IBM Storage

Storage Needs for Blockchain Technology - Point of View

Document version 4

Contents

List of figures

INTRODUCTION

Pioneered by cryptocurrencies such as bitcoin and Ethereum, blockchain is a model for a distributed, trusted, shared, ledger in which you can store transactions. Each transaction is encrypted and identified by a cryptographic signature that is chained to the previous block through additional cryptographic keys. This allows for the blockchain to be a shared, immutable source of truth as each block stores a hash or signature of the previous block. Thus, any change to a block would require recalculating all subsequent blocks and updating the entire chain. Figure 1 shows this concept.

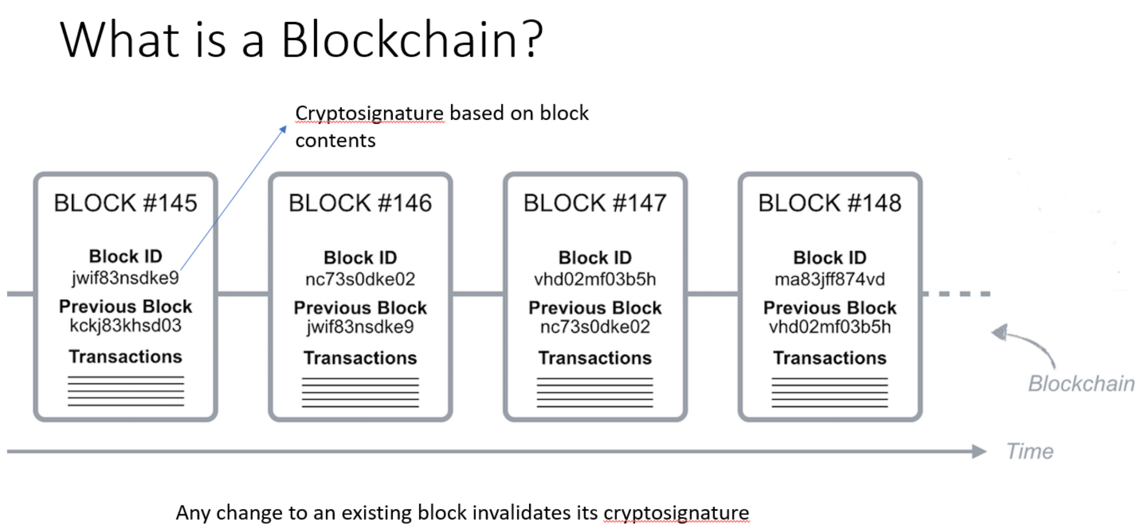

Figure 1: What is a blockchain?

The cryptocurrencies use either proof of work (PoW) or proof of stake (PoS) to be a member of the chain. PoW is a set of complex calculations that result in the creation of a new bitcoin and some estimates suggest that the computations require an equivalent of 5 to 7 kilowatt (kW) of energy be used per bitcoin. Obviously, PoW blockchains are not a high-performance option. PoS blockchains allow you to purchase into a blockchain, slightly more performant but still not a good way to perform thousands of transactions per second (TPS). A PoW blockchain can currently perform about 10 TPS under test and around 3 in reality, whereas a PoS can perform about 20 TPS under test and 6 to 8 TPS in reality. Both of these are non-permissioned blockchains where no one knows the other participants other than their hashed, encrypted identifiers.

Bitcoin is a PoW and Ethereum is a PoS. In PoW, you must utilize a cryptocurrency that you have mined to be a member, and in PoS, each blockchain transaction requires that you have cryptocurrency you have purchased. Each Ethereum transaction requires that you provide gas, which is a subset of an Ether, the cryptocurrency of Ethereum. Currently, Ethereum is one of the only cryptocurrency-based blockchains that provide (through their ERC721 specification for non-frangible tokens) non-cryptocurrency assets to be traded. A frangible token is a single ether or single bitcoin that can be subdivided. A non-frangible token would be a car, house, diamond, or other asset that is normally sold as a complete unit.

The Hyperledger blockchain model is a permissioned, non-POW, non-POS model. Essentially, you can become a member by being approved by the other members and can fully participate without onerous calculations or having to purchase each transaction. Because no POW or POS is required to use Hyperledger, it can support transaction rates that are impacted by the server computing power and speed of the underlying storage, aside from the wrapping code such as the blockchain smart contract. The current performance is 3500 TPS for a Hyperledger-based, performance-optimized blockchain. It is projected that this could soon reach as high as 30,000 TPS. Hyperledger based blockchains can be utilized to deal with both frangible (fungible) and non-frangible items. Figure 2 shows the Hyperledger structure.

