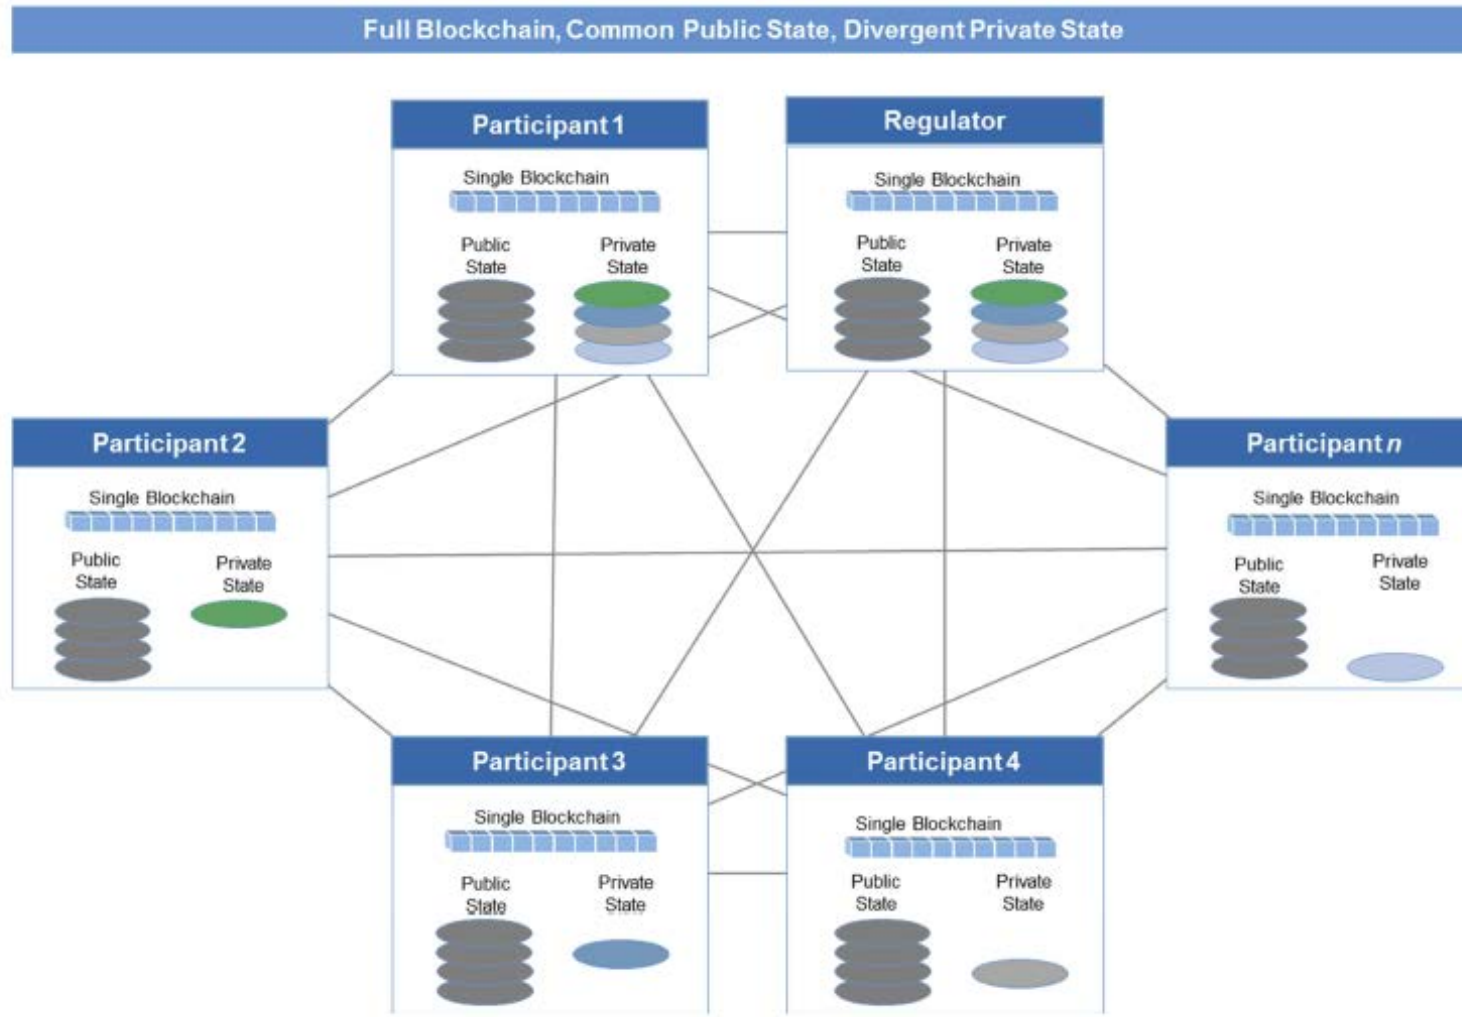
All Yellow Paper, Public Ethereum, and Application Layer components may be extended for Enterprise Ethereum as required.

