



Education

# ADVANCED DATA REDUCTION CONCEPTS

Thomas Rivera, Hitachi Data Systems  
Gene Nagle, BridgeSTOR

Author: Thomas Rivera, Hitachi Data Systems  
Author: Gene Nagle, BridgeSTOR

- ❖ The material contained in this tutorial is copyrighted by the SNIA.
- ❖ 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 Education Committee.
- ❖ 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.**

- This tutorial has been developed, reviewed and approved by members of the Data Protection and Capacity Optimization (DPCO) Committee which any SNIA member can join for free
- The mission of the DPCO is to foster the growth and success of the market for data protection and capacity optimization technologies
- 2012 goals include educating the vendor and user communities, market outreach, and advocacy and support of any technical work associated with data protection and capacity optimization



## Check out these SNIA Tutorials:

- **Understanding Data Deduplication**
- **Deduplication's Role in Disaster Recovery**



## Managing Big Data Hands-On Lab HOL

*Since arriving on the scene ~20 years ago, the adoption of data reduction has become widespread throughout the storage and data protection community. This tutorial assumes a basic understanding of data reduction techniques and covers topics that attendees will find helpful in understanding today's expanded use of this technology.*

Topics will include:

- ◆ Trends in data reduction design and usage
- ◆ Practical data reduction of primary storage
- ◆ Using data reduction techniques to reduce storage network traffic
- ◆ Pervasive data reduction across storage tiers

## **Capacity Optimization Methods [Storage System]**

Methods which reduce the consumption of space required to store a data set, such as compression, data deduplication, thin provisioning, and delta snapshots

## **Data Deduplication [Storage System]**

The replacement of multiple copies of data—at variable levels of granularity—with references to a shared copy in order to save storage space and/or bandwidth.

## **Compression [General]**

The process of encoding data to reduce its size. Lossy compression (i.e., compression using a technique in which a portion of the original information is lost) is acceptable for some forms of data (e.g., digital images) in some applications, but for most IT applications, lossless compression (i.e., compression using a technique that preserves the entire content of the original data, and from which the original data can be reconstructed exactly) is required.

- Original value and justification has not changed:
  - ◆ Satisfy ROI/TCO requirements
  - ◆ Manage data growth
  - ◆ Increase efficiency of storage and backup
  - ◆ Reduce overall cost of storage
  - ◆ Reduce network bandwidth requirements
  - ◆ Reduce operational costs including:
    - › Infrastructure costs: space, power and cooling
      - Movement toward a greener data center
  - ◆ Reduce administrative costs
- New OSes and file systems are now integrating deduplication (e.g., Windows 8 Server ReFS, ZFS)

- Compression
- Deduplication
  - File (Single Instance Storage)
  - Block (Hash-Based or Delta Block)
  - Application aware
- In-line vs. post process
- Thin Provisioning

Note: Some techniques may be combined

- Dedupe and compression are similar
  - ◆ Both are dependant on data patterns
  - ◆ Both consume system resources
  - ◆ Both can optimize required storage capacity
  
- Dedupe and compression are different
  - ◆ Some data can only be optimized via dedupe
  - ◆ Some data can only be optimized via compression
  - ◆ Some data can be optimized via dedupe **and** compression
  - ◆ Some data cannot be optimized at all
  
- Dedupe and compression are complementary
  - ◆ But some knowledge about the data pattern is required




The scope of data reduction is broadening:

- **Primary Storage**
  - ◆ Reduced physical capacity for storage of active data
- **Replication**
  - ◆ Reduced capacity for disaster recovery and business continuity
- **Data Protection**
  - ◆ Reduced capacity for backup with longer retention periods
- **Archivals**
  - ◆ Reduced capacity for data retention and preservation
- **Movement/Migration of data**
  - ◆ Reduced bandwidth requirements for data-in-transit

## Effective for specific workloads

- Acceptable performance - a factor for compression or deduplication
  - Post-processing
  - Inline ingestion
  - Network-based
- Deduplication works best with applications with high data redundancy
  - Virtual servers and desktops
  - Collaborative file “sharing”
  - Email (Software SIS replacement)




