SNIA NSF | NETWORKING STORAGE

Everything You Wanted to Know About Storage But Were Too Proud to Ask: Data Reduction

Live Webcast

August 18, 2020

10:00 am PT

SNIA-At-A-Glance

SNIA-at-a-Glance

185
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

Learn more: **snia.org/technical** **@SNIA**

Technologies We Cover

- Ethernet
- iSCSI
- NVMe-oF
- InfiniBand
- Fibre Channel, FCoE
- Hyperconverged (HCI)
- Storage protocols (block, file, object)
- Virtualized storage
- Software-defined storage

SNIA NSF | NETWORKING STORAGE

SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

What We Promised to Cover

- How companies end up with so many copies of the same data
- Difference between deduplication and compression – when should you use one vs. the other?
- Where and when to reduce data: application-level, networked storage, backups, and during data movement. Is it best done at the client, the server, the storage, the network, or the backup?
- What are snapshots, clones, and thin provisioning, and how can they help?
- When to collapse the copies: real-time vs. post-process deduplication
- Performance considerations

The Age of Loads of Stuff

- Why So Much Data?
  - Facebook photos
    - 2008: 10 billion
    - 2014: 240 billion, +340 million/day
    - 2020: couldn't find any stats; perhaps no-one knows or cares to count any more
  - Google search on "google" returns 9.8 billion page hits
- Most companies' growth less than this, but multiPB of data with multiTB/day growth not uncommon

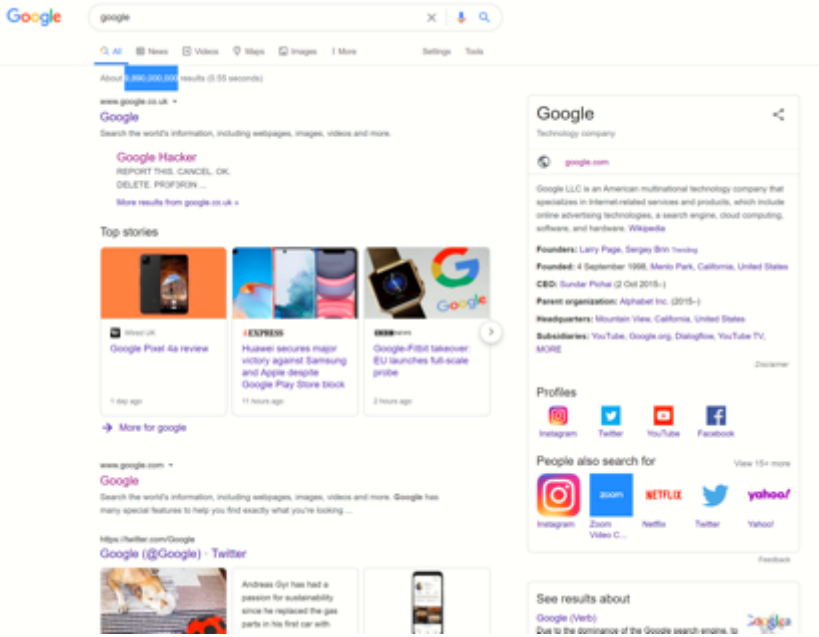

Data Generates Data

- Data tends to spontaneously multiply
  - Create the data here, copy it there to process it
  - Back it up
    - Back that up too
  - Put it on the cloud
    - Three times just in case
  - Copy it again to process it somewhere else
  - Change it, if only very ~~slightlt~~ slightly
  - Lather, rinse, repeat
- Which is fine, but that much repetition is expensive to keep
  - Per bit cost is reducing, but the number of bits is growing faster*
  - Difficult to impossible to delete data
- How can we reduce the impact?

*Anecdotal, but the total storage market value is increasing even though the per bit cost is decreasing, therefore the claim

Data Compression, Deduplication, and Data Compaction

- Deduplication
  - Removing duplicates from the data
  - Sparse data; removing zero or repetitive bit pattern data
- Compression
  - Squeezing data into a smaller space
- Compaction
  - Filling the gaps
- Snaps & clones
  - Copying only the changes
- Thin provisioning & hole punching
  - Appearing smaller and recovering space
- Techniques can be applied to all storage types
  - File, block and object

