VLDB2020TOKYO

LedgerDB : A Centralized Ledger Database for Universal Audit and Verification

- Ant Financial Services Group

- Alibaba Group

Alibaba 阿里巴巴 蚂蚁金服 ANT FINANCIAL

Terminologies

- DLT (Decentralized Ledger Technology)
- CLT (Centralized Ledger Technology)
  - CLD (Centralized Ledger Database): LedgerDB, QLDB, Oracle BC Table, ProvenDB, etc.
- Immutability: Any piece of data, once committed into the system, cannot be modified by subsequent operations and becomes permanently available.
- Verifiability: The capability of validating specific data integrity and operation proofs.
- Auditability: The capability of observing a serial of user actions and operation trails based on predefined audit rules.
  - Internal audit: an internal user of the ledger can observe and verify the authenticity of all actions.
  - External audit: an external third-party entity can observe and verify the authenticity of all actions.

Why **CLD** is important & valuable ?

- Motivations
  - Decentralization is not proved to be indispensable for permissioned blockchain.
  - Conventional permissioned blockchain and CLD systems:
    - Low performance, storage overhead, regulatory issues, limited external auditability
- Gartner Forecast Gartner
  - Gartner Strategic Vision 2019

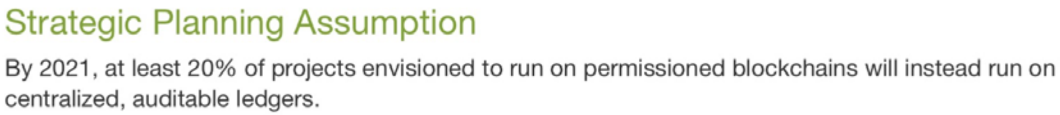

Strategic Planning Assumption

By 2021, at least 20% of projects envisioned to run on permissioned blockchains will instead run on centralized, auditable ledgers.

  - Gartner Strategic Vision 2020

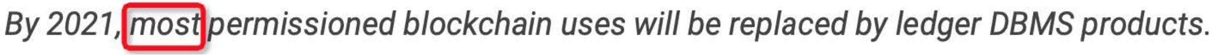

By 2021, most permissioned blockchain uses will be replaced by ledger DBMS products.

Highlight and Comparison

- LedgerDB – a ledger database that provides tamper-evidence and non-repudiation features in a centralized manner (CLD), which realizes strong auditability, high performance, and data removal support.

- Key comparisons between LedgerDB and other systems.

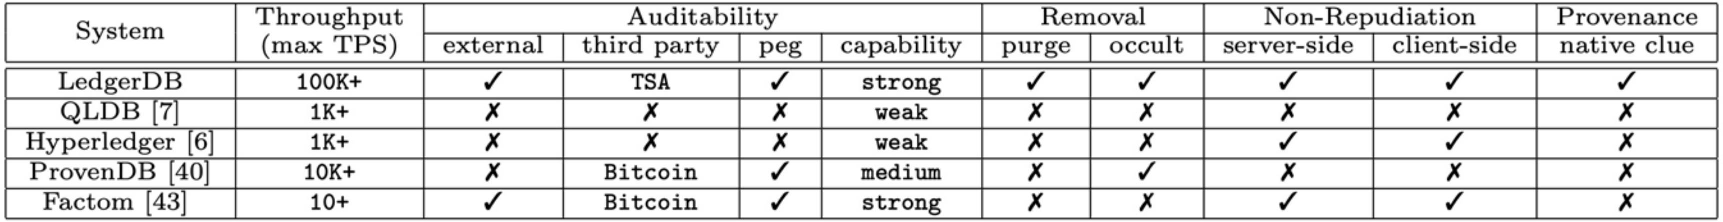

System	Throughput	Auditability				Removal		Non-Repudiation		Provenance
	(max TPS)	external	third party	peg	capability	purge	occult	server-side	client-side	native clue
LedgerDB	100K+	✓	TSA	✓	strong	✓	✓	✓	✓	✓
QLDB [7]	1K+	✗	✗	✗	weak	✗	✗	✗	✗	✗
Hyperledger [6]	1K+	✗	✗	✗	weak	✗	✗	✓	✓	✗
ProvenDB [40]	10K+	✗	Bitcoin	✓	medium	✗	✓	✗	✗	✗
Factom [43]	10+	✓	Bitcoin	✓	strong	✗	✗	✓	✓	✗

LedgerDB system architecture.