# Deduplication Savings Expectation

Deduplication Savings Depend on Use Case and Time

Primary/replicated storage has less duplicate data

Periodic archives have moderate duplicate data

Repeated backups have significant duplicate data



## ➤ Array Cache

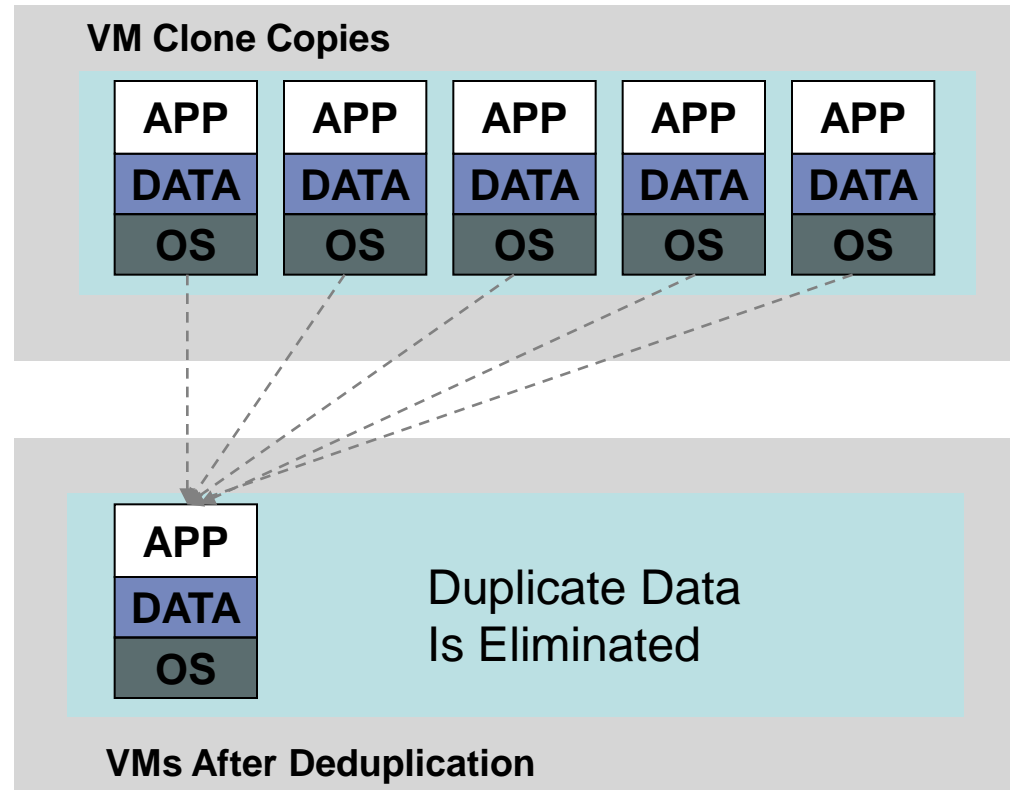
- ◆ Intelligent cache can be “dedupe-aware”
- ◆ Hot data is cached with dedupe attributes
- ◆ Reduces rotating media latencies
- ◆ Example: Virtual Desktop “boot storms”

## ➤ Solid State Drives


- ◆ Deduplication/Compression helps offset the higher cost/GB of SSD's
- ◆ High performance applications with highly redundant data
- ◆ Random nature of (block-level) deduplicated data I/O a good match to SSD

# Primary Storage Considerations

- Balance the tradeoff between cost savings and performance impact
- Some workloads lend themselves better to data reduction
- Walk before you run
  - ◆ Use estimation tools
  - ◆ Perform POCs
  - ◆ Implement one workload at a time

