



*BY Developers FOR Developers*

**Storage Developer Conference**  
**September 22-23, 2020**

# **Offchain Storage for Blockchain**

**Ramya Krishnamurthy**

**Ajay Kumar**

**HPE**



# Agenda

**Blockchain & Off-chain storage**

**Overview & fundamentals**

**NVME/Flash & Scale-up/Out**

**Performance & Scalability**

**Data reduction & protection**

**De-duplication/ Compression  
Backup & Replication**

**HPE storage for off chain**

**Vendor specific example**



# Blockchain – A Quick Overview

- Blockchain is a distributed database of records, stored in blocks. The data stored in Blockchain cannot be modified or deleted and remains immutable
- Storage infrastructure for Blockchain requires massive scale and performance resulting from petabytes of records and distributed access
- Major use cases include secure, automated transactions, and tracking in the sectors like financial services , healthcare, and supply chain

***Scope of this session is centered around aspects of storage infrastructure needed for Blockchain***

# Off-chain Storage

- A large chunk of the information contained in block chain implementations is unstructured data in the form of files, contracts, documents, emails and even images. The large unstructured data need not reside in block chain nodes and needs a different storage mechanism known as off-chain storage
- As the name suggests, An off-chain storage solution offloads the heavy lifting of large datasets from the Blockchain nodes




- Every Blockchain implementation must have high-performance, highly scalable, secure and easily accessible off-chain data storage
- Even if data resides off the block chain, Immediate low-latency access to shared data remains a requirement

***So, What's needed to provide an efficient off-chain storage?***

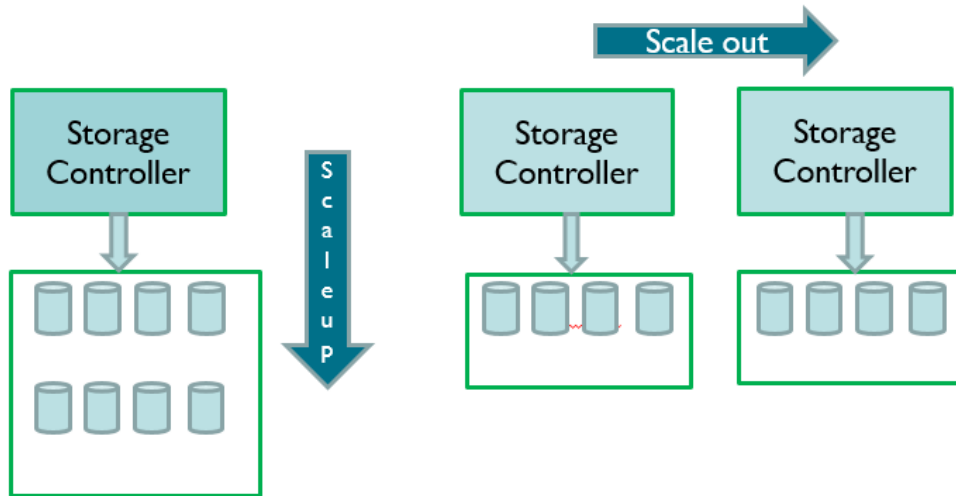
# NVMe/Flash Technology

- Off-chain storage unburdens the block chain of storing large datasets but this can impact performance since physical disks are slow
- NVMe technology can help in reducing processor cycle for storage and provide the desired performance at scale. NVMe bypasses the slow SATA bus and communicates directly with the CPU, providing higher bandwidth, IOPS and lower latency to improve the off-chain workload performance
- NVMe with storage class memory (SCM) ensures persistent data even if there is a power failure or system crash
- NVMe with SCM also provides increased performance as compared to NAND flash drives



# Scalability


- Few transactions in Blockchain can easily lead to generation of a document in unstructured form. Hence the storage requirement of Blockchain increases significantly with the increase in the number of transactions
- Likewise, Performance also needs to be scaled up with an increase in the unstructured data
- We can serve the storage requirements with either scale-up by adding disks/disk enclosures or by adding nodes (Scale-Out) w/o disruption to the availability and accessibility of blockchain



# DATA Reduction

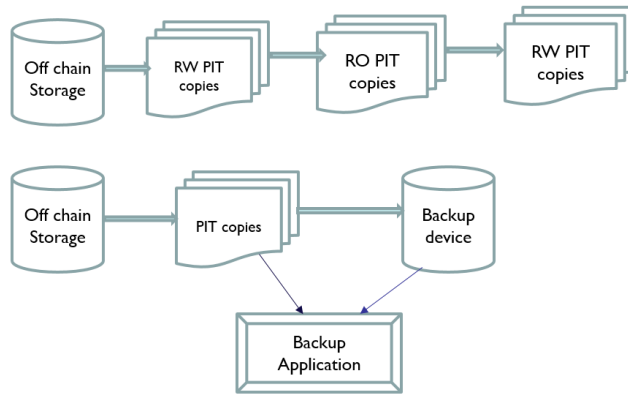
- Total cost of ownership can be reduced with deduplication and compression technology
- In an all-flash storage array, inline deduplication reduces the overall need for capacity
- Depending on the application, deduplication rates of 1:4 can be achievable. With deduplication enabled, a 17.5 terabyte array can perform at the same level as a 50-70 terabyte array. You can see how this levels the capacity playing field between all-flash and traditional storage arrays. This will also help if multiple copies of the same data are stored with the same node as required in Blockchain
- To visualize this data, let's simplify the off-chain storage data by taking an average