Ledger master - manage the runtime metadata of the entire cluster (e.g., status of servers and ledgers) and coordinate cluster-level events (e.g., load balance, failure recovery).
Ledger proxy - receive client requests and preprocesses, and then dispatch them to the corresponding ledger server.
Ledger server - complete the final processing of requests, and interact with underlying storage layer that stores ledger data.

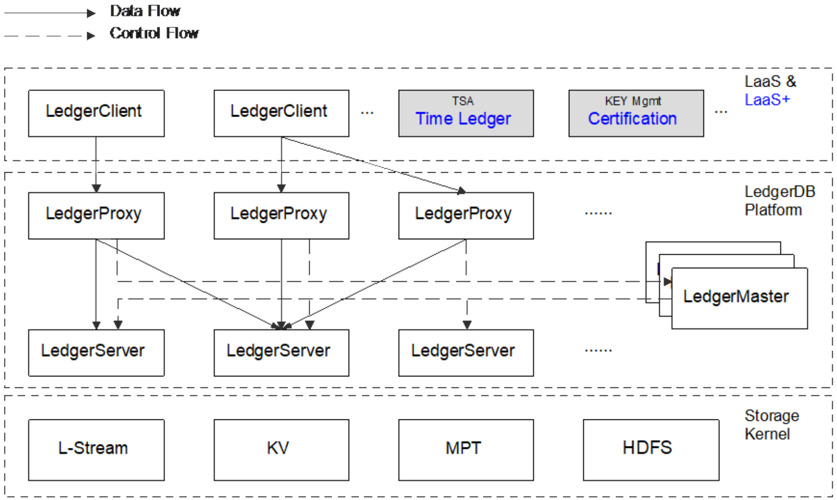

Journal Management

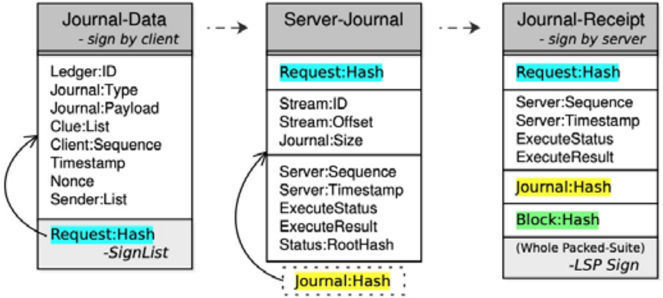

Journal-Data
- sign by client

Ledger:ID
Journal:Type
Journal:Payload
Clue:List
Client:Sequence
Timestamp
Nonce
Sender:List

Request:Hash
-SignList

Server-Journal

Request:Hash

Stream:ID
Stream:Offset
Journal:Size

Server:Sequence
Server:Timestamp
ExecuteStatus
ExecuteResult
Status:RootHash

Journal:Hash

Journal-Receipt
- sign by server

Request:Hash

Server:Sequence
Server:Timestamp
ExecuteStatus
ExecuteResult

Journal:Hash

Block:Hash

(Whole Packed-Suite)
-LSP Sign

LedgerDB adopts an *execute- commit-index* transaction management approach:

① execute - a transaction first enters the execute phase based on its transaction type. It runs on ledger proxy for better scalability.

② commit - collect multiple executed transactions, arranges them in a global order (jsn), and persist them to the storage system. It runs on ledger server.

③ Index - start on ledger server to build indexes for subsequent data retrieval and verification.

