Flash Storage for Backup, Recovery and DR

Ron Herrmann/IBM

SNIA Legal Notice

- The material contained in this tutorial 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 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.

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Abstract

- Flash storage for backup, recovery and DR
  - This session will appeal to Data Center Managers, Storage Managers, Server Administrators and those that are seeking a fundamental understanding of how flash storage can help with backup, recovery and disaster-recovery. The session will provide an overview of backup, recovery, DR, deduplication and virtual-tape libraries and discuss the common pain points in these environments. Also, this session will show how the introduction of flash storage to the backup and recovery realm can dramatically improve performance and reliability.

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Data Protection and Recovery (DP&R) Overview
- Common Bottlenecks
- Examples of using Flash Storage in the DP&R Environment
- Flash Storage and Deduplication/VTLs

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Backup Environment Components

- Backup Application (BUA).
  - A package of software designed to process and catalog files backed up from clients Larger vendors are Symantec (Netbackup), IBM (TSM), EMC (Networker) and CommVault
- Backup Server.
  - The server(s) hosting the backup application. This server has network connections to clients and media (disk and tape)
- Backup application database.
  - The index where all of the client backup records are stored
- Backup Client.
  - The server, VM, or application being backed up. The client has a network connection to the backup server and an agent is typically installed
- Backup media.
  - Typically tape, disk or virtual tape library. This is where the client backup data is stored
- Onsite and Offsite.
  - Typically, two copies of all backups are made. The onsite copy is available to do normal local application restores and the offsite copy is created to keep at a remote location for disaster recovery (DR)