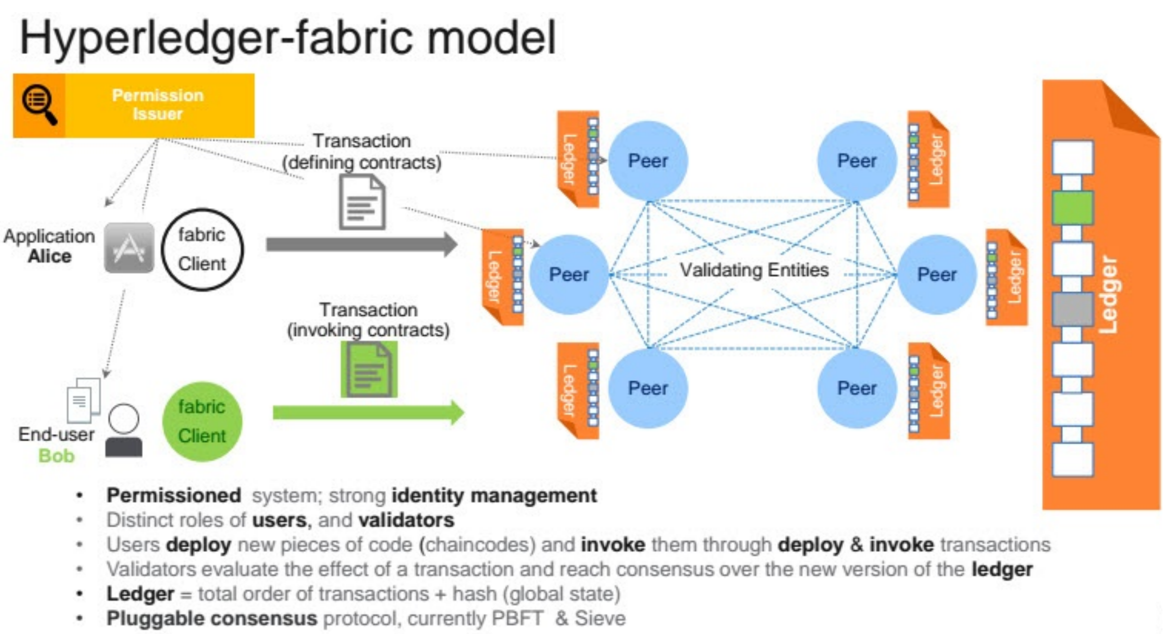

Figure 2: Hyperledger structure

Blockchain requires both on-chain storage of the core ledger data and off-chain storage of data required by smart contracts for verification and documentation. This is demonstrated in the IBM® Blockchain architecture example leveraging IBM Cloud Private as shown in Figure 3.

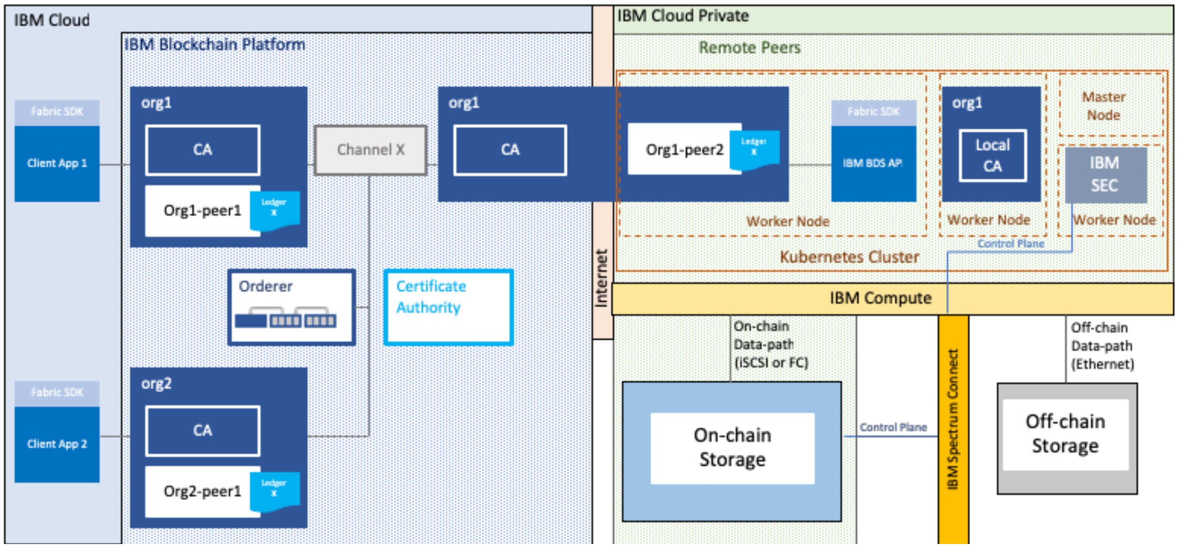

Figure 3: IBM Blockchain Architecture Example

Blockchain is becoming a disruptive technology in payments, logistics, medicine, and anywhere a trusted, provenanced, immutable ledger is required. According to IDC, blockchain spending is projected to reach an upwards of USD $11 billion by the year 2022. This is shown in Figure 4.