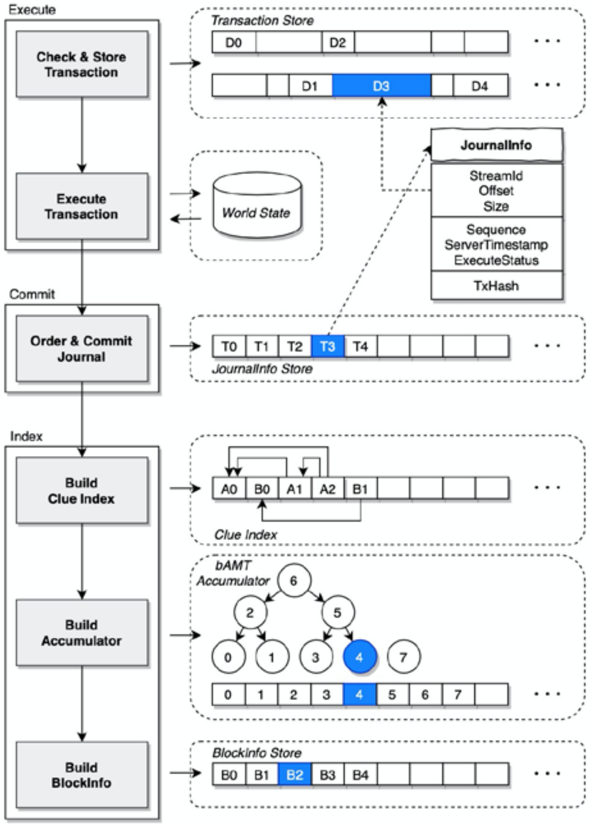

Execute

Check & Store Transaction

Execute Transaction

Commit

Order & Commit Journal

Index

Build Clue Index

Build Accumulator

Build BlockInfo

Transaction Store

World State

JournalInfo

StreamId
Offset
Size

Sequence
ServerTimestamp
ExecuteStatus

TxHash

JournalInfo Store

Clue Index

bAMT Accumulator

BlockInfo Store

Two-way peg TSA notary journals

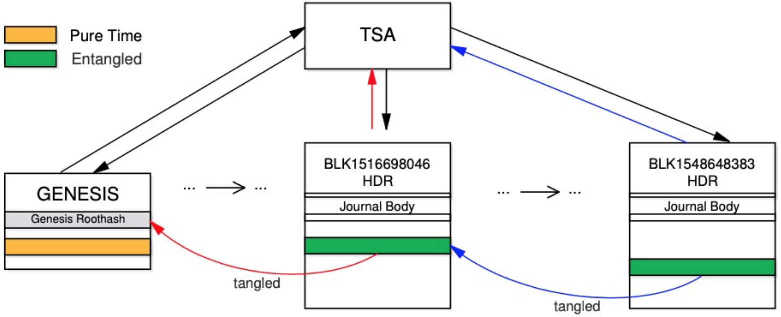

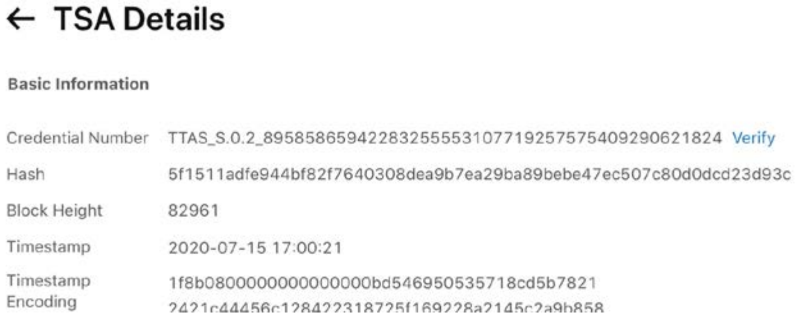

← TSA Details

Basic Information

Credential Number	TTAS_S.0.2_89585865942283255553107719257575409290621824 Verify
Hash	5f1511adfe944bf82f7640308dea9b7ea29ba89bebe47ec507c80d0dcd23d93c
Block Height	82961
Timestamp	2020-07-15 17:00:21
Timestamp Encoding	1f8b0800000000000000bd546950535718cd5b7821 2421c44456c128422318725f169228a2145c2a9b858