© 2018 Enterprise Ethereum Alliance

ENTERPRISE ETHEREUM HIGH LEVEL ARCHITECTURE




Quorum Network J.P. Morgan



Hyperledger: A Linux Foundation project

- A collaborative effort created to advance cross-industry blockchain technologies for business
- Founded February 2016; now more than 230 member organizations
- Open source, open standards, open governance
- Five frameworks and five tools projects
- IBM is a premier member of Hyperledger

Hyperledger Momentum



www.hyperledger.org

Hyperledger Fabric Roadmap

V1 Alpha

- Docker images
- Tooling to bootstrap network
- Fabric CA or bring your own
- Java and Node.js SDKs
- Ordering Services - Solo and Kafka
- Endorsement policy
- Level DB and Couch DB
- Block dissemination across peers via Gossip

V1 GA

- Hardening, usability, serviceability, load, operability and stress test
- Chaincode ACL
- Chaincode packaging & LCI
- Pluggable crypto
- HSM support
- Consumability of configuration
- Next gen bootstrap tool (config update)
- Config transaction lifecycle
- Eventing security
- Cross Channel Query
- Peer management APIs
- Documentation

V1.1

- Node.js smart contracts
- Node.js connection profile
- Smart Contract APIs:
 - Encryption library
 - Txn submitter identity
 - Access control (using above)
- Performance & Scale
 - More orderers at scale
 - Parallel txn validation
 - CouchDB indexes
- Events
 - Per channel vs global
 - Block info minimal events
- CSR for more secure certs
- Serviceability
 - Upgrade from 1.0
- **Technical Preview features**
 - Private channel data
 - Finer grained access control on channels (beyond orgs)
 - ZKP features (ID Mixer)
 - Java for Smart contracts

V1.2

- V1.1 Technical Preview features
 - Finalize Side DB - Private Data
 - Finalize Java chaincode
 - Finalize Fabric ACL mechanism
- Usability Features
 - e.g. Service discovery
- Technical Debt/Hygiene
 - e.g. testing frameworks
 - Parallel testing
 - More modular code
- Pluggable endorsement and validation
- State-based Endorsement
- Privacy-preserving state-based endorsement

Based on <https://wiki.hyperledger.org/projects/fabric/roadmap>

March 2017

July 2017

March 2018

June 2018 (quarterly)

* Dates determined by the Hyperledger community, subject to change

Hyperledger Fabric Project

- Initiated by IBM with IBM open source ledger contribution (Feb 2016)
<http://hyperledger-fabric.readthedocs.io/en/latest/>
- Significant change in architecture from V0.6 to V1
 - Chaincode trust flexibility
 - **Channel** concept for Scalability & Confidentiality enhancement
 - Consensus modularity
 - Pluggable State DB APIs
 - 2 types of peer nodes: Endorsing, non-endorsing/committing
- Used PBFT for consensus before V1
- Other Hyperledger Projects: Iroha, Sawtooth, Composer, Quilt, ...

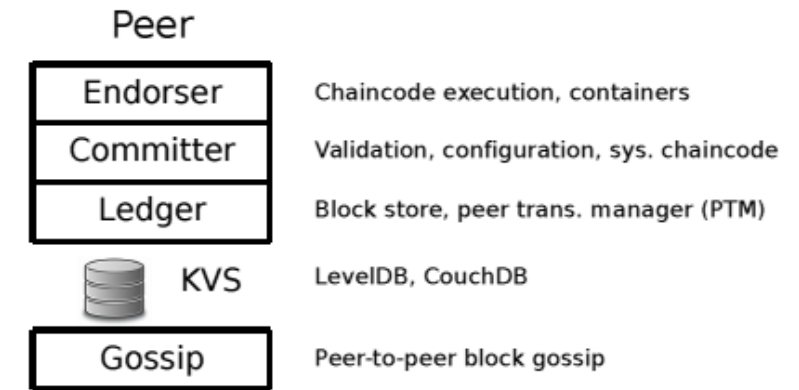
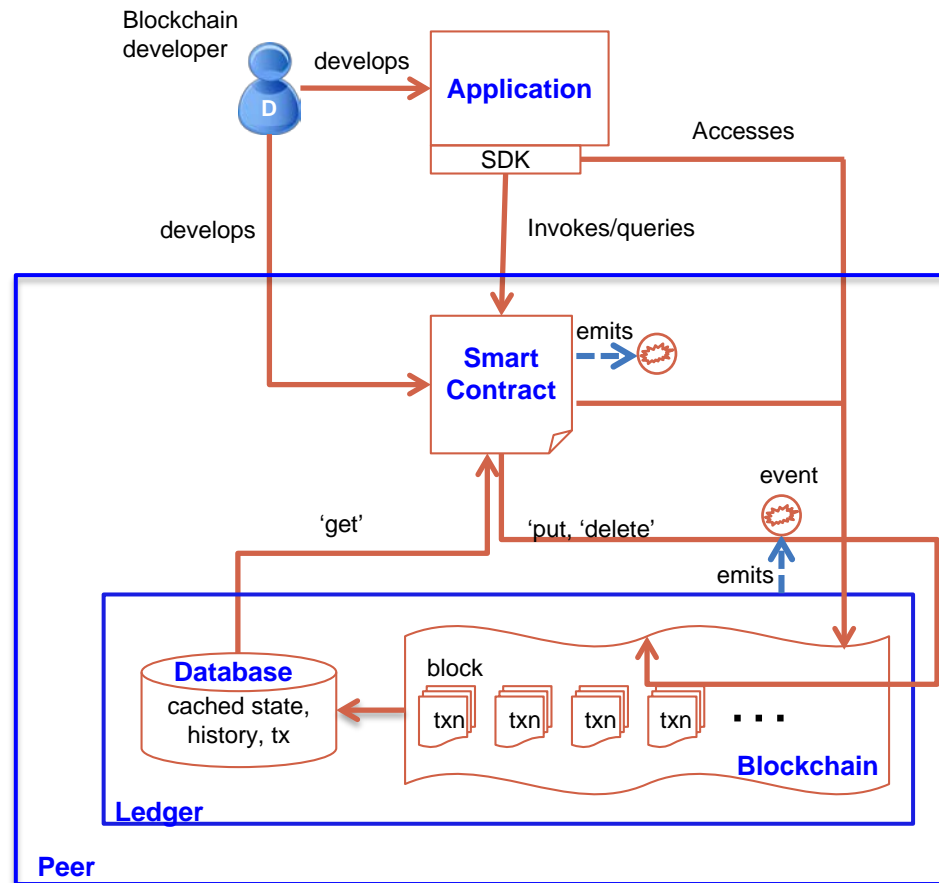


Figure 5: Components of a Fabric peer.

