

The Curious Case of Database Deduplication

Gurmeet Goindi

Oracle

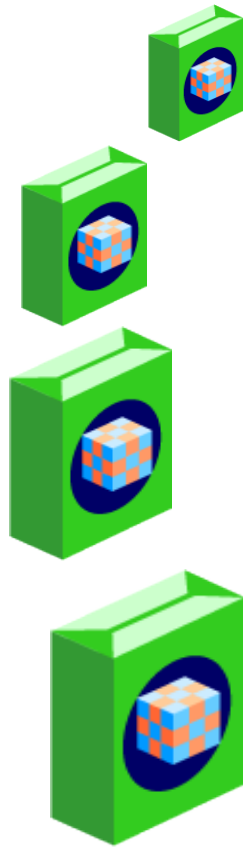
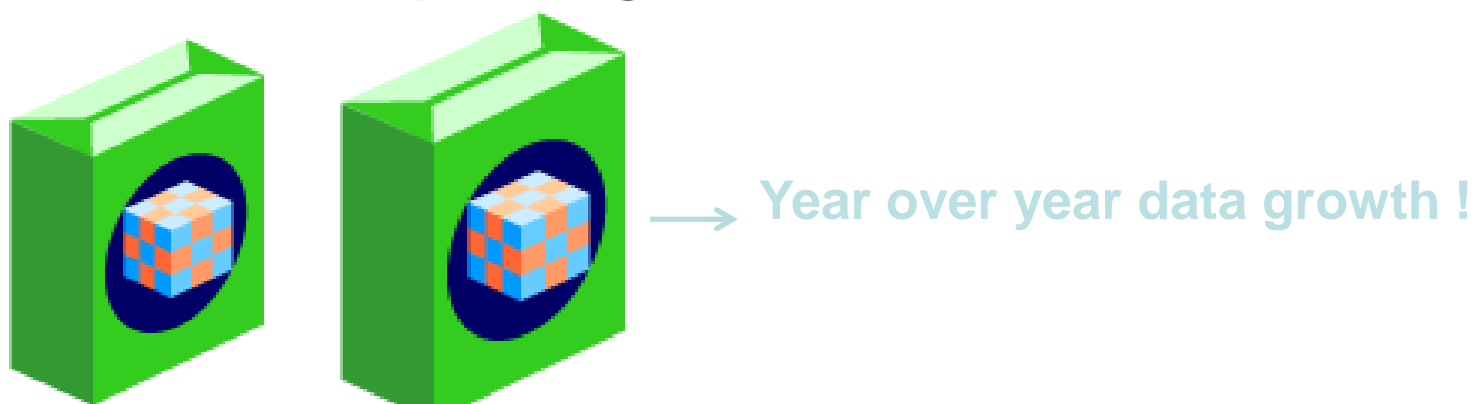
Agenda

- Introduction
- Deduplication
- Databases and Deduplication
- All Flash Arrays and Deduplication

Quick Show of Hands


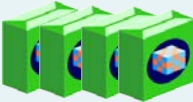

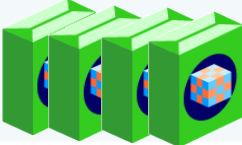
- How many Storage administrators?
- How many System administrators ?
- How many Network administrators ?
- How many Database administrators?
- 😊 Neither? Example, Managers?

Enterprise Data Growth

- 
- Enterprise data is growing 10 - 20% per year
 - ◆ Some industries upwards of 50%+ data growth / annually
 - Increase in production data is magnified in the backup storage infrastructure
 - ◆ More production data expands exponentially into the backup storage infrastructure
- 
- Year over year data growth !