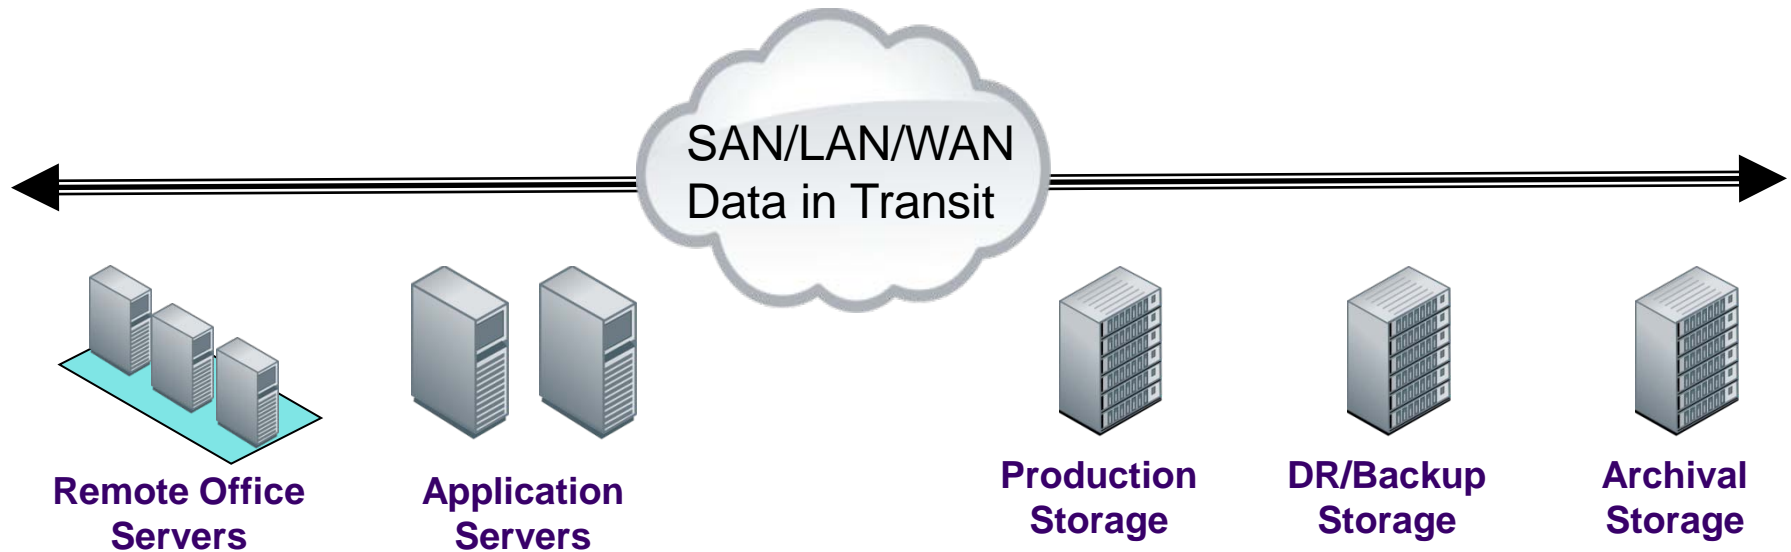


# Data Reduction and Replication



- Can be one-way, two-way, multi-node, or cascade
- Optimized location(s) can be determined by Users
- Data reduction makes replication more affordable
- Data reduction enables replication on constrained networks

- Focus on your Service Level Agreements (SLAs) first
  - ◆ Needs to meet window for *Replication*
  - ◆ Needs to meet SLA for *System Recovery or Data Restore*
  
- Is it Necessary to Optimize All Data?
  - ◆ Mission-critical applications
  - ◆ May have regulatory issues for some data
  - ◆ Some data types not conducive to data reduction
  - ◆ Replicate incremental changes only, without other optimization



## ➤ Increased SAN/LAN/WAN Efficiency

- ◆ Compression/deduplication for data in flight
- ◆ Transfer data references instead of data objects
- ◆ Shorten data transfer times by sending less data



## The Original Promise


- Faster data recovery from disk
- Reduction in D2D cost per terabyte stored
- Reduction in D2D backup storage footprint
- Less network bandwidth required for D2D backups
- Makes longer retention possible

## What's New?

- Wide use as part of backup software
- Scalability of deduplication appliances
- Deduplication across appliances
- Deduplication to tape