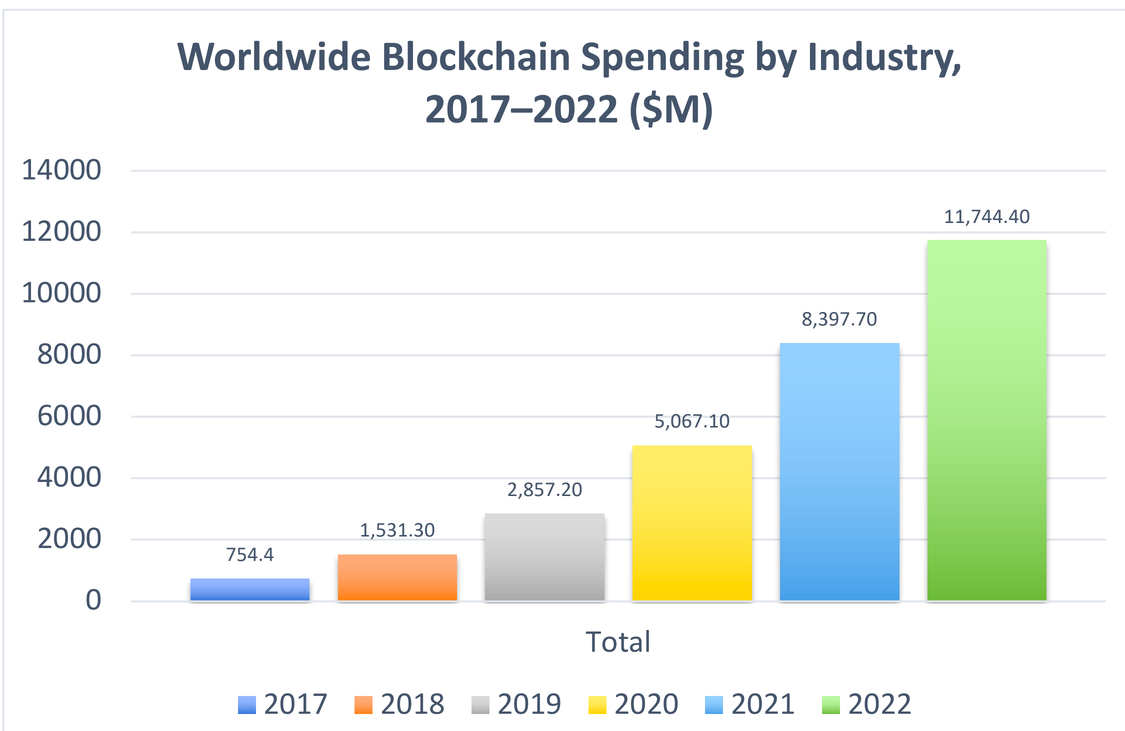

Figure 4: World-wide blockchain projected spending

Recent studies have shown that permanent storage for blockchain data can cost as much as USD $100 per gigabyte (GB). Within the Hyperledger environment, it has been estimated that storage for the blockchains alone, in medium to high-level transaction environments, could be as high as 197 terabytes (TB) or more per year. Figure 5 shows the estimated storage required per transaction per second, per year.

Attempting to find hard data for the storage required for Hyperledger blockchain, for either on-chain or off-chain data is difficult due to the newness of the technology and lack of hard documentation by users who are using it in their daily operations.

As we previously stated, the blockchain ledger is currently immutable. This means that it is append-only and cannot be pruned. Calculating the blockchain ledger storage requirement may not be as trivial as everyone assumes. While pruning may come into play in the future for certain industries, it is expected that some other industries, such as financial and medical, may need to adhere to proscribed retention requirements (some for 50 years). So, even if ledger data is *pruned* from the primary chain, it will still have to be maintained in an archive storage somewhere either in the cloud or on premises.

1 Blockchain ledger storage calculations

Given a basic understanding of storage and some basic understanding of how blockchain stores data, you can derive a back-of-the-envelope calculation on how much storage is required for the blockchain ledger.