Exponential Growth in Backup Infrastructure

Data Growth Multiplies in Backup Copies:

Production Data Size		One Month of Backups
Last Year:		
This Year:		

- Backup retention periods remain constant based on business needs
 - ◆ Often expanding due to regulatory requirements
- Storage budget doesn't scale with data growth
- Data center space and power constraints hamper storage scale-outs

Backup Storage Optimization Strategies

➤ Backup Compression

- ◆ Software (e.g. native utilities)
- ◆ Hardware (e.g. tape or disk drives)

➤ Deduplication

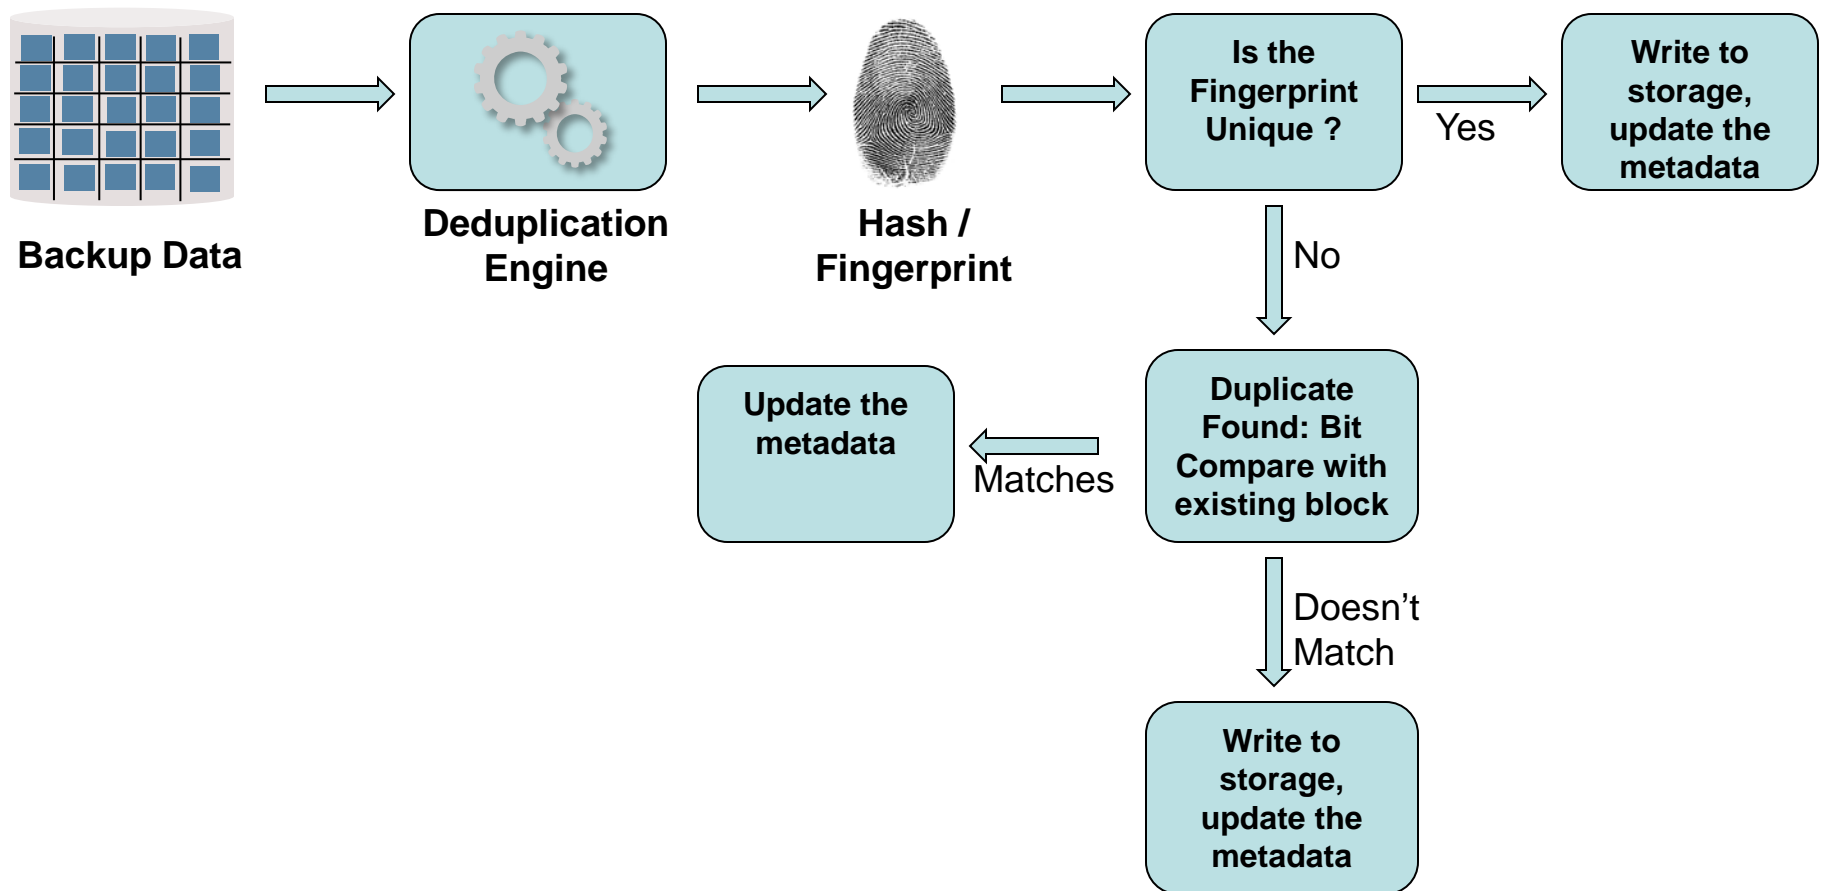
- ◆ Source Side (software)
- ◆ Purpose Built Backup Appliances
- ◆ Hybrid Approach