Prerequisite; Virtualized Storage

- Storage vendors use many of these techniques, but…
  - Don't assume that all systems work the same way; they don't, as there are a variety of ways of performing data reduction
  - This session reviews some of the commonest techniques, but it's not exhaustive and it's not authoritative for all
  - Performance varies between solutions; we'll discuss only briefly
  - Specific techniques; ask the vendor directly
- So, in general…
  - Requires a "virtualized" storage systems
    - Because data reduction technologies generate meta-data (data that describes data)
    - One-to-one relationship between data and layout of the data on the device is lost (by necessity)
  - We'll use a simplified conceptual model of storage to demonstrate the techniques, based on block storage

Physical Layouts

- Device level: based on blocks, sectors, clusters & pages
- SSD have different layout characteristics than hard disks
- However…
  - Detailed driver knowledge is really one (or two) layers of technology below what we need to know & discuss here

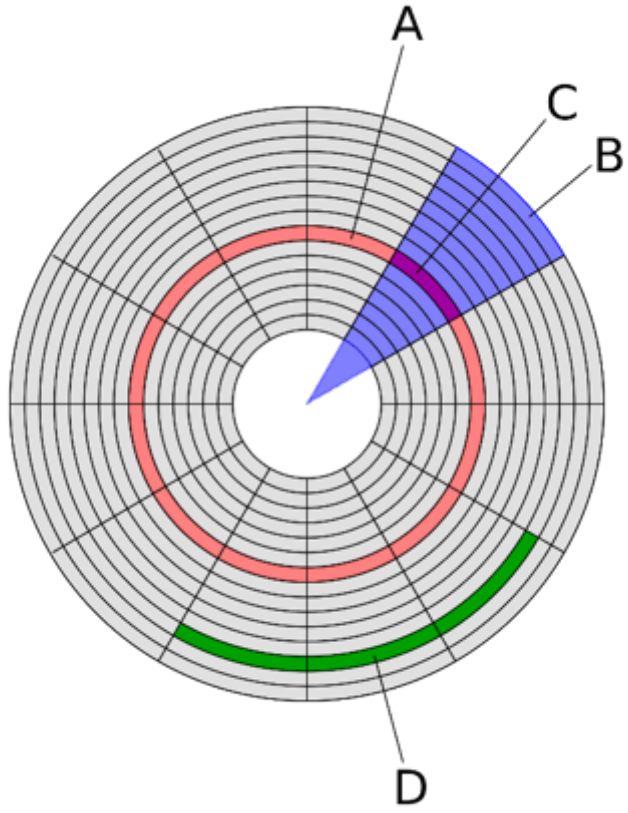

Conceptual Virtual Model

- Pseudo file system virtualizes block layout storage
- inodes (meta data) store information about the data
  - Looks like a filesystem, but can be used to manage files, blocks and to describe objects
- Mapping to blocks done by layers below virtualization level