1.1 Blockchain ledger assumptions

The following assumptions are made in the blockchain ledger storage calculations:

- All calculations will be in actual bytes, that is 1024 bytes is 1 KB.
- All Hyperledger blocks are 1 megabyte (MB).
- Only hash, signature, or key data is stored in the blockchain ledger.
- The company works 8 hours per day. The company has roughly 240 processing days per year and the number of transactions per second are averaged over those days only (allows for holidays and weekends.)
- Currently bitcoin is storing about 1400 basic transactions per MB block. So, Hyperledger blockchains have larger headers and a bit more robust transaction size so 1000 transactions per block was used as a mark in the sand.

1.2 Calculation

Using the assumptions in the previous section, the calculation to find out the amount of storage required per TPS becomes (example for 1000 TPB):

(1 TPS/1000 transactions per block)*1024 KB/block*3600 sec/hr*8 hr/day*240 days/year=
7,077,888 KB of data per transaction per year
= 6,912 MB

= 6.75 GB
= 0.00659 TiB/transaction/yr

Figure 6 shows the application of this value at various transaction levels.

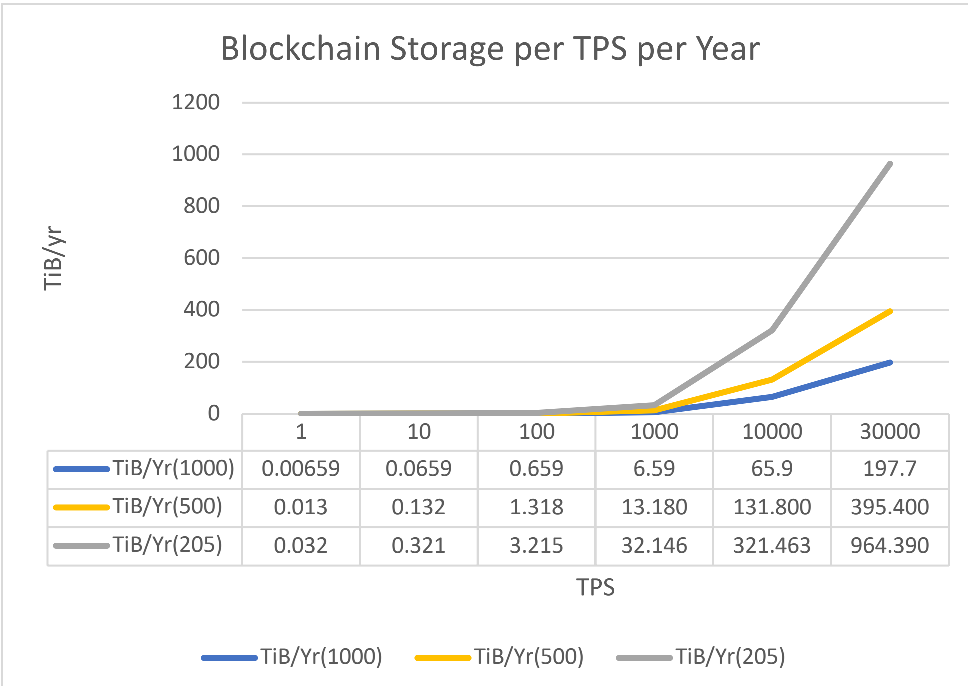

	1	10	100	1000	10000	30000
TiB/Yr(1000)	0.00659	0.0659	0.659	6.59	65.9	197.7
TiB/Yr(500)	0.013	0.132	1.318	13.180	131.800	395.400
TiB/Yr(205)	0.032	0.321	3.215	32.146	321.463	964.390

Figure 5: Graph of storage required per TPS per year for blockchain ledger (IBM calculation)

So even at a modest transaction rates, storage for Hyperledger blockchain ledger is in the terabyte or multi-terabyte size. Using the most conservative estimate of 1000 transactions per block for the 100 TPS level of blockchain activity, 0.659 TiB of storage per year is required. Going with the estimate of 5 KB per transaction, which is based on the actual test values from the Hyperledger team, only 205 transactions per block has been achieved. At 205 transactions per block, the same 100 TPS would require 3.215 TiB of storage per year.

To put this all in perspective, based on the total size of the ledger and total number of transactions stored, bitcoin is averaging close to 555 bytes per transaction or 1,889 transactions per block and Ethereum is close to 2 KB per transaction or 512 transactions per block.

2 Off-chain storage requirements

Since it is not recommended to store non-transaction data, such as pictures, contracts, PDFs, and personal information in the actual blockchain ledger, some form of off-chain or sideDB storage is required. Off-chain data can be structured or unstructured. A hash or signature for the off-chain item will be generated and that is what is stored in the blockchain ledger. The actual item is either stored in the cloud, on premises, or in a near-cloud storage. It is expected that the required storage for off-chain data will exceed the needs of blockchain ledger storage.