- Introduction
- Deduplication Concepts
- Databases and Deduplication
- All Flash Arrays and Deduplication

Deduplication 101

- Deduplication: Replace redundant backup data with pointers to shared copy
 - ◆ Redundant data identified by unique algorithms
- Lowers storage costs by reducing capacity requirements
- Can be done at source, inline or post-process
- Mileage varies for deduplication of database blocks

Deduplication Internals



Deduplication & Compression

- Compression is about shrinking the dataset
 - ◆ Encoding information using fewer bits than the original representation
 - ◆ Examples: Lempel-Ziv (LZ) compression
- Deduplication is about finding unique datasets and reducing redundant dataset
- Deduplicated dataset can be further compressed
- Compressed dataset will yield poor and unpredictable deduplication ratios

Source Side Deduplication

- Deduplication software is installed on each host
- Software reads the data to be backed up and looks for duplicates
- Transmits the deduped backup to a storage appliance, for most cases
- Performance impact:
 - ◆ Host : High
 - ◆ Network: Low
- Usually the backups need to be reconstituted when copied to tape

Target Side Deduplication

- Deduplication occurs at backup/storage appliance
- Two flavors:
 - ◆ Inline – Backup is deduplicated prior to being written to disk
 - ◆ Post-process – Backup is written to a staging area, then deduplicated
- Performance Impact:
 - ◆ Host : High
 - ◆ Network: High
- Backups must be reconstituted when copied to tape

Agenda


- Introduction
- Deduplication Concepts
- Databases and Deduplication
- All Flash Arrays and Deduplication

- Leading vendor – 30x times deduplication for File Systems 6x for Relational Databases
- Emerging all Flash array vendor – 10x saving for file systems – 2.1x for Relational Databases
- Most vendors recommend full backups to achieve better deduplication ratios

What drives low deduplication ratios for databases ?