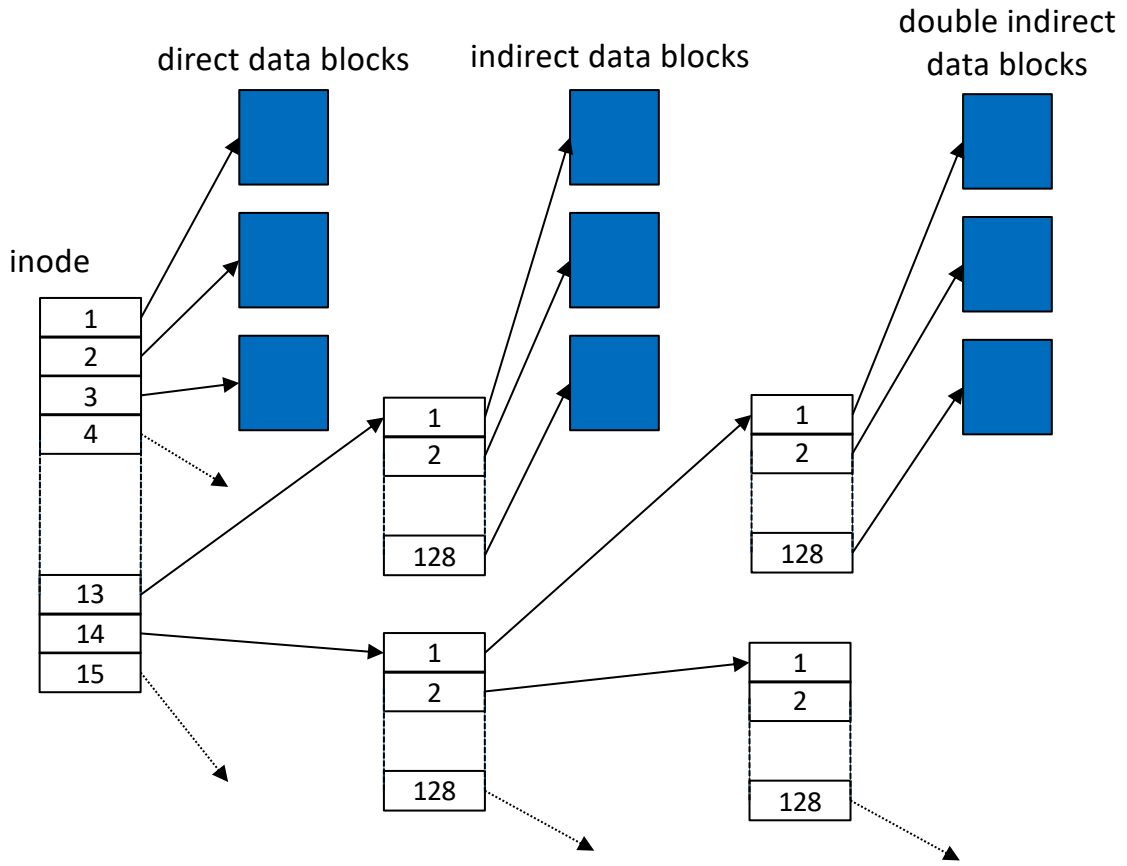

Techniques for Data Reduction

- Deduplication
  - Removing duplicates from the data
  - Sparse data; removing zero or repetitive bit pattern data
- Compression
  - Squeezing data into a smaller space
- Compaction
  - Filling the gaps
- Snaps & clones
  - Copying only the changes
- Thin provisioning & hole punching
  - Appearing smaller and recovering space

Deduplication

- General principle
  - Don’t write data you’ve seen before
- Works with files, blocks and objects
- Most techniques use “no false negatives” principle
  - Query returns either "possibly deduplicated" or "definitely not deduplicated“
- The bigger the deduplication space, the larger the metadata information required

Detecting Matches

- Simple hash & table lookup & Bloom filters
- Common pattern elimination
  - Blocks of all zeroes
  - Blocks with repetitive patterns
    - Some database technologies use non-zero blocks to indicate free space
- Techniques that work on partial matches; store only the changes
  - Can be done at file, part file, object, part object, block levels
  - Depends on the virtualization being used, and how the data is stored

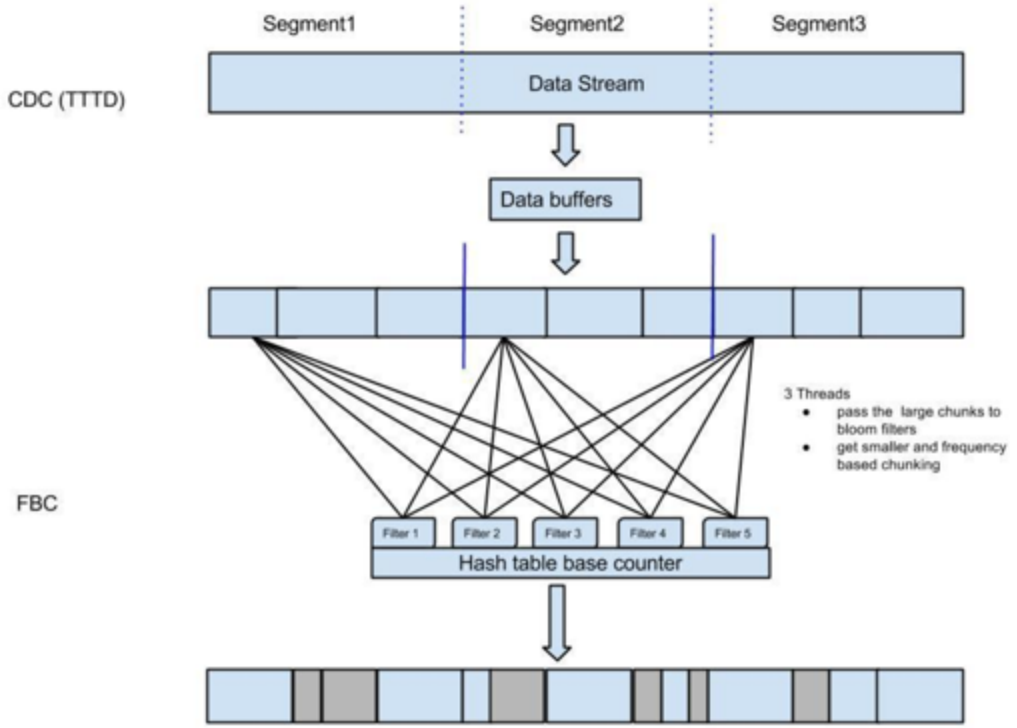

Before & After Deduplication

- Same data, but much less space
  - Metadata load likely to increase as a %age of the total space