Hyperledger **Premier** members include: Accenture, Airbus, American Express, Baidu, Change Healthcare, Cisco, CME Group, Deutsche Bank, Deutsche Borse Group, Daimler, Digital Asset, DTCC, Fujitsu, Hitachi, IBM, Intel, J.P. Morgan, NEC, R3, SAP, Tradeshift and Wanda FFan Technology

Hyperledger Fabric V1 Contributors - Engineers from: Arxan, Cloudsoft, CLS, d20 Technical Services, Depository Trust & Clearing Corporation (DTCC), Digital Asset, Fujitsu, GE, Gemalto, HACERA, Hitachi, Huawei Technologies, Hyperchain, ImpactChoice, IT People, Knoldus, Linux Foundation, Netease, Passkit, State Street Bank, SecureKey, IBM, SAP, Thoughtworks and Wanda Group. There were also contributions from 35 unaffiliated individuals. In total, 159 developers have contributed.

Overview of Application Flow



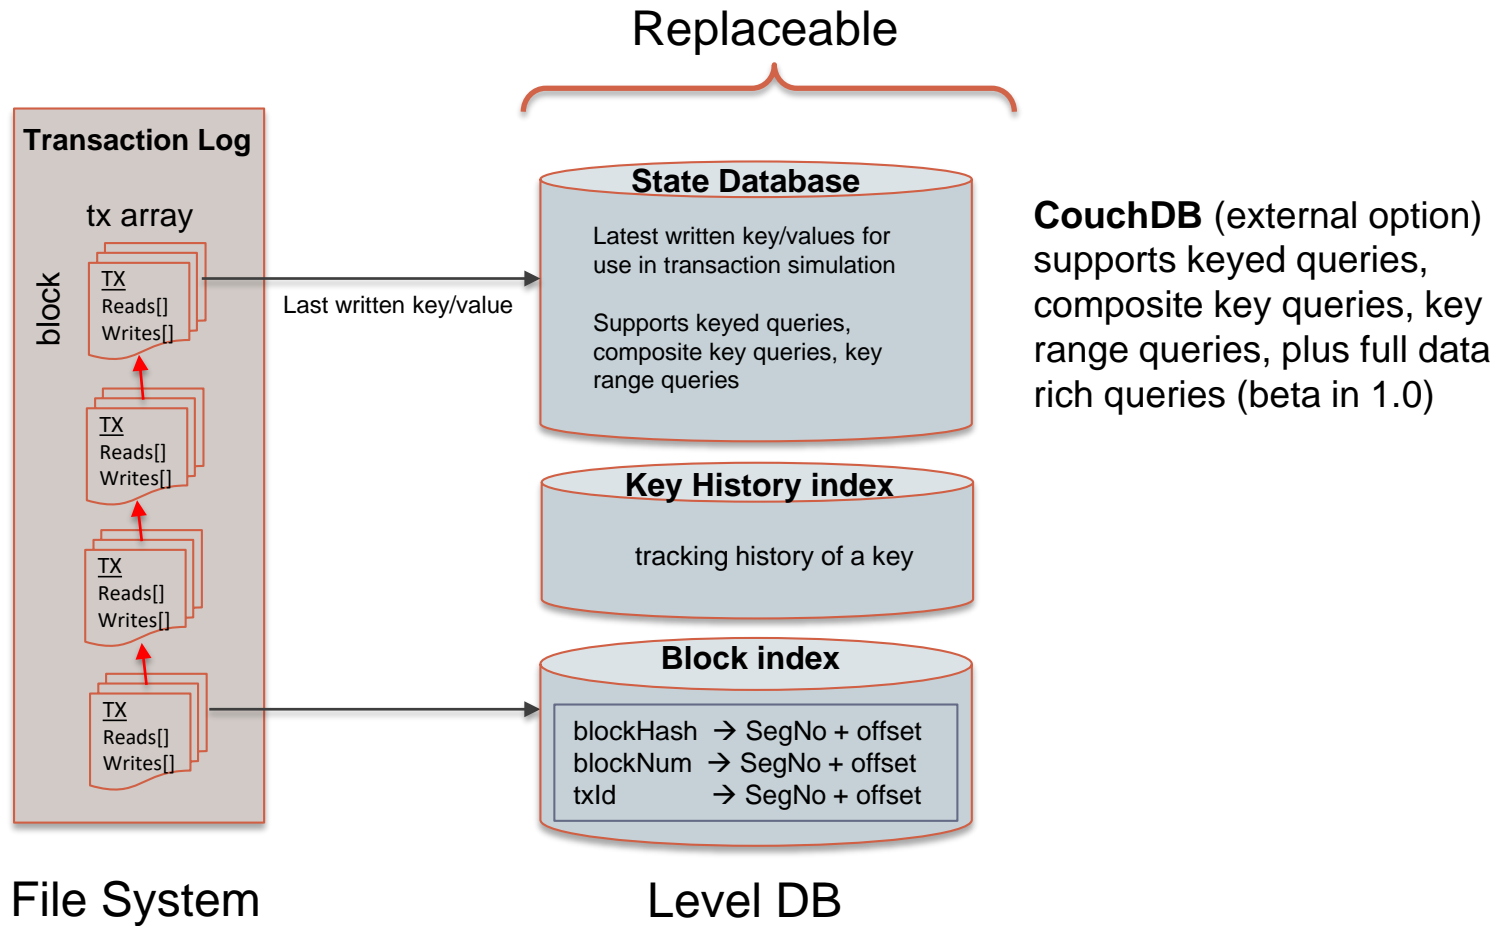
- Developers create **application** and smart contracts (**chaincodes**)
 - Chaincodes are deployed on the network and control the state of the **ledger**
 - Application handles user interface and submits **transactions** to the network which call chaincodes
- Network emits **events** on **block** of transactions allowing applications to integrate with other systems