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Typical Backup Environment

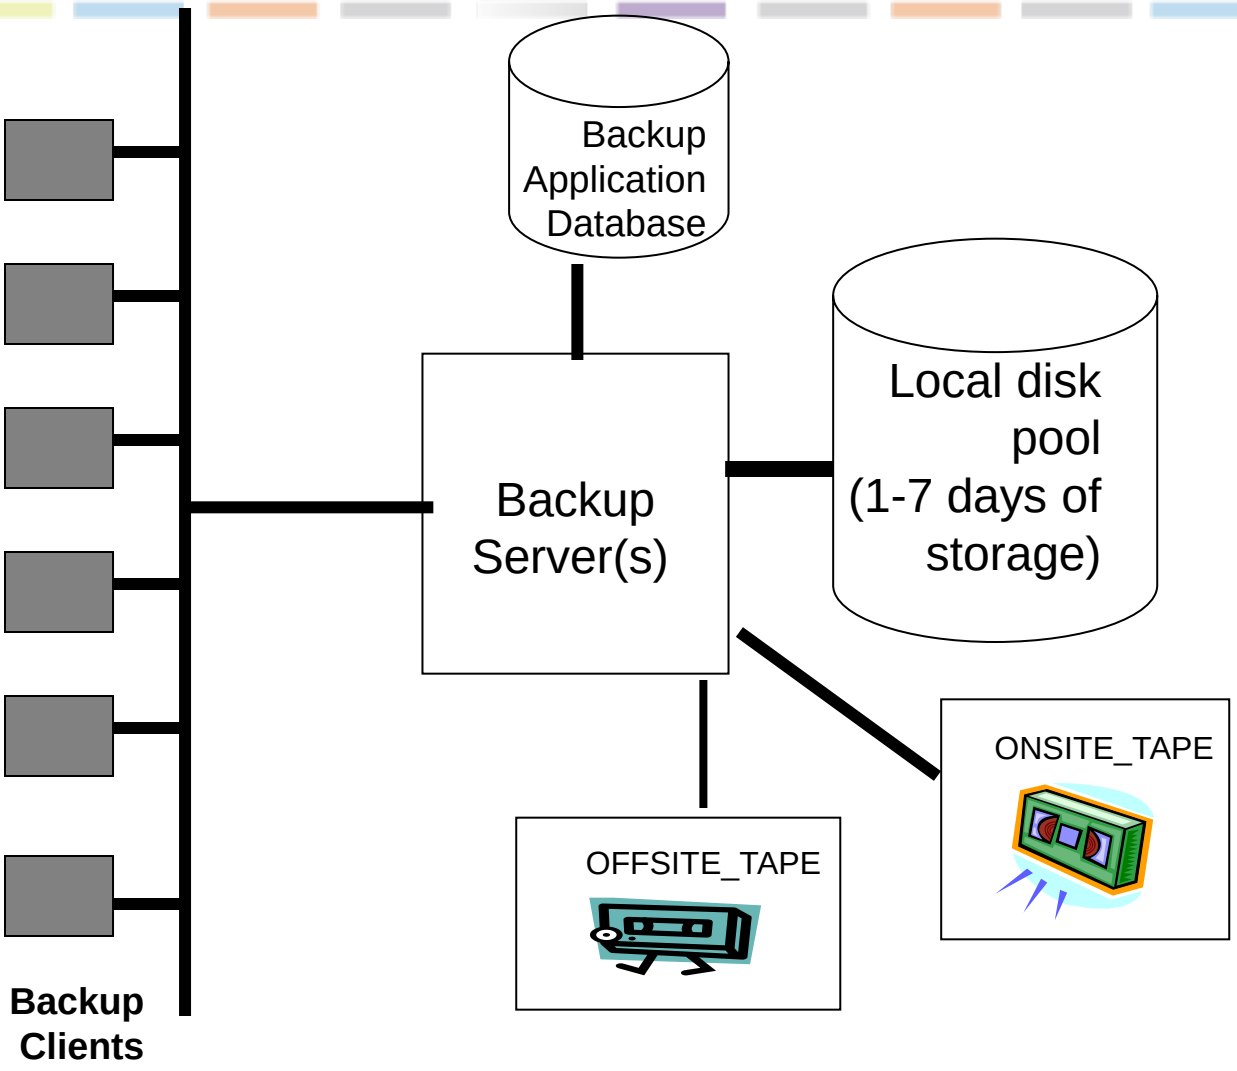

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Backup and Recovery Pain Points

- Backups do not finish in time to get tapes offsite in the morning
- Application performance suffers due to long backup times – backups still running during production hours
- Restore times can be painfully slow
- Disaster recovery testing is time consuming and unsatisfying
- Management and storage of multiple tape copies is costly

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Some Backup and Recovery Bottlenecks:

- Backup application database (or catalog) must be updated real-time for each file backed up
- Applications, such as large databases, can be slow to backup due to workload and underlying disk systems
- Disk pool speed and network is critical to accepting backups from Clients
- Disk pools feed the tape pools. The speed of database and disk pool is critical to the speed of creating tapes

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Flash Storage to the Rescue!