Relational Databases – Quick Intro

- ◆ Relational database: repository and management of data organized in a relational model *
- ◆ Database Management System (DBMS): the software infrastructure to maintain these key characteristics:
 - ◆ Structures: Well-defined objects to store data
 - ◆ Operations: Well-defined actions to access & update data & structure
 - ◆ Integrity rules: Well-defined rules to govern operations
- ◆ Data organized in a table: a two-dimensional relation with rows (tuples) and columns (attributes)
 - ◆ Each row in a table has the same set of columns
 - ◆ E.g. employee, department, salary tables



➤ Database Objects

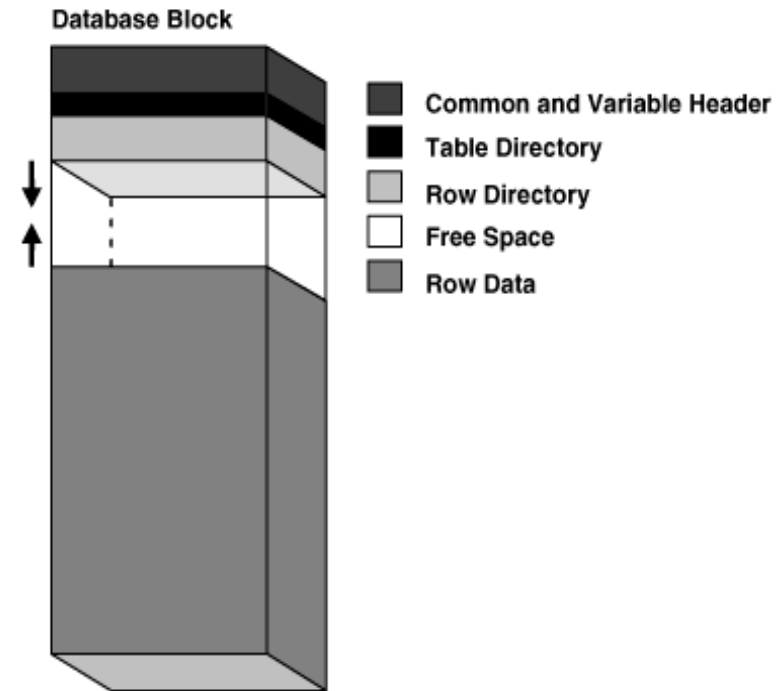
- ◆ Reads and Writes
- ◆ Meta data associated with these Objects

➤ Database Logging

- ◆ Mechanism to deliver transactional consistency
- ◆ Critical functionality for database operation
- ◆ Availability features


Database Block Format

- Database manages its logical storage in data blocks
 - ◆ Minimum unit of I/O
- A data block has a well-defined structure
 - ◆ Block header is kept consistent with payload, rows do not overlap, metadata in its place ...
 - ◆ More than simple bits: can always verify logical consistency



Database Block Format .. Contd

- Even the row (user) data is very carefully formatted within a database data block
- Blocks have a fixed size (usually 8K)
- All blocks are unique
- Update to a block will result in metadata updates to adjacent blocks as well



Understanding Database Backup Options

Full Backups

Image copy:

- Fairly efficient for restore
- Inefficient use of Server, Network and Storage resources
- Might Provide Better deduplication ratios

Application aware intelligent image copy

- Better storage and network resource utilization, but not for server workload
- Faster restore times compared to incremental copies


Incremental Backups

- Transmit only changes on a periodic basis
- Efficient use to storage, network and server resources
- Time to recover depends on the number of incremental backups needed to be restored
- Faster than any full backup strategy, reduces backup window


Most successful Database backup strategies include a combination of the two approaches to meet your Recovery Objectives