Fabric V1 Architecture

- Elements of the Architecture
 - Smart Contract (*Chaincode* in Go): **System** and **regular** ones. Deploy/Invoke latter.
 - Assets represented as Key-value pairs in binary and/or JSON form
 - State: Versioned KV model - stored in Level DB or CouchDB
 - Ledger data: Blockchain has full state, including history
 - Transactions
- Kafka for Ordering:
 - No Byzantine fault handling
 - Done to improve performance
 - Pluggable consensus permits other methods

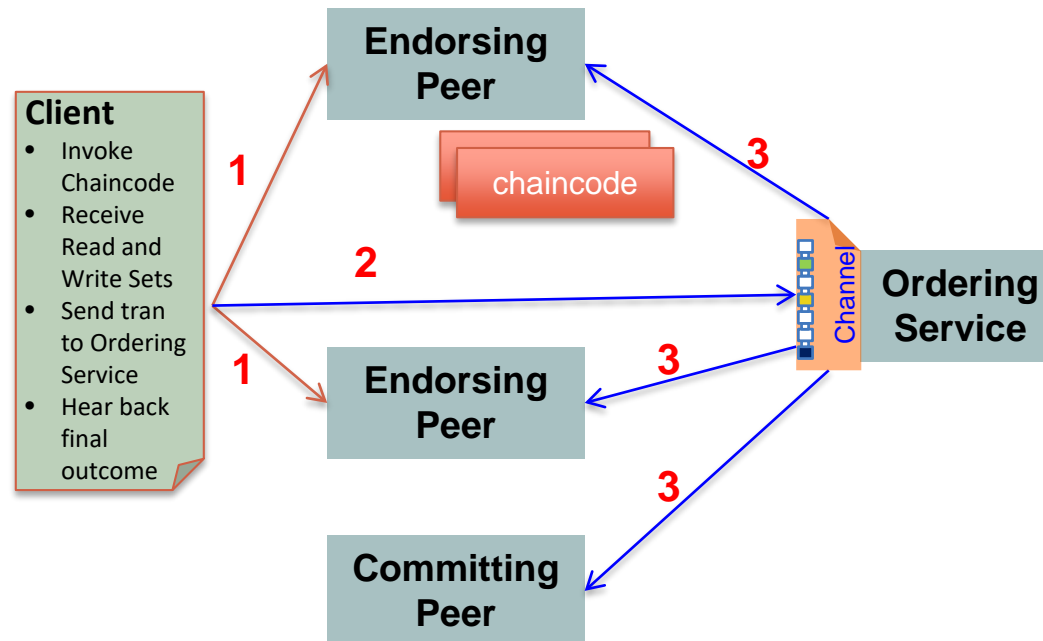
IBM Blockchain Platform and Oracle Blockchain Cloud Service based on it

Fabric V1 Ledger



Transaction Execution Overview Fabric V1

Endorsement, Ordering, Validation/Commit



- Transaction is sent to the counter-parties represented by **Endorsing Peers** on their **Channel**
- Each Peer **simulates** transaction execution by calling specified **Chaincode** function(s) and signs result (**Read-Write Sets**)
- Each Peer may participate in multiple channels allowing concurrent execution
- **Ordering Service** accepts endorsed transactions and **orders** them according to the plug-in consensus algorithm then delivers them on the channel
- All (**Committing**) peers on channel receive transactions: on successful **validation**, **commit** to ledger. No chaincode execution.

Channel.SendTransactionProposal (Step 1) and channel.SendTransaction (Step 2)

Architectural Differences




Figure 1: Order-execute architecture in replicated services.

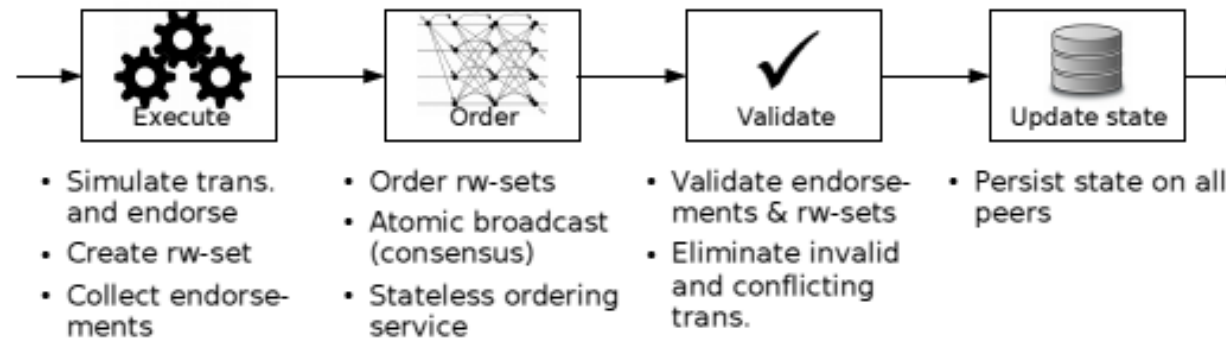



Figure 2: Execute-order-validate architecture of Fabric (*rw-set* means a readset and writeset as explained in Sec. 3.2).