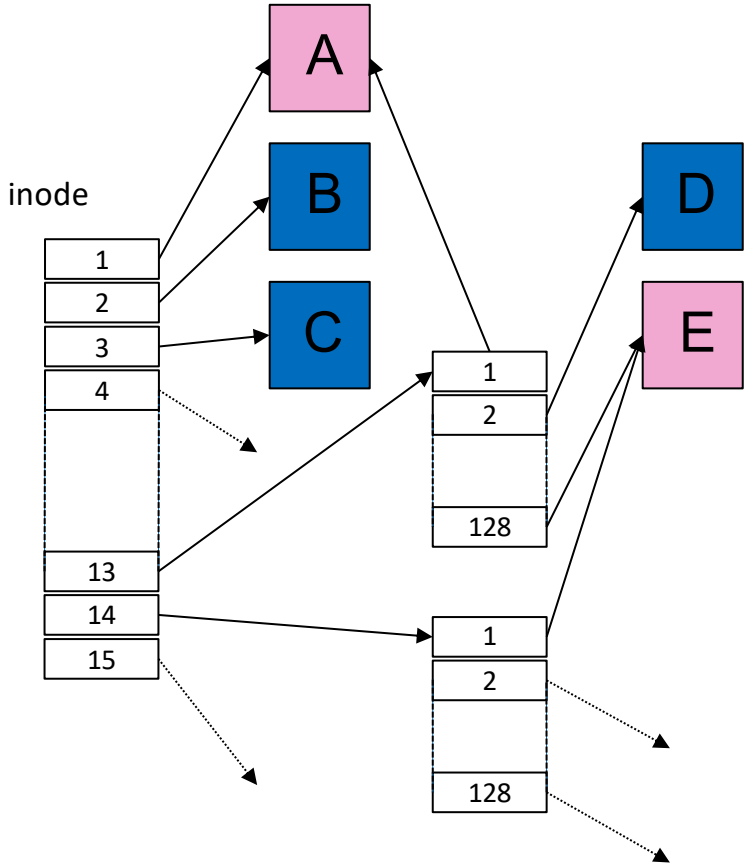

Compression

- Lossy vs lossless
  - Storage systems use lossless (for obvious reasons…)
- High compression ratios have high overheads
  - Tend to use fast & simple especially for inline
- Already compressed data generally compresses badly
  - Usually inflates the data
- Should we deduplicate then compress or compress then deduplicate?
  - Normally compress then deduplicate

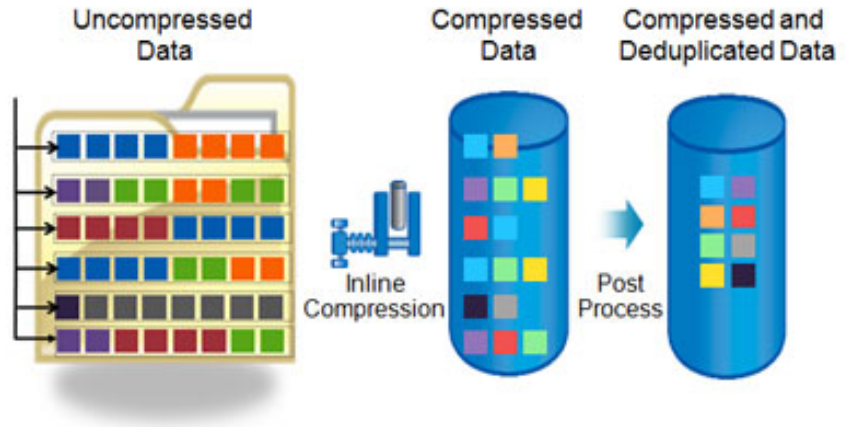

Example Compression Process

- Blocks broken into compression groups
  - Otherwise random I/O would require reading & decompressing the entire set of data first, then recompressing & writing the entire set
  - Compress block(s) in the compression group
  - If some %age smaller than the original, use compressed block(s)
  - Otherwise use original block(s)