Incremental Backups And Deduplication


■ Changed data




Full Backup
Sunday




Incremental
Monday



Incremental
Tuesday



Incremental
Wednesday



Full Backup
Thursday

Blocks
Transmitted

All

1 2 3 4

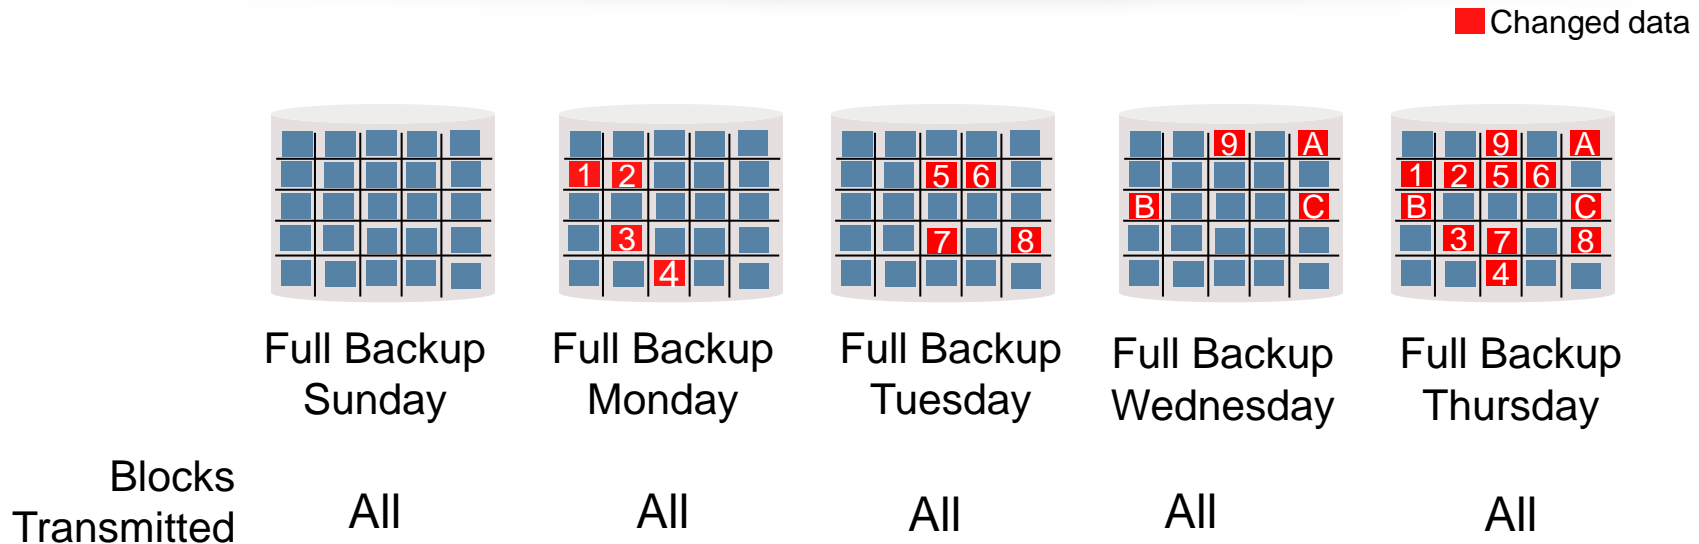
5 6 7 8

9 A B C

All

- A database aware incremental backup only has unique blocks
- Deduplication ratios are **minimal / unpredictable**
 - ◆ Limited to duplicates in the user data within the block
- Low impact on database server and network

Full Backups And Deduplication



- Huge number of redundant blocks transmitted
 - ◆ In the above example: almost 80% of the traffic is redundant
 - ◆ With a typical 5%-10% churn rate nearly **90-95% traffic will be redundant**
- Hence deduplication ratios are high
- Significant impact on database server and network

Database Compression and Deduplication

- Several database integrated compression techniques:
 - ◆ OLTP data compression on user data
 - ◆ Columnar compression
 - ◆ Backup compression
- Most will yield unpredictable deduplication performances on the backup stream
- Similar implications with encryption as well

What About Database Logs ?