Some regulations, such as General Data Protection Regulation (GDPR) specifications for the European Union will also affect how data is stored within a blockchain network. Quoting **Bertrand Portier -** IBM Distinguished Engineer, Blockchain in Financial Services:

"One of the **GDPR** requirements is the right to erasure which is when an individual asks an organization that has their personal data to completely remove that data. The organization then has a limited time to comply. Well, if you know **blockchain**, you know that the blockchain ledger is **append-only and immutable** — there is no "undo" button after a write, and the chain of blocks contains historical transaction information that goes all the way back to when the blockchain was created. That can be a challenge for applying blockchain to GDPR. To comply with GDPR, **no personal data should be put on the blockchain directly**. Techniques exist to deal with this, which consist of putting a cryptographic hash on the chain or the "evidence" instead of the actual data. More guidance and expertise needs to be collected in this space. And, as my Promontory colleagues would say, "Be sure to check with your legal counsel!"

2.1 Off-chain assumptions

The following assumptions are made in the off-chain storage calculations:

- Each transaction does not produce a document. For calculation, a ratio of three transactions for each document is used so a figure of 0.3 is used.
- 8 hours per day is the workday.
- 240 days per year are actual processing days.
- Indexing and other storage requirements are not included.

2.2 Calculation

Using the assumptions in the previous section, the calculation becomes:

1 TPS*docs per transaction*3600 sec/hr*8 hr/day*240 day/year = Documents per year
= 3*3600*8*240 = 814,301 Documents per year for 1 TPS

2.2.1 Storage required for documents

According to current references 1,000,000 documents, of mixed formats, requires 333 GB of storage. This means, 814,301 documents will require:

1000000X/(333*814301) = GB/year/TPS

X = (333*814,301)/1,000,000
= 271 GB/TPS/year
= 0.264 TiB/Year/TPS

Figure 6 shows how this storage required per TPS will result in increased storage requirements as TPS increases.

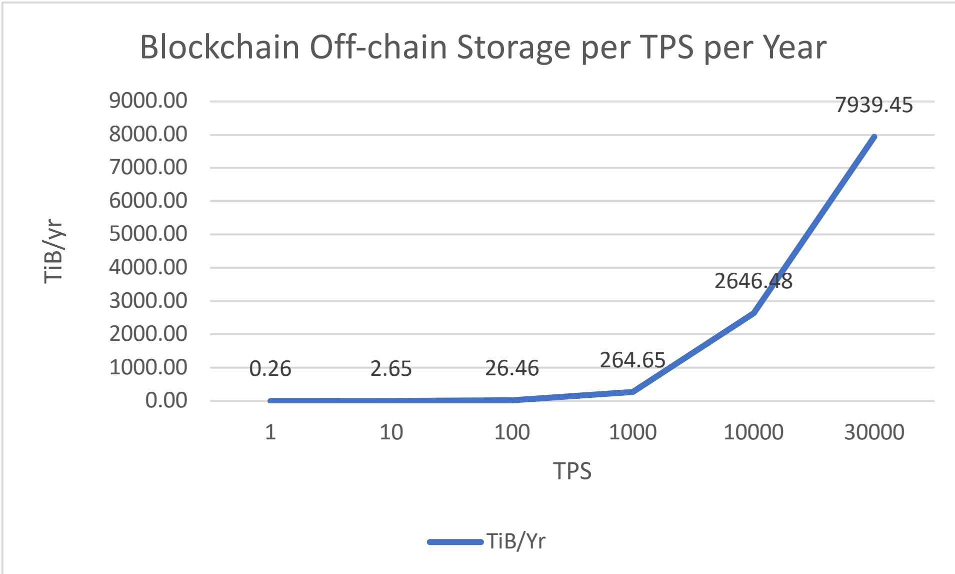

Figure 6: Required storage for off-chain data based on transactions per second

So, as expected, you can see if the transactions produce 1 document for every three transactions recorded, the off-chain storage requirement for blockchain will soon exceed those required for the blockchain ledger by an order of magnitude.

Some blockchain implementations leverage already-stored support information for their off-chain data store. This leads to an unrealistic expectation that no additional off-chain data store will be required for blockchain networks. Even when leveraging existing data, some peers are choosing to move the data they want to share to a separate storage before they present it to their blockchain network. However, many blockchain implementations are requiring additional new off-chain data, be it on premises, or, in the cloud (or a combination of both). As the technology matures, more of this off-chain data will be generated and, for the peer to satisfy the blockchain SLA, its off-chain data will need to be stored in dedicated data stores designed to provide optimal access to this data for all network peers. Moreover, some network peers are choosing to duplicate off-chain data, even if for a short duration, during which they experience a high usage of this data.