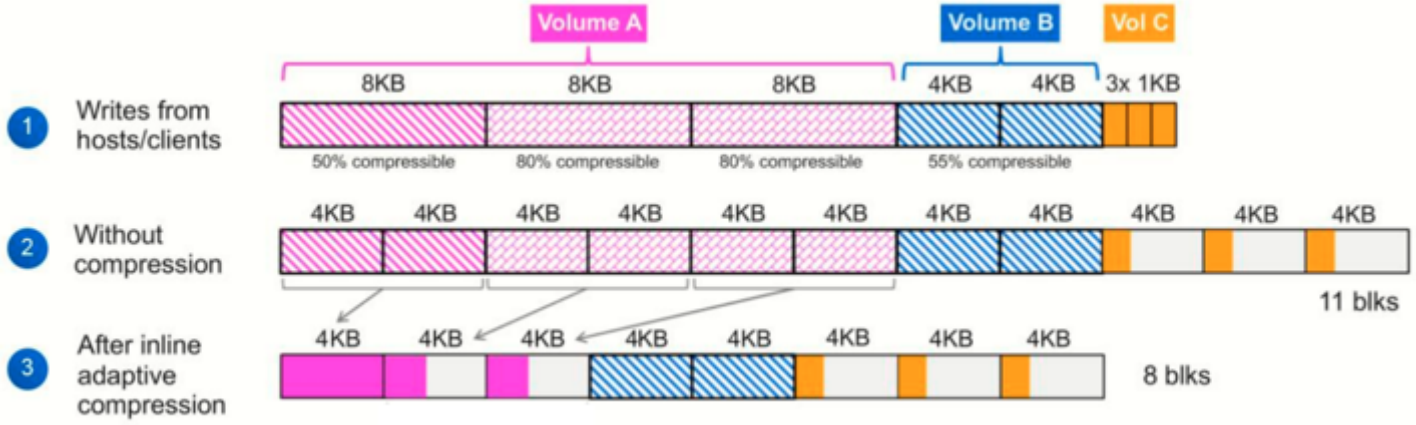

Compaction

- Filling the gaps

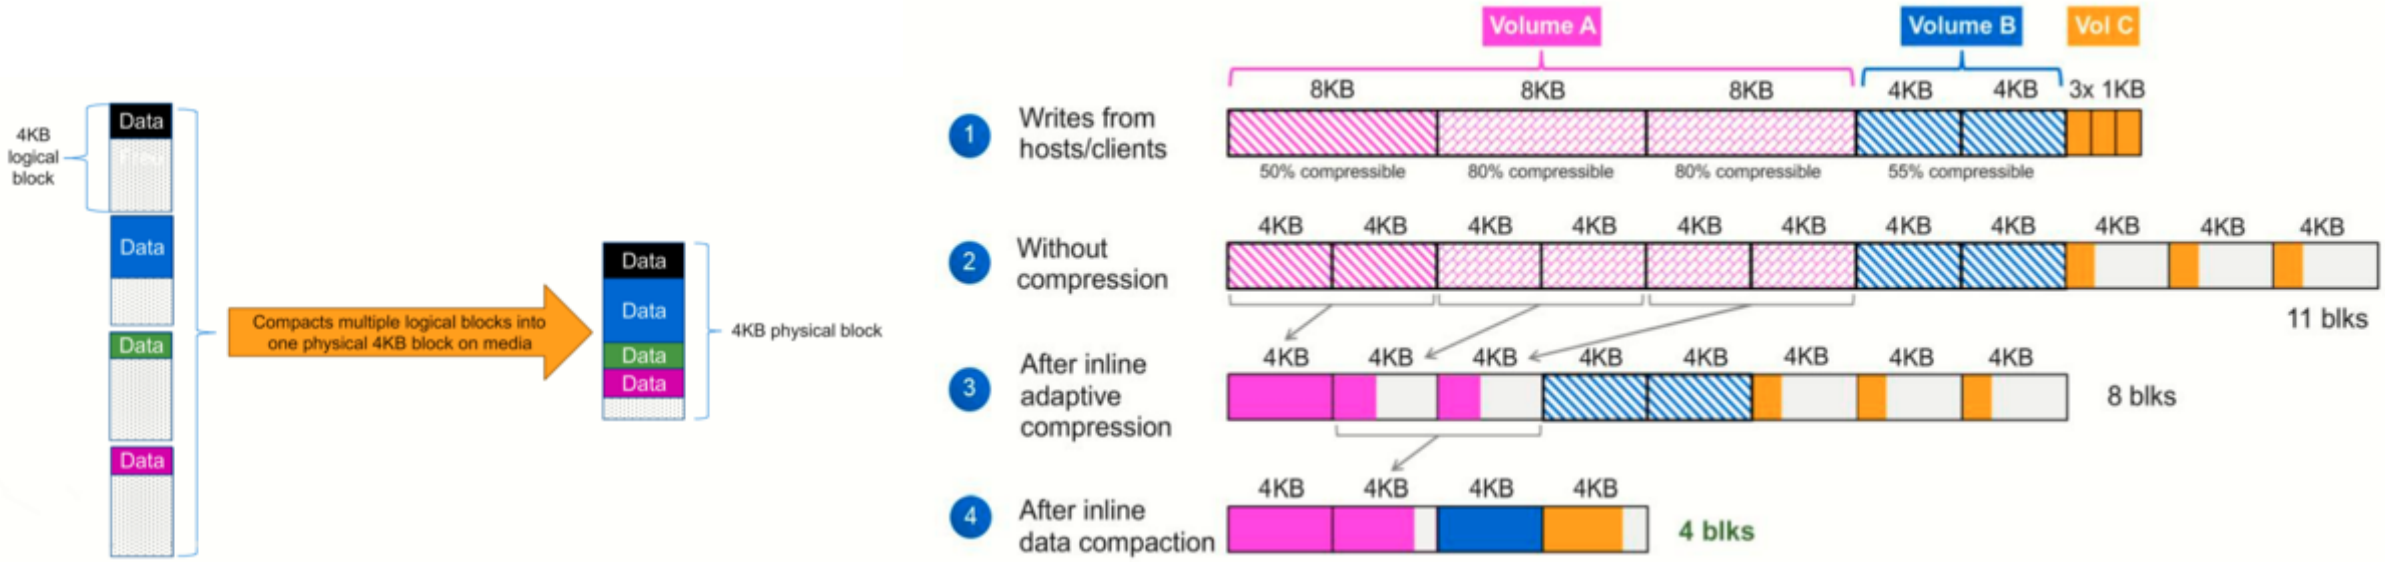

Snapshots & Clones

- Snapshots (static copies) and clones (modifiable copies) provide significant data reduction
  - Copy on Write or Redirect on Write
  - Size is original data + delta of updated blocks