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

CPU performance **up 10x** this last decade

Storage has grown capacity but unable to keep up in performance

Systems are now Latency & IO bound resulting in significant performance gap

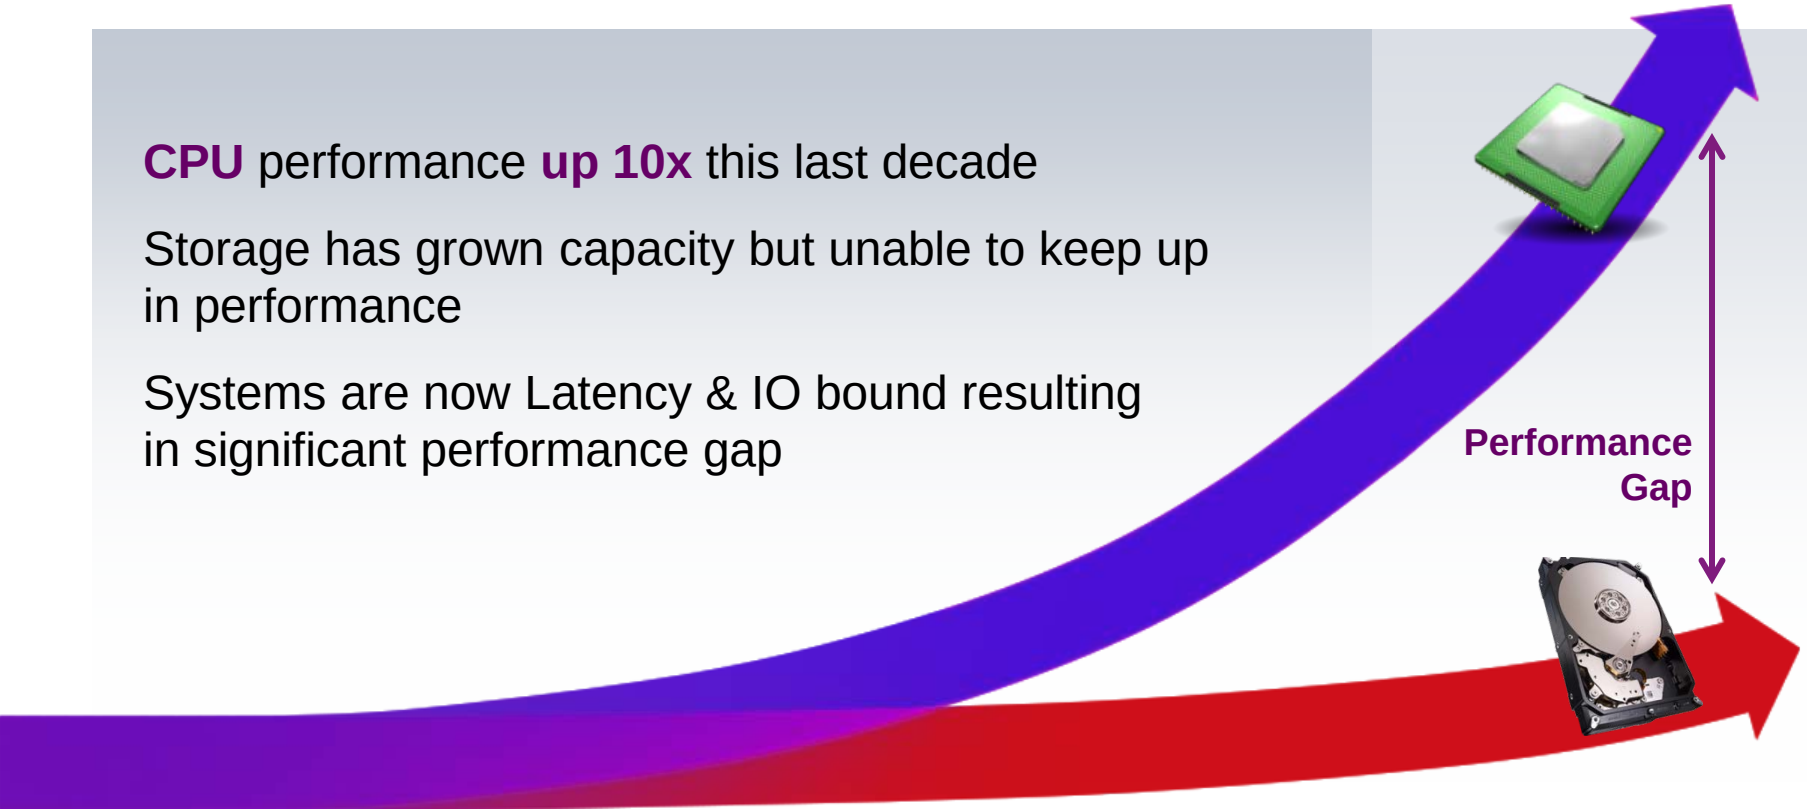

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Flash Keeps up with the CPU!

Servers, Applications and Databases are FASTER!

Benefits & economics outweigh disk

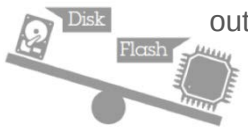

No application Or architecture Changes

Reduce floor space, power & cooling

200µs = 50% CPU – BEST!
500µs = 30% CPU
1000µs = 20% CPU

CPU Utilization & App. Efficiency **4%**

Total Application Processing Time **5,200µs (5.2ms)**

5,000µs (5ms)

200µs (.2ms)

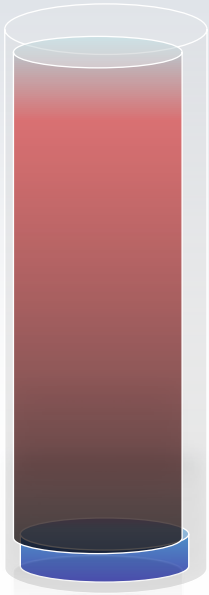

Disk/Hybrid/SSD

Increased CPU equals more work for less money!

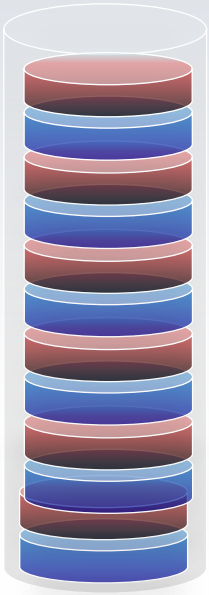

Flash Storage

CPU Utilization & App. Efficiency **50%**

Total Application Processing Time **400µs (.4ms)**

200µs (.2ms)

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Backup/Recovery Accelerated with Flash Storage