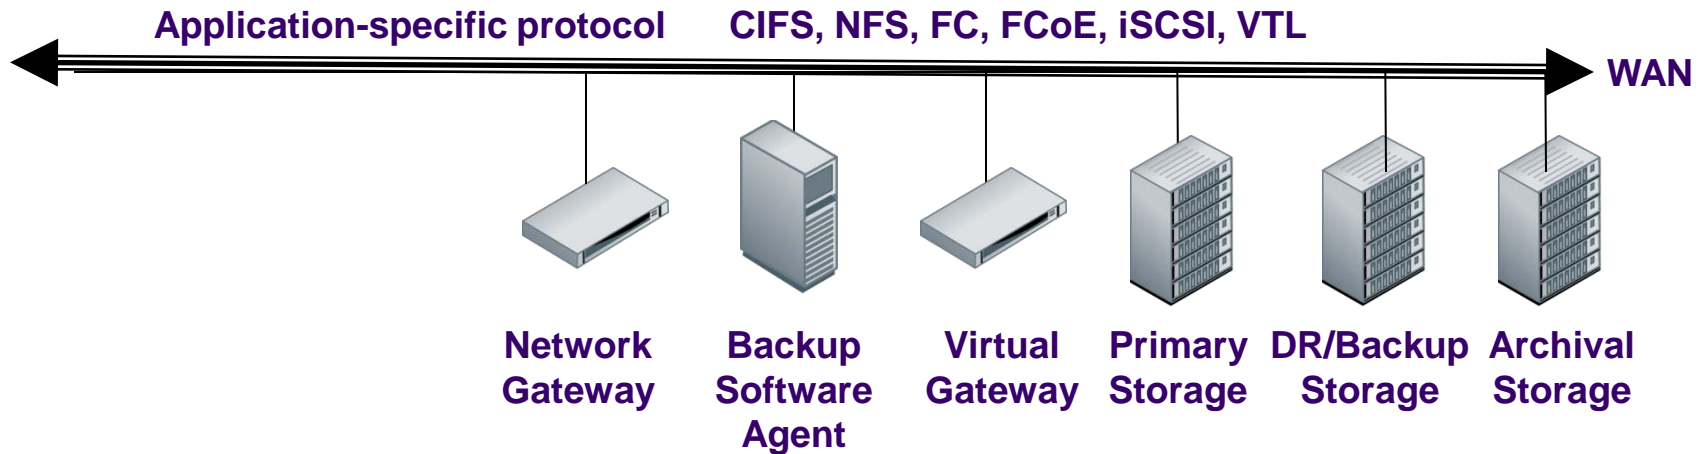
# Backup Considerations

- Hardware or software deduplication?
- Source or target deduplication?
- Variable or fixed-length deduplication?
- File or sub-file deduplication?
- Compression WITH deduplication?
  
- Answers depend on the problem you are trying to solve



- Data reduction can reduce the cost of online archive repositories
- These repositories are often required for regulatory compliance
- No standard exists today for “approved” use of data reduction techniques with regulatory data
- Vendors should provide assurances that the ability to retrieve data in its original form is not impaired

# Consideration Matrix for Data Reduction Location



	Network Gateway	Backup Software Agent	Virtual Gateway	Primary Storage	DR/Backup Storage	Archival Storage
Reduce Network Traffic	✗	✗	✗		✗	
Reduce Physical Capacity		✗		✗	✗	✗
Reduce Backup Time		✗				
Reduce Recovery Time		✗			✗	
Reduce Replication Time	✗		✗	✗		
Reduce Media Latency				✗		

- The primary value of data reduction is in reducing costs and helping manage data growth
- The scope of data reduction is broadening
  - All storage tiers
  - Various levels of granularity
  - Bandwidth reduction
  - Data subject to regulatory rules
- New use cases and new technologies bring new challenges
  - And new opportunities!

- Please send any questions or comments on this presentation to SNIA: [tracktutorials@snia.org](mailto:tracktutorials@snia.org)

**Many thanks to the following individuals  
for their contributions to this tutorial.**

**- SNIA Education Committee**

**Data Protection and Capacity Optimization (DPCO) Committee:**

**Mike Dutch  
Larry Freeman  
Tom McNeal  
Gene Nagle**

**Ron Pagani  
Thomas Rivera  
Tom Sas  
SW Worth**



It's easy  
to get  
involved  
with  
the DPCO !

- Join us
- Find a passion
- Gain knowledge & influence
- Make a difference

[www.snia.org/dpco](http://www.snia.org/dpco)