  - Of course, we can deduplicate & compress the new data too

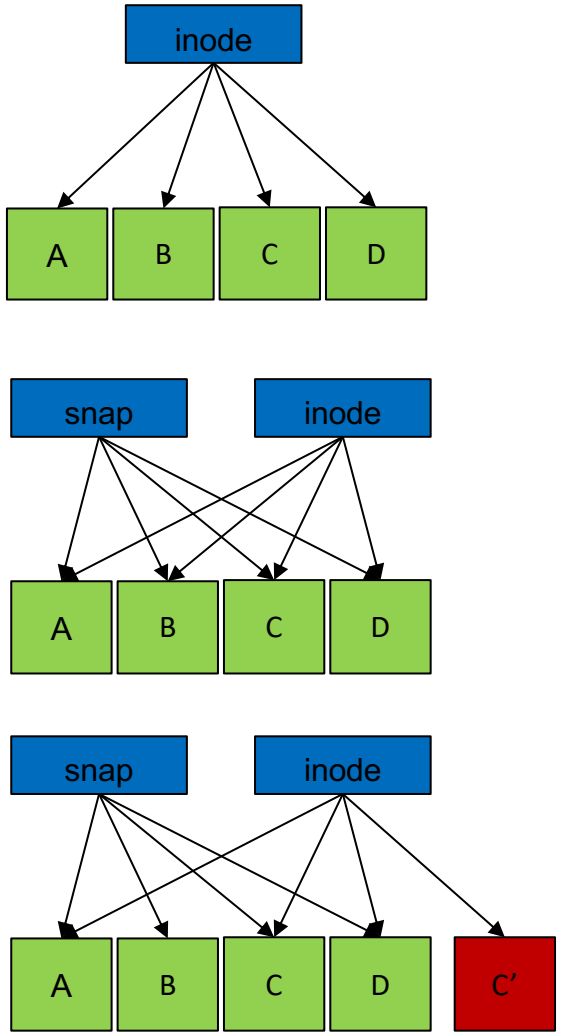

Thin Provisioning & Hole Punching

- Thin Provisioning
  - Allow unlimited virtual free space (in a file for example)
  - Only allocate blocks when required
- Hole Punching
  - “Squish” out deleted data and return it to the free pool