<http://bit.ly/HFpaper>

Endorsement Policies

An endorsement policy describes the conditions by which a transaction can be endorsed. A transaction can only be considered valid if it has been endorsed according to its policy.

- Each chaincode is deployed with an Endorsement Policy
- **ESCC** (Endorsement System ChainCode) signs the proposal response on the endorsing peer
- **VSCC** (Validation System ChainCode) validates the endorsements



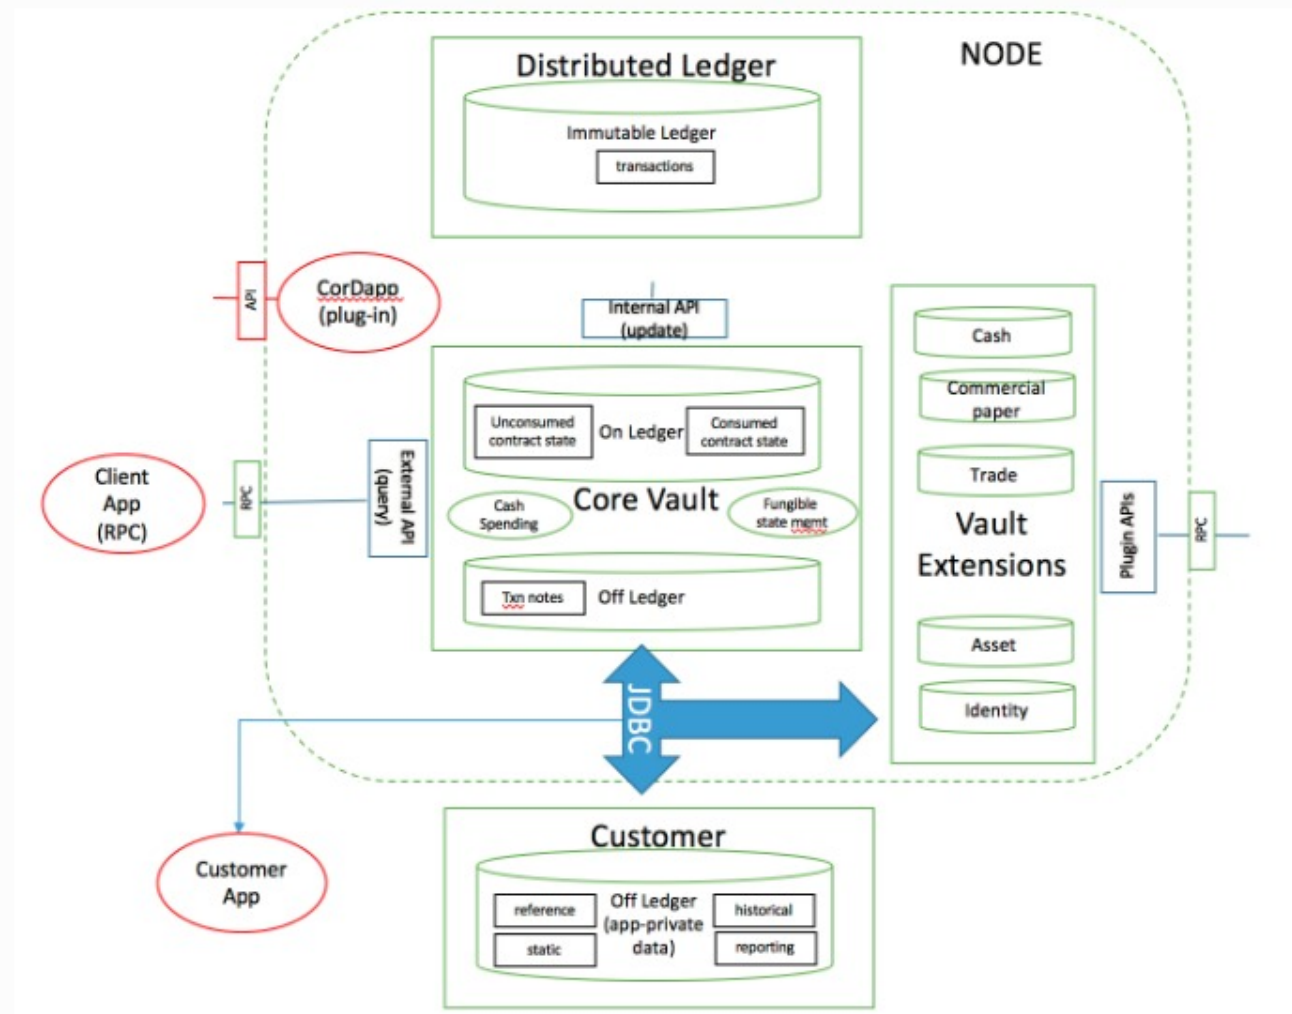
DBMS Implications

- Simulation concept requires layer between chaincode and State DB having to take on analysis of DBMS calls
 - Update statements split into two: read part and write part
 - Read alone sent to DBMS with modifications to retrieve version #s for items read
 - Writes not sent to DBMS but processed and cached locally – doesn't allow for read your own write by chaincode transaction
- During Commit phase, read sets validated by retrieving each item's version # individually and then, if validation succeeds, writes also done one at a time
- Dealing with phantoms requires reexecution of query during commit phase to be sure simulation read set same as read set at Commit time
- Chaincode portability across different State DBMSs hard to do
- Lots of open questions and research issues in this area