**Offchain Storage comparison with data reduction on AFF**



# Data Protection


- Secure backup and recoverability are the table stakes for off-chain data. Off chain storage needs to have a mechanism for backing up volumes
- Backups for off-chain data can be accomplished by using point in time (PIT) copy -The capability to generate a PIT copy of data while an application is online is vital to creating backups
- Operations related to point in time copy should not adversely impact the block chain storage performance and should consume only minimal additional capacity
- Like any other backup, An ability to recover/restore quickly to any previous point in time copy is critical





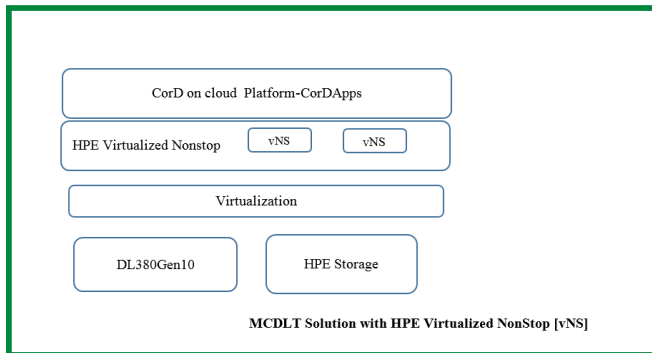
# Data Protection- Replication

- Off chain data needs to be replicated to another storage device at the same site or remote location for disaster recovery, backup, or for data migration
- Choice of replication technology is dependent on the recovery point objective/recovery time objective which can be aggressive for off chain storage
- Synchronous replication writes data to primary storage and the replica simultaneously. That way, the primary copy and the replica always remain synchronized
- Asynchronous replication products write data to primary storage first and then copy the data to the replica, so there is a delay before data is copied to a secondary site



# Proposed Vendor Example

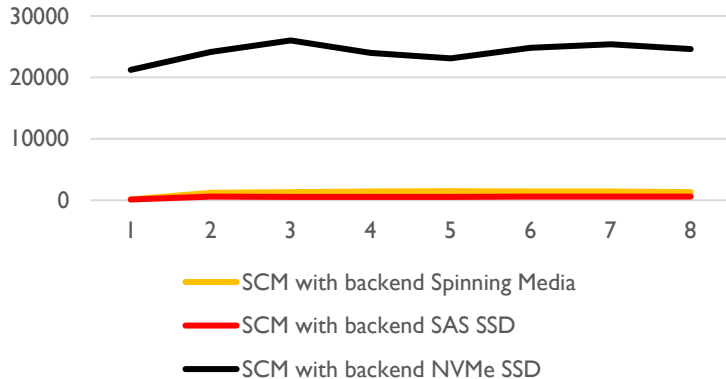
- Hewlett Packard Enterprise (HPE) 'Mission Critical Distributed Ledger Technology' (MCDLT) enables running distributed ledger workloads using Corda- A permissioned distributed ledger platform
- HPE storage will serve as an off-chain storage provider for MCDLT. With MCDLT, Blockchain nodes run on HPE Virtualized NonStop (vNS) system deployed in a virtualized IT environment. Virtualization stack runs on HPE ProLiant servers
- The scale out nature of HPE storage is best suited for mission-critical apps providing desired performance with the NVME technology



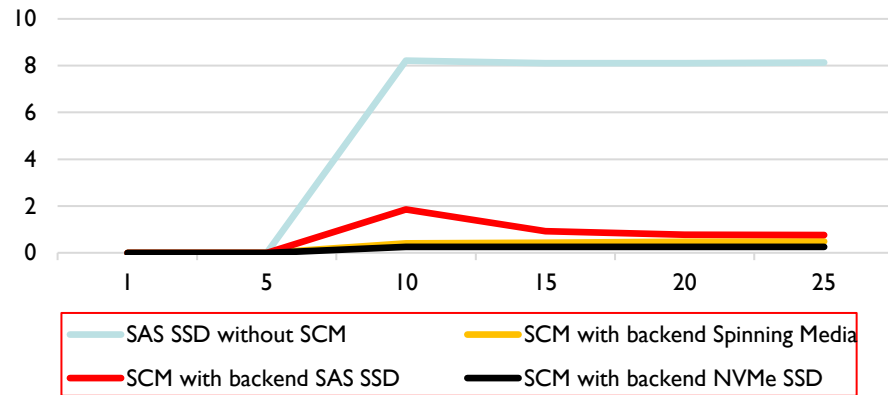
# Proposed HPE storage solution for Offchain

HPE All Flash array along with memory-driven flash combines software intelligence, storage class memory (SCM) and non-volatile memory express (NVMe)- Improves the performance in both read/write intensive workloads with high IOPS and decreased latency.

## Write IOPS



## Read Service Time



# References

HPE Storage

[www.hpe.com/storage](http://www.hpe.com/storage)

3PAR Architecture

<https://h20195.www2.hpe.com/v2/getpdf.aspx/4aa3-3516enw.pdf>

HPE Primera Architecture

<https://h20195.www2.hpe.com/v2/getmobile.aspx?docname=a50000189enw>

HPE virtual copy

<https://h20195.www2.hpe.com/v2/GetPDF.aspx/4AA6-4486ENW.pdf>

HPE Virtualized NonStop Software

[www.hpe.com/us/en/product-catalog/detail/pip.hpe-virtualized-nonstop-software.1009703816.html](http://www.hpe.com/us/en/product-catalog/detail/pip.hpe-virtualized-nonstop-software.1009703816.html)

Corda


[www.r3.com/technology](http://www.r3.com/technology)

HPE Virtualized NonStop Deployment and Configuration Guide

[https://support.hpe.com/hpsc/doc/public/display?docId=a00005448en\\_us](https://support.hpe.com/hpsc/doc/public/display?docId=a00005448en_us)

HPE Servers

[www.hpe.com/servers](http://www.hpe.com/servers)





**THANK YOU**