- A TSA journal contains a ledger snapshot (i.e., a ledger digest) and a timestamp, signed by TSA in entirety. These journals are mutually entangled between each other, which provide external auditability for timestamps.
- Two-way peg protocol: ① a ledger digest is first submitted and then signed by TSA;
  ② TSA journal is recorded back on ledger as a TSA journal.
- We offer T-Ledger service on Alibaba Cloud LaaS+ (Ledger-as-a-Service).

Verifiable Data Removals

- Purge

A purge operation deletes a set of contiguous (obsolete) journals starting from genesis to a designated jsn on ledger

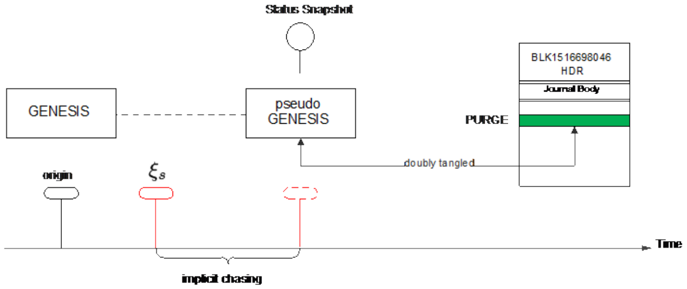

```
DELETE FROM ledger_uri
 WHERE jsn < pur_jsn;
```

- Occult

An occult operation converts the original journal to a new one that only keeps its metadata, and retains its digest.

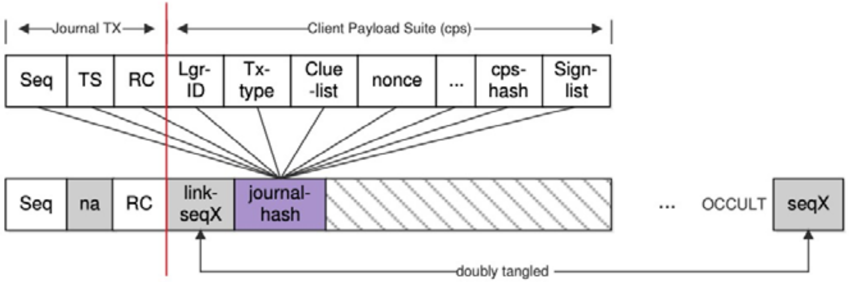

```
UPDATE ledger_uri
 SET TS = na, cps = CONCAT(
  seqX, journal_hash, blanks)
   WHERE jsn = Seq
      OR cid = des_cid;
```

Clue – Native lineage in LedgerDB

- A clue is a user-specified label (key) that carries on business logic for data lineage.
- A typical clue use case of copyrights ledger of NCAC:

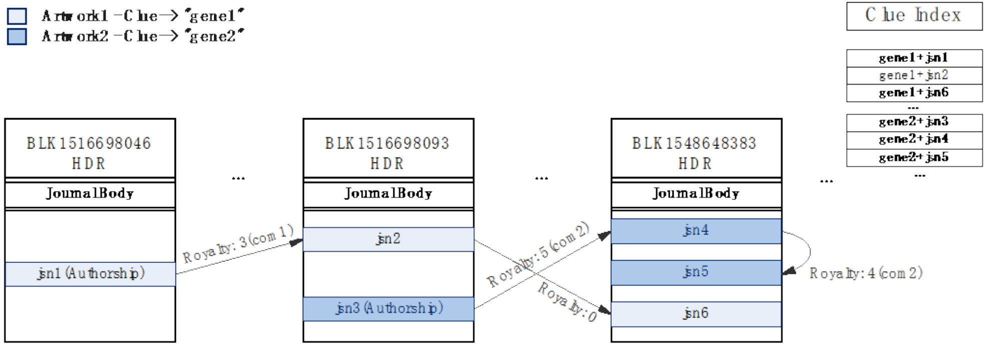

- LedgerDB conducts a write-optimized clue index structure by a reversed *clue Skiplist* (cSL).
- For clue verification, we apply a dedicated verification protocol combining a *clue-counter MPT* (ccMPT) (*n*) and each journal verification based on *n*.

Evaluation – cSL & bAMT

bAMT vs. Libra accumulator

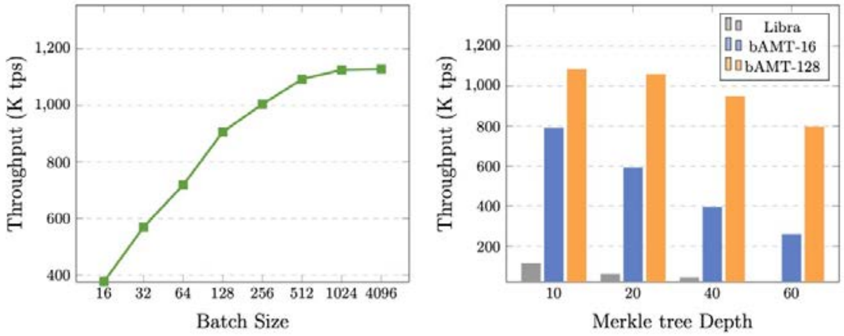

(a) bAMT root calculation

(b) bAMT vs. Libra

cSL vs. RocksDB

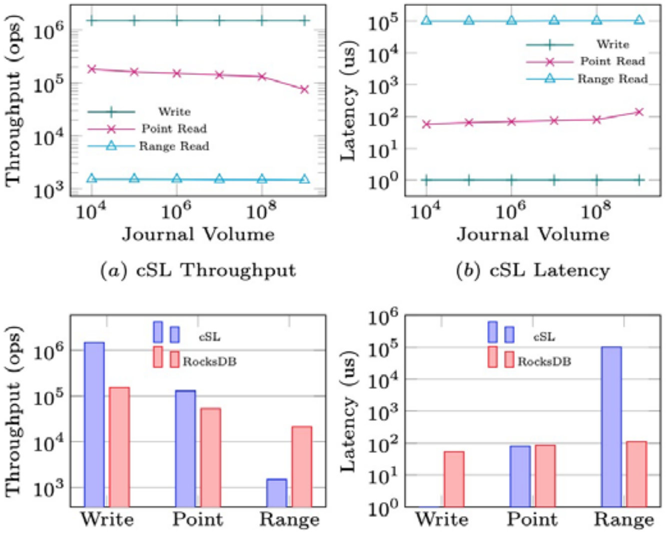

(a) cSL Throughput

(b) cSL Latency

(c) Throughput comparison

(d) Latency comparison

VLDB2020Tokyo

Evaluation – performance and appl

LedgerDB end-to-end performance

LedgerDB is 80× faster compared to Hyperledger Fabric in the same notarization application