R3 Alliance & Corda

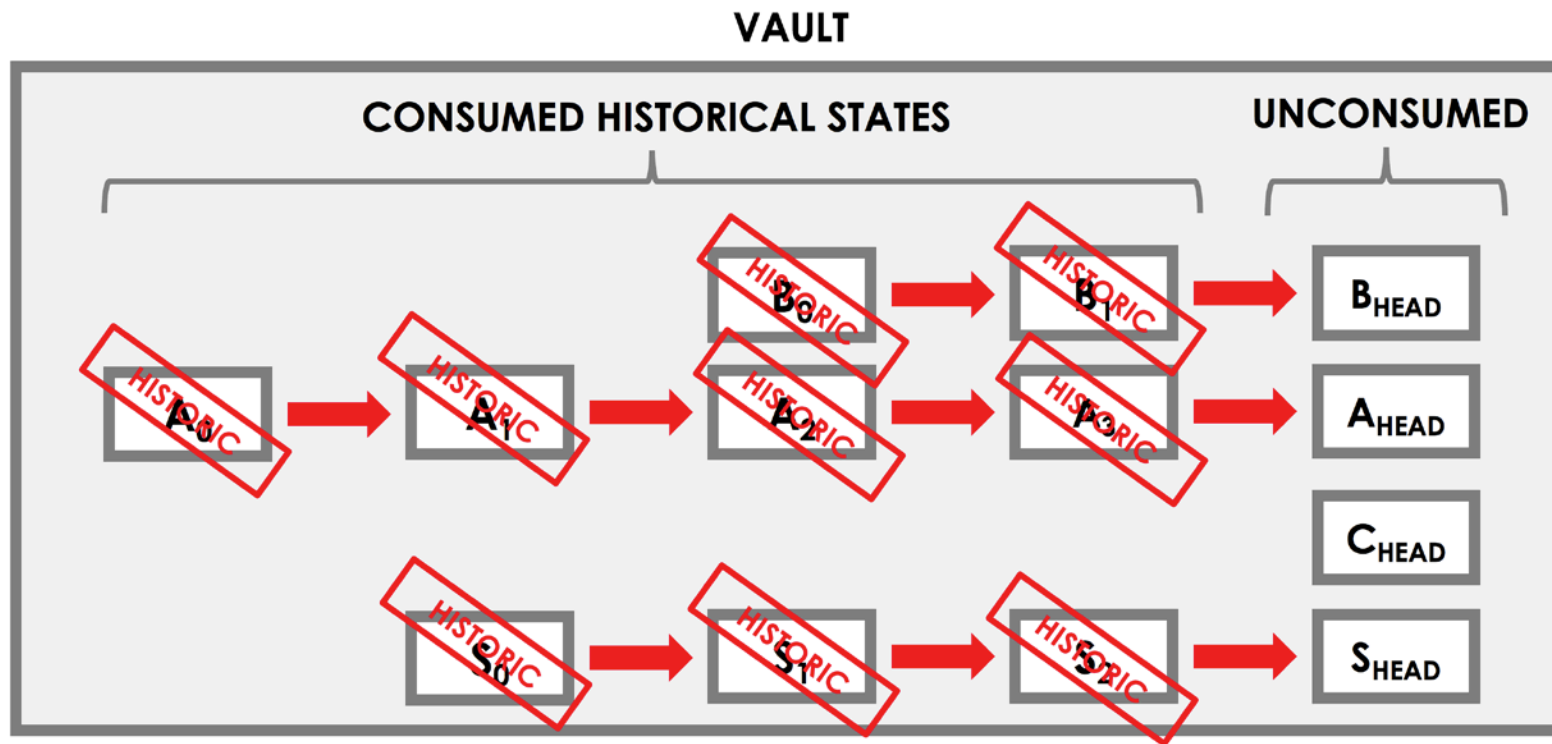
- Barclays, BBVA, Commonwealth Bank of Australia (CBA), Credit Suisse, J.P. Morgan, State Street, Royal Bank of Scotland, UBS
- Special features for JVM to guarantee deterministic behavior
- Hearn, M. Corda: A distributed ledger, Version 0.5, November 2016.
https://docs.corda.net/_static/corda-technical-whitepaper.pdf
- Nodes backed by RDBMS, ledger data SQL queryable and joinable with private tables
- Corda written in Kotlin (simpler Scala with much better Java interoperability) from JetBrains – contracts in Kotlin/Java
- Contract execution is deterministic and its acceptance of a transaction is based on the transaction's contents alone. A transaction is only valid if the contract of every input state and every output state considers it to be valid

R3 Corda Vault




Corda Vault and State

Each node on the network maintains a *vault* - a DB where it tracks all the current and historic states that it is aware of, and which it considers to be relevant to itself