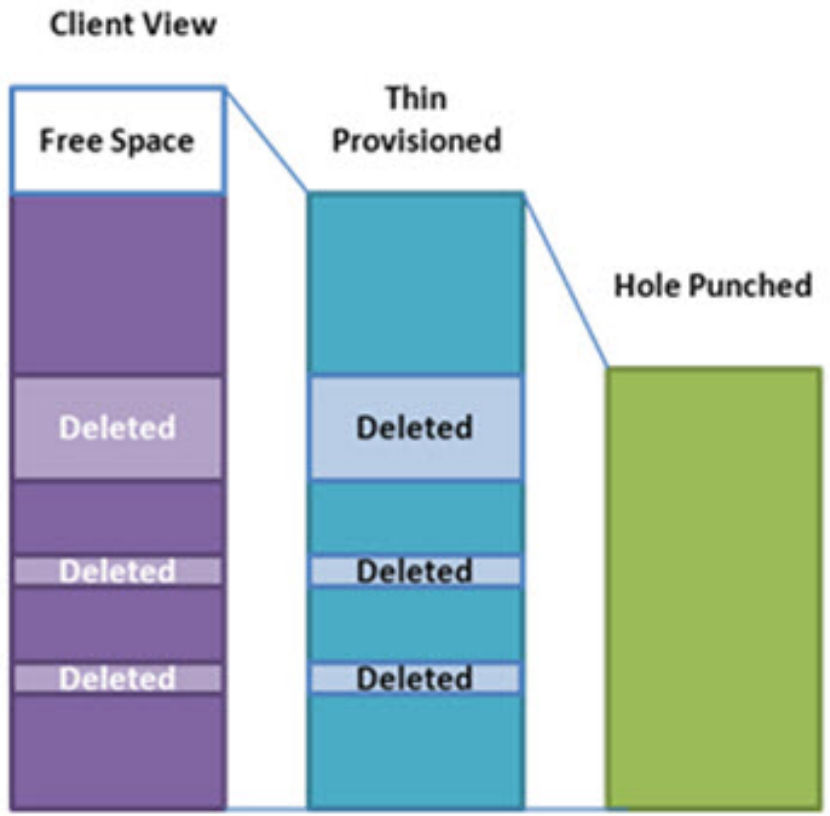

Timing & Scope; The When & Where

- Timing (applies to deduplication, compression, compaction)
  - Do reduction inline or out-of-line; i.e now or post-processed
  - Inline requires
    - Horsepower (CPU/FPGA/ASIC processing)
    - Increases latency (although aggressive caching can help)
  - Post processed requires
    - More space until the reduction is performed
    - Increase I/O activity to read and rewrite data
- Static techniques
  - Thin provisioning; doesn't cost at I/O time
- Scope (applies to deduplication)
  - Local; reduces the search space looking for duplicates
  - Global; increases the effectiveness of deduplication

Where Is All This Done?

- On the host
  - E.g compressed file formats on your laptop (zip files, JPEGs etc)
  - Application level
- At the file system level
  - NAS devices can compress and deduplicate at the file level
  - NFS and SMB offer “hole punching” and zero block/pattern block suppression
- On the storage system controllers
  - Sees all the I/O and data and can make global decisions
  - Deduplication, compression & compaction at block level
- On the device (“computational storage”)
  - Smart devices that can look like standard block devices but deduplicate, compress and compact locally
- In the cloud
  - Although this shouldn’t be apparent
- On the network
  - Packet & session based compression
- All of the above

What Savings Should I Expect?

- Depends on the data
  - Some data deduplicates well, some does not
    - Same or similar files; databases with large amounts of freespace
  - Some data compresses well, some does not
    - Precompressed data does very badly
  - Some data doesn't readily compact, other data does
    - Lots of small or empty files compact well
    - Fragments caused by compression are candidates too
- Vendors are well placed to answer this question
  - Experience across a wide variety of customer data & applications
  - Systems are often quoted in "effective space" sizes; that is TBs of data reduced space
  - Ratios of 3:1 or more possible

Is There An Effect On Performance?