- Fundamental structure of the database:
transaction logs
 - ◆ Most crucial structure for data recovery
 - ◆ Contents: change data, undo data, schema / object management statements, etc.
 - ◆ Well-defined data fields, relationships
- Almost zero scope for deduplication

```
REDO RECORD - Thread:1 RBA: 0x00003e.0000d188.01b0 LEN: 0x07c8 VLD: 0x01  
SCN: 0x0000.00d6ca18 SUBSCN: 1 04/30/2007 21:06:42  
CHANGE #1 TYP:0 CLS:67 AFN:3 DBA:0x00c1fbb1 OBJ:4294967295 SCN:0x0000.00d6ca15 SEQ: 1 OP:5.2  
ktudh redo: slt: 0x0007 sqn: 0x00000000 flg: 0x000a siz: 120 fbi: 0  
uba: 0x00c05098.0004.01 pxid: 0x0000.000.00000000
```

To Summarize

- Is deduplication not the correct design choice for Database backups? Well, it depends
- Better deduplication ratios possible:
 - ◆ With daily full backups
 - › May not be practical, based on database size
 - ◆ By disabling database compression / encryption
 - › Consistent with your security / capacity policies?
 - ◆ After customizing backup parameters
 - › Need to balance database / dedupe best practices
- For storage efficiency, do the testing + math

Daily incremental backups +
native compression

vs.