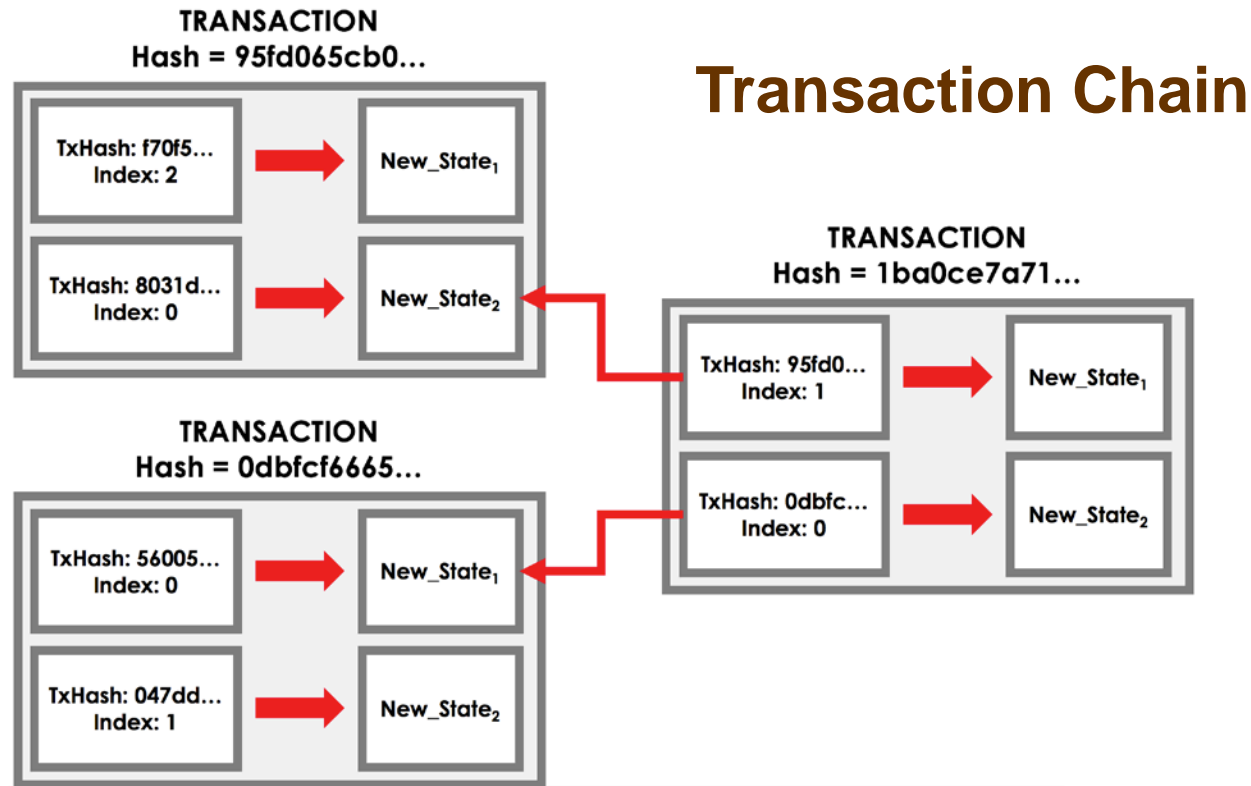


Corda Contract Validity



A transaction is only valid if it is digitally signed by all required signers. However, even if a transaction gathers all the required signatures, it is only valid if it is also **contractually valid**.

Corda Transactions



Notary and Regular Transactions

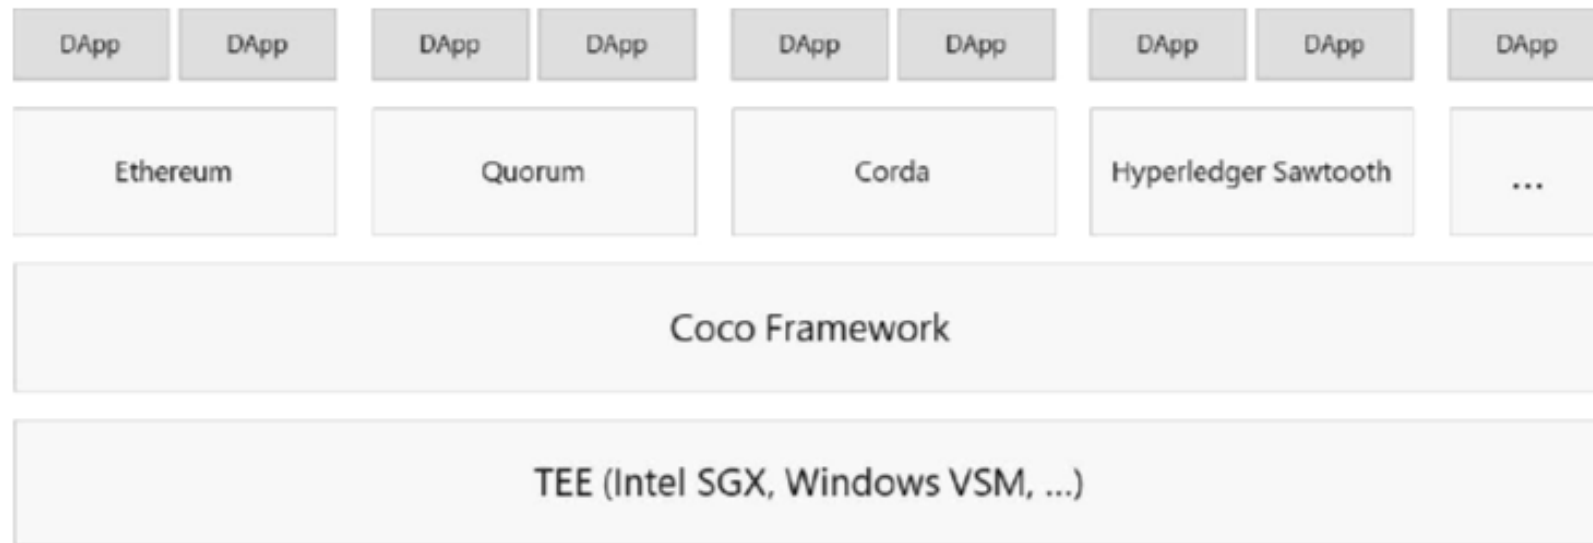
Every state has an appointed notary, and a notary will only notarize a transaction if it is the appointed notary of all the transaction's input states.