- Negatives
  - Trade off between processing power and space
  - CPU overheads & added latency may be expensive to prohibitive
  - Cache sizes may be a significant factor
- Positives
  - Reduced I/O per block can improve performance & elapsed time to write/read data
  - Space savings may be substantial
- On balance
  - Techniques are being constantly refined
  - Plenty of industry experience
  - Systems that support data reduction techniques are widely available

What We Promised to Cover

- How companies end up with so many copies of the same data
- Difference between deduplication and compression – when should you use one vs. the other?
- Where and when to reduce data: application-level, networked storage, backups, and during data movement. Is it best done at the client, the server, the storage, the network, or the backup?
- What are snapshots, clones, and thin provisioning, and how can they help?
- When to collapse the copies: real-time vs. post-process deduplication
- Performance considerations

In Summary

- Deduplication
  - Removing duplicates from the data
  - Sparse data; removing zero or repetitive bit pattern data
- Compression
  - Squeezing data into a smaller space
- Compaction
  - Filling the gaps
- Snaps & clones
  - Copying only the changes
- Thin provisioning & hole punching
  - Appearing smaller and recovering space
- Techniques can be applied to all storage types
  - File, block and object

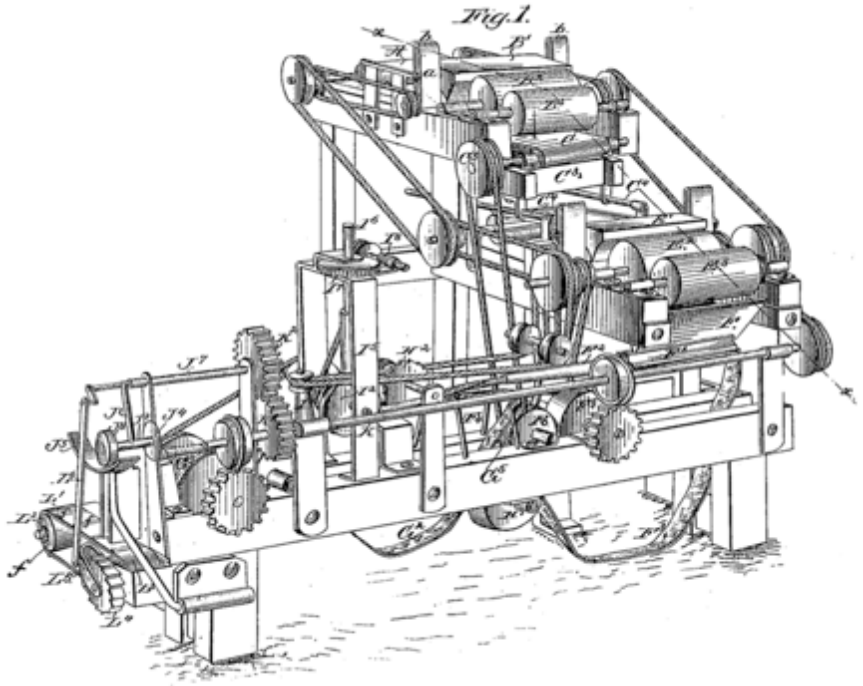

Everything You Wanted to Know About Storage But Were Too Proud to Ask Series

- There are 10 more presentations in this “Everything You Wanted to Know About Storage But Were Too Proud to Ask” series
- Topics include:
  - Storage Basics
  - Buffers, Queues and Caches
  - Storage Controllers
  - Memory
  - Storage Management
  - And more
- View them all and get access to the presentation slides at the SNIA Educational Library https://www.youtube.com/playlist?list=PLH_ag5Km-YUZhjnj8vl1u-dEaO6zyGIb6
- Watch them on our SNIAVideo YouTube channel https://www.youtube.com/playlist?list=PLH_ag5Km-YUZhjnj8vl1u-dEaO6zyGIb6

After this Webcast

- Please rate this webcast and provide us with your feedback
- This webcast and a copy of the slides will be available at the SNIA Educational Library https://www.snia.org/educational-library
- A Q&A from this webcast, including answers to questions we couldn't get to today, will be posted on our blog at https://sniansfblog.org/
- Follow us on Twitter @SNIANSF

Thank You