One final point is that each peer that participates fully in the blockchain will own a node and will require off-chain storage. The amount of storage specified in these sections on on-chain and off-chain storage are per node. Figure 7 demonstrates this concept.

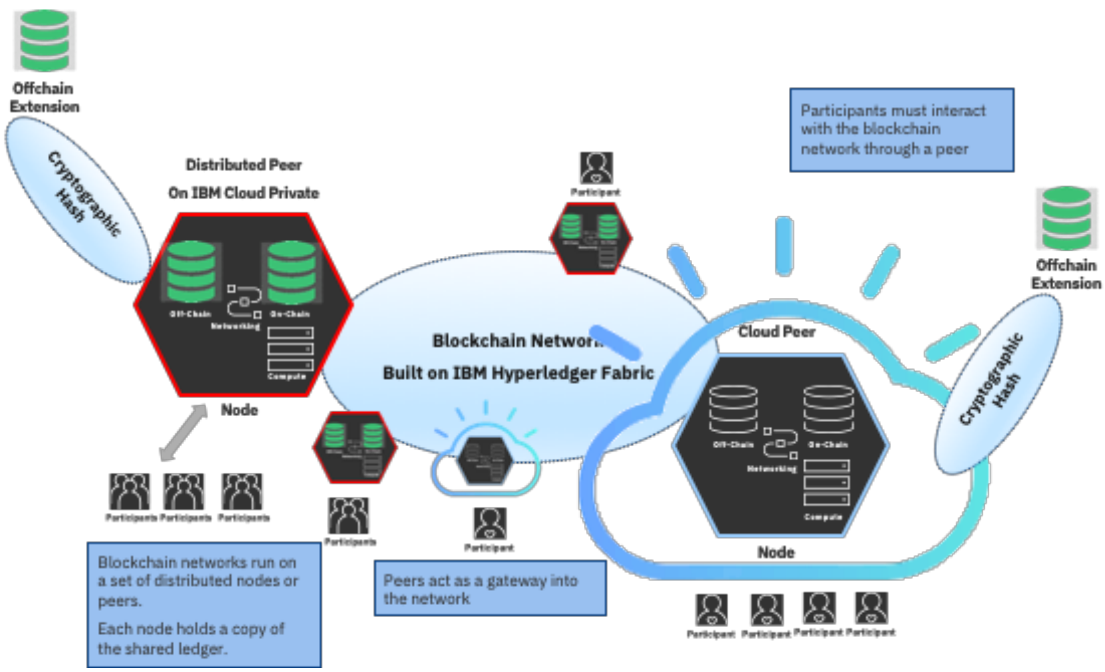

Figure 7: Blockchain peer storage

3 Performance

There are various types of storage configurations that can be implemented for the blockchain off-chain data store. Network speed, peer location, cloud vs on premises blockchain network implementation, will play a role in the overall performance of a blockchain network. Aside from the performance delivered by the compute and the software that makes up and supports the blockchain environment, off-chain data store performance becomes another focus item. Having to cobble together access links to existing data stores or setting up the wrong off-chain data store might cause unacceptable performance to the whole blockchain network. Figure 8 below depicts 3 various example of off-chain data store configurations.

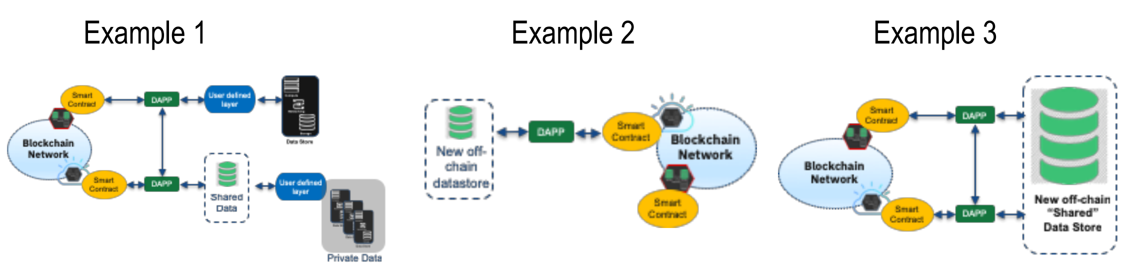

Figure 8: Off-chain data store configurations

Aside from on premises or in the cloud, the off-chain data can also be stored in near-cloud storage in dedicated data stores. This is known as cloud-edge storage. In file systems, such as Internet Protocol (IP) file system or the InterPlanetary File System (IPFS) which forms the basis for the

InterPlanetary Data Base (IPDB), each member is responsible for storing and sharing their off-chain data with all other members in the blockchain network. In such a file share system, the weakest link will be the member with the slowest off-chain storage, assuming that the consortium is using a central shared-storage for the blockchain ledger data.