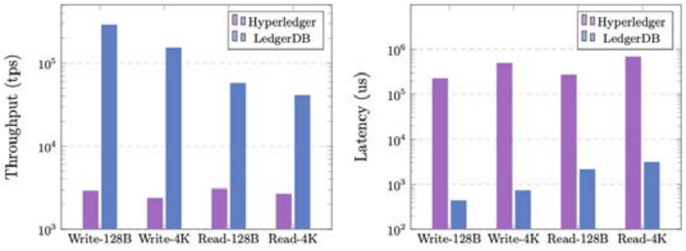

(*a*) Throughput comparison

(*b*) Latency comparison

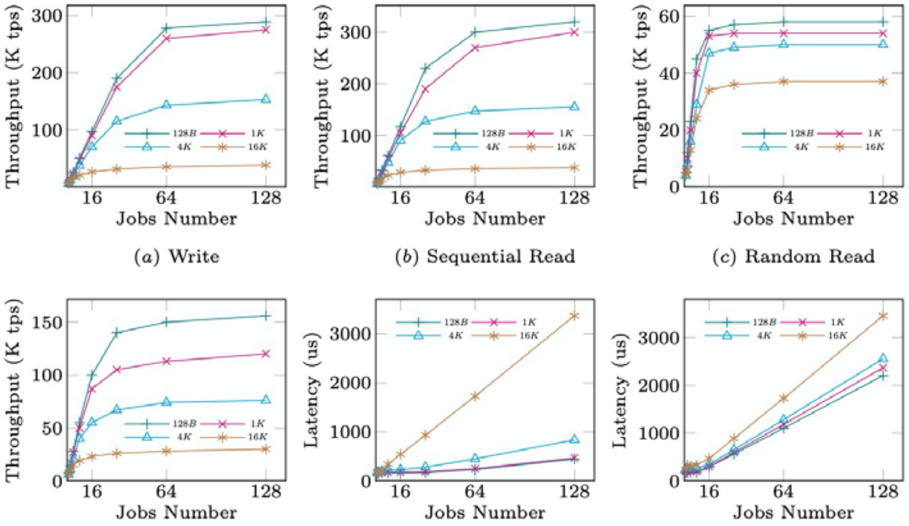

(*a*) Write

(*b*) Sequential Read

(*c*) Random Read

(*d*) Latest Random Read

(*e*) Write

(*f*) Random Read

Alibaba 阿里巴巴

蚂蚁金服 ANT FINANCIAL

LedgerDB in Production

Federated ledger vs. permissioned blockchain

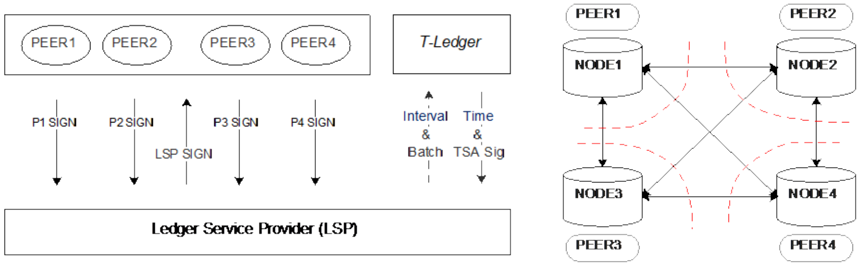

(a) (b)

LedgerDB customer use cases

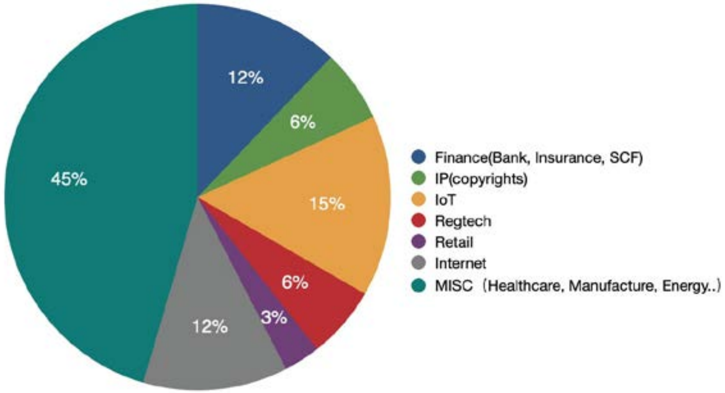

Decentralized vm-like exec is just an implementation, the soul of consensus in ledger technique is dancing with time and cryptographic theorem.

- LedgerDB

https://www.aliyun.com/product/ledgerdb	domestic
https://www.alibabacloud.com/product/ledgerdb	international

Thanks!