Sawtooth (Intel)



- Project of Hyperledger; 1.0 release announced in 1/2018
- Proof of Elapsed Time (PoET) – Consensus Protocol
 - Every validator requests a wait time from a trusted function
 - Validator with shortest wait time for a particular transaction block is elected leader
 - Guaranteed wait time
 - Randomness in leader election (~ to lottery algorithm)
- Intended to run in a Trusted Execution Environment (TEE), e.g., Intel's Software Guard Extensions (SGX)
- Concept of Transaction Family and Transaction Dependencies
- Transaction Scheduling: Serial or Parallel
- Same block can contain multiple transactions which modify same value!
- <https://sawtooth.hyperledger.org/docs/core/releases/latest/contents.html>

Microsoft Coco Framework/System

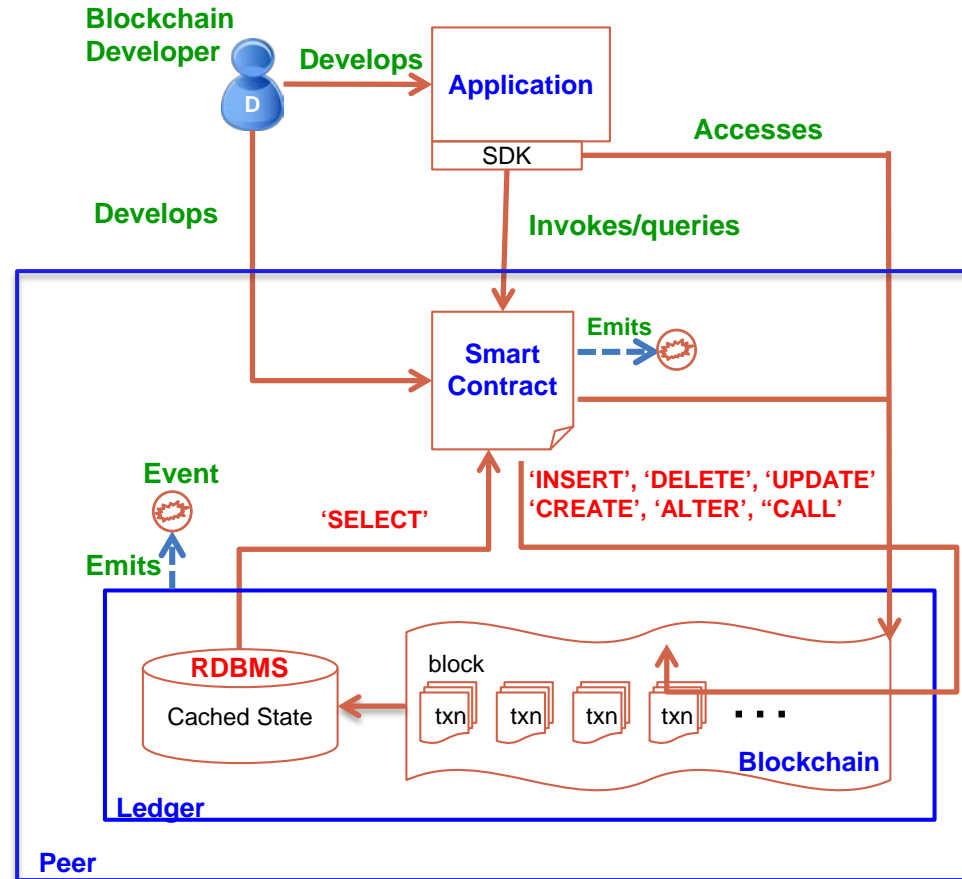


Designed to be open and compatible with any blockchain protocol

Blockchain Architecture/Feature Choices

- Cryptocurrencies Vs Generalized Assets
- Permissionless/Public Vs Permissioned/Private
- Byzantine Vs Non-Byzantine fault model
- Consensus approach: PoW, PoA, PoET, PBFT, ...
- SQL Vs NoSQL data stores
- Transactional stores Vs Non-transactional stores
- Versioned/Unversioned state database
- On-Chain Vs Off-Chain data
- Parallelism exploitation during different phases of transaction execution


Application Flow with RDBMS (In Progress)



- Developers create application and smart contracts (chaincodes)
 - Chaincodes are deployed on the network and control the state of the ledger
 - Application handles user interface and submits transactions to the network which call chaincodes
- Network emits events on block of transactions allowing applications to integrate with other systems

Futuristic Topics

- Smart Contract portability & power of data APIs
- DBMS enhancements to add BC features
- Standards across BC systems
- Cross channel transactions
- Non-deterministic actions
- Analytics on chaincode data
- Many app design issues
- Design tools for endorsement decisions
- NL contracts -> formal contracts -> executable contracts



Numerous research possibilities for database and distributed systems people in this new era of distributed computing!

More Information

Links to Videos, Slides, Bibliography, Twitter Handles

<http://bit.ly/CMbcDB>

Follow me on

Telegram, Twitter, WeChat, Instagram: @seemohan

Facebook: <http://www.facebook.com/cmohan>

LinkedIn: <http://www.linkedin.com/in/seemohan/>