While much of the focus for Ethereum, and other, blockchain development efforts have been in the realm of open source solutions such as BigDB and IPDB, based no doubt on their non-establishment bias, enterprise-level blockchain users will be looking for proven enterprise-level databases such as Microsoft® SQL, Oracle, and IBM DB2® as well as Postgres. Instead of reinventing the wheel and having to take significant reliability and performance hits as the technology matures, enterprise developers utilize existing, proven technologies. For enterprise-level blockchain users, the focus for off-chain data stores should be on utilizing the best existing, tested, and proved database technology for the type of off-chain data being stored: structured, unstructured, or both.

Newer storage technologies such as NVMe will reduce processor cycle requirements for storage thus increasing the cycles available for cryptosecurity-related calculations. Therefore, block storage is ideal for on-chain deployments. Figure 9 shows a Non-Volatile Memory Express (NVMe) or NVMe-like storage system.

Figure 9: Example NVMe-based storage system

A historical reference may serve to demonstrate the need for performance. When SAP HANA was first introduced, it was portrayed as a memory-only database and storage was not that important. It was later found that utilizing the proper, low-latency, high-speed storage under the SAP HANA database improved performance by at least 30%. If performance had been a key concern from the beginning, many projects would have been more successful.

3.1 Service-level agreements

All members of a blockchain network, Hyperledger, Ethereum, or Corda, will have to commit to a service level agreement (SLA) specifying minimal latency and bandwidth requirements for both on-chain and off-chain storage. Network consideration will have to be given about transactional logic and efficiency to ensure optimal benefit is acquired from processor and memory resources.

However, the off-chain storage would be at the discretion of each member, and thus the need for an agreement to performance-based SLA between members.

4 Storage specification

Ideally, the off-chain storage is a self-contained processing and storage system capable of tiering non-critical data and providing optimal speed for current or critical data. Sharing existing non-blockchain application data could result in sub-optimal performance and might lead to SLA violations and ejection from the network. Any storage must also be scalable, as far as storage volume, with invisible addition of storage capabilities. What is meant by invisible is that adding additional storage should be completely non-impactful to existing operations and transactions. Use of a virtual machine environment such as VMWare or Oracle VM would allow deploying identical systems to each member of a blockchain network, and, as long as identical hardware was deployed to host the VM environment, similar if not identical performance would be obtained.

IBM proposes that the storage for off-chain data be architected accounting for the needs of the data, whether it is structured, unstructured, or both. This means there is no one-size-fits-all solution, but instead a portfolio of storage and compute options must be supplied, allowing the blockchain developer or user to choose the best fit for their environment.

5 Processing requirements for off-chain data stores

Unlike the blockchain peer environment itself where the processing and storage requirement will be met by the peer resources, the off-chain processing required by the off-chain storage (whether it is enterprise or open source) will have to be met by each participant in the blockchain network. Each participant needs to meet the blockchain network SLA to ensure that the performance of the overall blockchain is not compromised. Therefore, the processing requirements of off-chain or sideDBs is a critical factor for the overall performance of the blockchain environment.

The IBM POWER® series of processors, the latest to date being the IBM POWER9™, utilizing IBM architecture, has shown superior performance, reliability, and security to the x86 processor-based commodity servers.

6 Blockchain data management

Throughout this document, we repeatedly mentioned both on-chain and off-chain data for each peer on the blockchain network. Although it’s not required, properly managing all data on the blockchain is highly recommended.

6.1 On-chain data management

Your peer's on-chain data, i.e. your ledger, has built in availability in that it can be recovered from other ledger peers that are on the same network. However, that is assuming the blockchain has enough peers to do so. Thus, managing the on-chain data can be considered a "native" function of the blockchain itself. However, proper data protection principles should still be followed such as setting up backup and recovery protocols as saving copies of your own ledger will someday save your business. Keep in mind that recovering those copies back into the blockchain network is not an easy task to do and would require consensus from the rest of the blockchain participants. Your on-chain storage hardware may already have the right management tools to copy and protect your ledger, for example; the flashcopy tools that comes native with our IBM FlashSystem 9100 family of arrays. Moreover, IBM's Spectrum Copy Data Management (CDM) software is an ideal way to schedule and manage these copies. This provides all the functionality for data recovery due to corruption, sabotage or local failure. Full disaster recovery can also be enabled with the additional use of replication services to the cloud, whether it be the IBM Cloud™, Amazon Web Services (AWS), or Microsoft Azure.