Daily full backups + 3rd
party dedupe

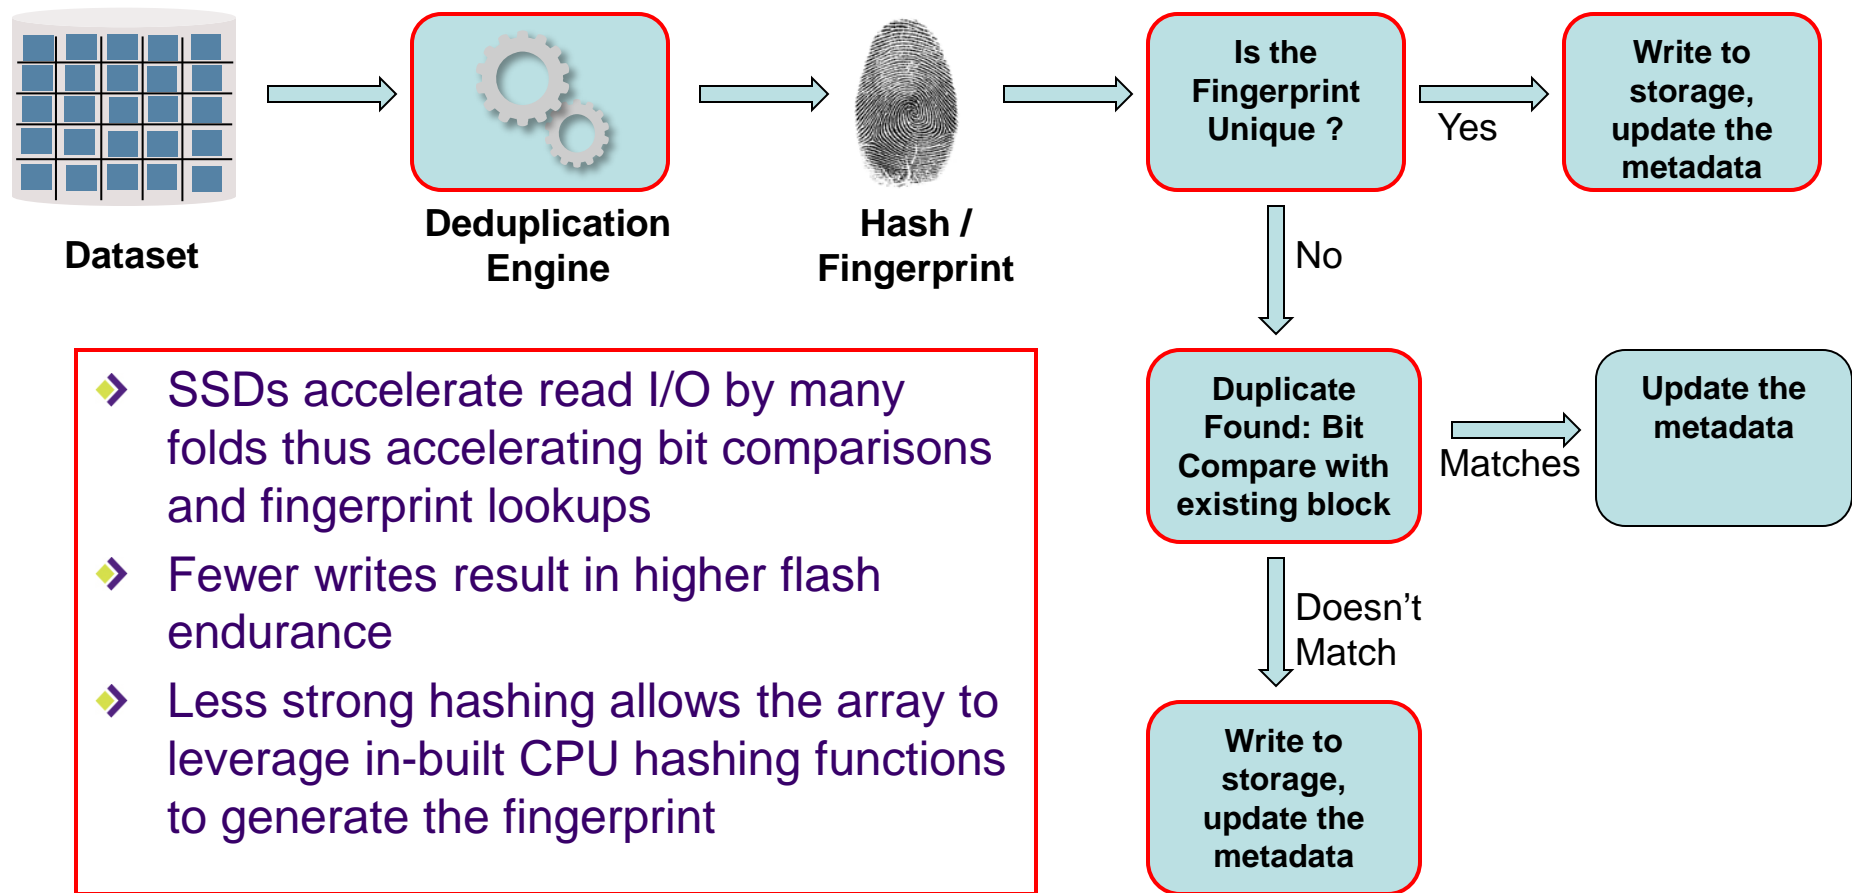
Agenda

- Introduction
- Deduplication Concepts
- Databases and Deduplication
- All Flash Arrays and Deduplication

Deduplication And All Flash Arrays

- Solid State Drives are still 5-10x more expensive than Hard Disk Drives
- Economics of SSDs improve as the footprint of the dataset to be stored is reduced
- Some All Flash Array vendors are leveraging deduplication to increase the storage efficiency of SSDs
- Deduplication ratios are fairly modest compared to that seen in secondary storage or Purpose Built Backup Appliances

Deduplication in All Flash Arrays – Technology Enablers



- ▶ SSDs accelerate read I/O by many folds thus accelerating bit comparisons and fingerprint lookups
- ▶ Fewer writes result in higher flash endurance
- ▶ Less strong hashing allows the array to leverage in-built CPU hashing functions to generate the fingerprint

Databases, Deduplication and All Flash Arrays

- Most All Flash Arrays use fixed length chunk size
- By definition each database block update results in a unique block and changes to metadata of adjacent blocks
 - ◆ Algorithms that are unaware of database block boundary will generate a unique hash for each update
 - ◆ Thus defeating the purpose of deduplication

Questions ?