6.2 Off-chain data management

Off-chain data store of an individual's peer and its data management is the responsibility of that peer owner. Here lies the bigger need for proper data management processes and tools. Losing your supporting documentation's, i.e. your off-chain data, may result in failing audits or even being rejected from your blockchain network. Moreover, it is extremely important to properly manage the hashes of your off-chain data, as you create and recover your backups. Note that deleting any supporting data without going through the blockchain may be considered illegal within the blockchain network's terms and conditions.

Given the critical nature of off-chain data, trusted and secure means of backup, recovery, and replication must be provided. Solutions leveraging the IBM Spectrum® family of solutions, come highly recommended.

For block and file off-chain data transfers to the cloud, IBM Spectrum Virtualize™ on the IBM Cloud enables a true hybrid multicloud solution offering the ability to transfer data between on-premises data centers and any IBM Spectrum Virtualize base appliance (IBM SAN Volume Controller (SVC), IBM Storwize®, IBM FlashSystem® V9000 and VersaStack or IBM Spectrum Virtualize only and IBM Cloud).

For unstructured data, IBM Cloud Object Storage provides an innovative, secure, and cost-effective way to store and manage data without creating multiple copies. IBM's groundbreaking SecureSlice technology distributes data slices across storage systems for security and protection. This unique approach combines encryption, erasure coding, and information dispersal of data to multiple locations for protection, without complex or expensive copies. The service is continuously available and can even tolerate catastrophic regional outages without downtime or the need for customer intervention.

To achieve similar dispersed storage on other leading clouds, customers would have to invoke a replication technique, creating a second copy of the data in a separate region, which would then have to be managed. IBM Cloud Object Storage does not require replication techniques or complex configurations to make it fast and efficient to meet cloud data storage needs.

Some users might want to use some sort of stand-alone utility to interact with their storage. As the Cloud Object Storage API supports the most common set of S3 API operations, many S3-compatible tools can also connect to Cloud Object Storage using HMAC credentials.

Some examples include file explorers such as Cyberduck or Transmit, backup utilities such as CloudBerry or Duplicati, command line utilities such as s3cmd or Minio Client, and many others.

7 Does the source blockchain matter?

All blockchains that will be used for tokenized, non-frangible data, will require off-chain storage. When the ERC721 standard was released in November 2017, one of the first uses was for Cryptokitties. Cryptokitties are stylized cat pictures that contain a DNA code for their features. Cryptokitties can be mated to produce a unique offspring based on the DNA of the parents and a proprietary mixing algorithm. After the introduction, the traffic for Cryptokitties crashed the Ethereum network. If more robust storage and planning for off-chain storage and server assets had been used, they would have not had this issue. Within a few months, the Cryptokitties were generating transactions equivalent to USD16 million dollars and had grown to over 600,000 distinct Cryptokitties. While Cryptokitties itself was a great success, the underlying infrastructure failed. While the Cryptokitties example seems frivolous, it does show that you must always plan for success.

Figure 10: Cryptokitties was the first success for Ethereum ERC721 blockchain

Whether it is ERC721 data in Ethereum, pictures of diamonds in Hyperledger or whatever comes down the pike for other blockchain specifications, off-chain storage of some kind will be required. Therefore, this specification is good, no matter the source of the off-chain data.

8 Summary and recommendations

High-performance processors to calculate and confirm the needed crypto-hash and signature verifications, as well as handling off-chain database processing requirements, will necessitate the use of IBM System Z® LinuxONE or POWER9 processor-based computing assets.

When it comes to blockchain storage, backup, recovery, and replication services must be robust, tested, and trusted and must provide enterprise-level performance and reliability. Utilization of containerized, virtual environments will ensure performance when deployed on identical or more efficient hardware platforms.

In view of this need for tiered, high performance, self-contained storage, IBM offers a variety of storage products (both hardware- and software-defined storage) that can meet the on-chain and off-chain needs for blockchain. The IBM Storwize product, such as IBM FlashSystem 9100, delivers performance (NVMe), security (encryption), and controls (SVC) to meet the various on-chain multicloud peer requirements. IBM also offers a variety of Storage that can meet, and sometimes exceed, the off-chain data store and data management requirements for a blockchain network. For additional information, please visit our IBM Storage Solutions for Blockchain (www) page.

REFERENCES

Figure 4 Source: IDC's Worldwide Semiannual Blockchain Spending Guide, 2H17 (August 2018)

Trademarks