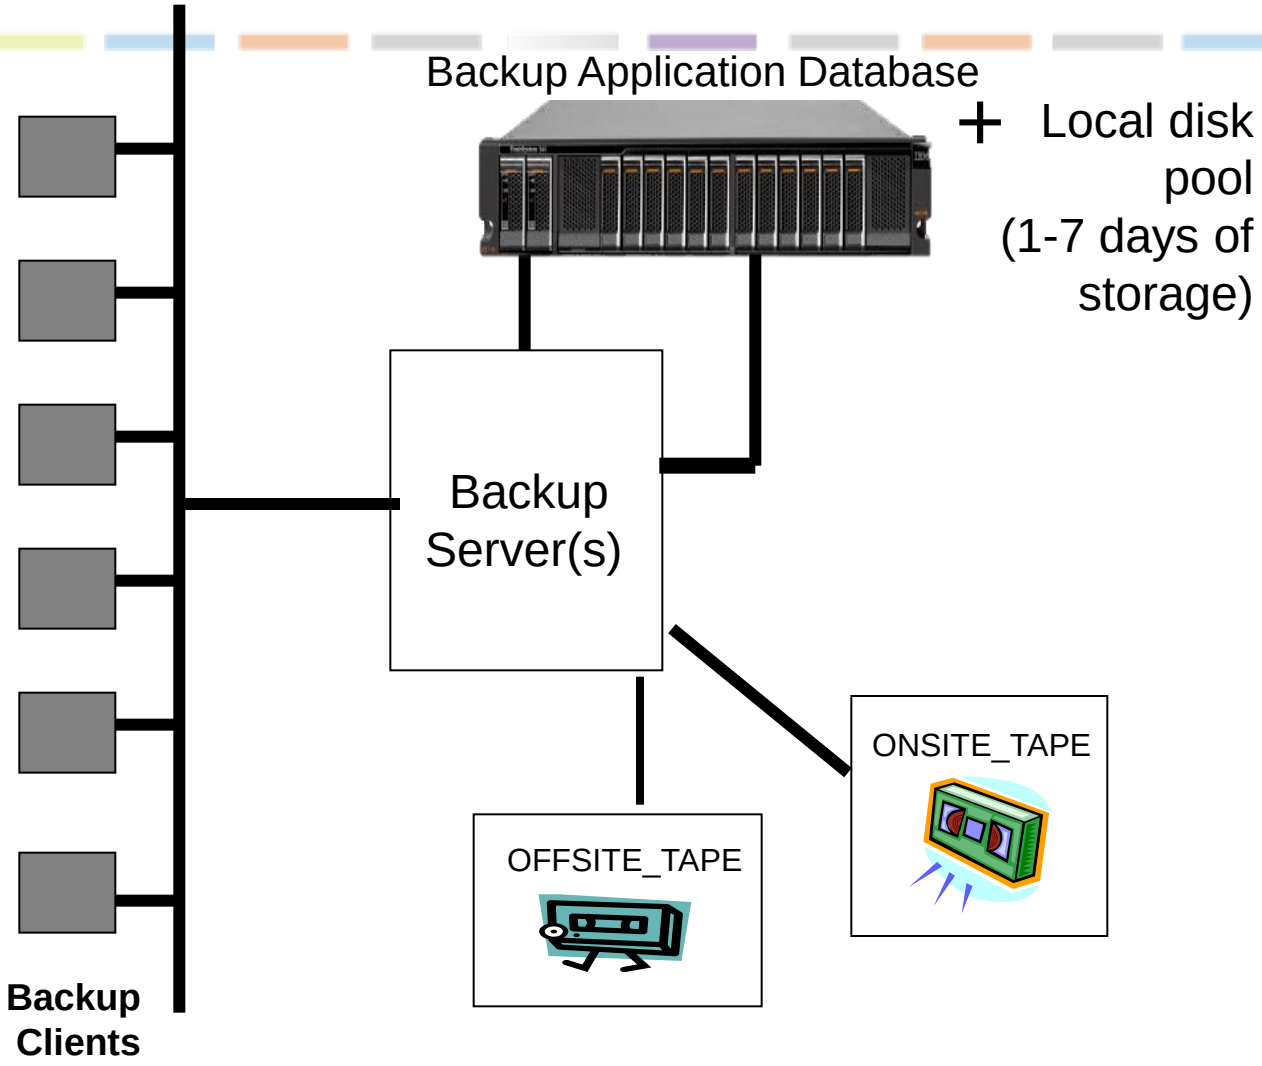

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Backup and Recovery Environment Accelerated with Flash Storage

- Move BUA database onto flash storage LUNs.
- Move BUA disk pool to flash storage LUNs.
- Increased client backup speeds into flash storage disk pool.
- Increased client restores from flash storage disk pool.
- Increased tape creation speeds – data offsite faster…
- Similar benefits for all major backup applications.

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Back-up and Recovery – Back-up at flash speeds

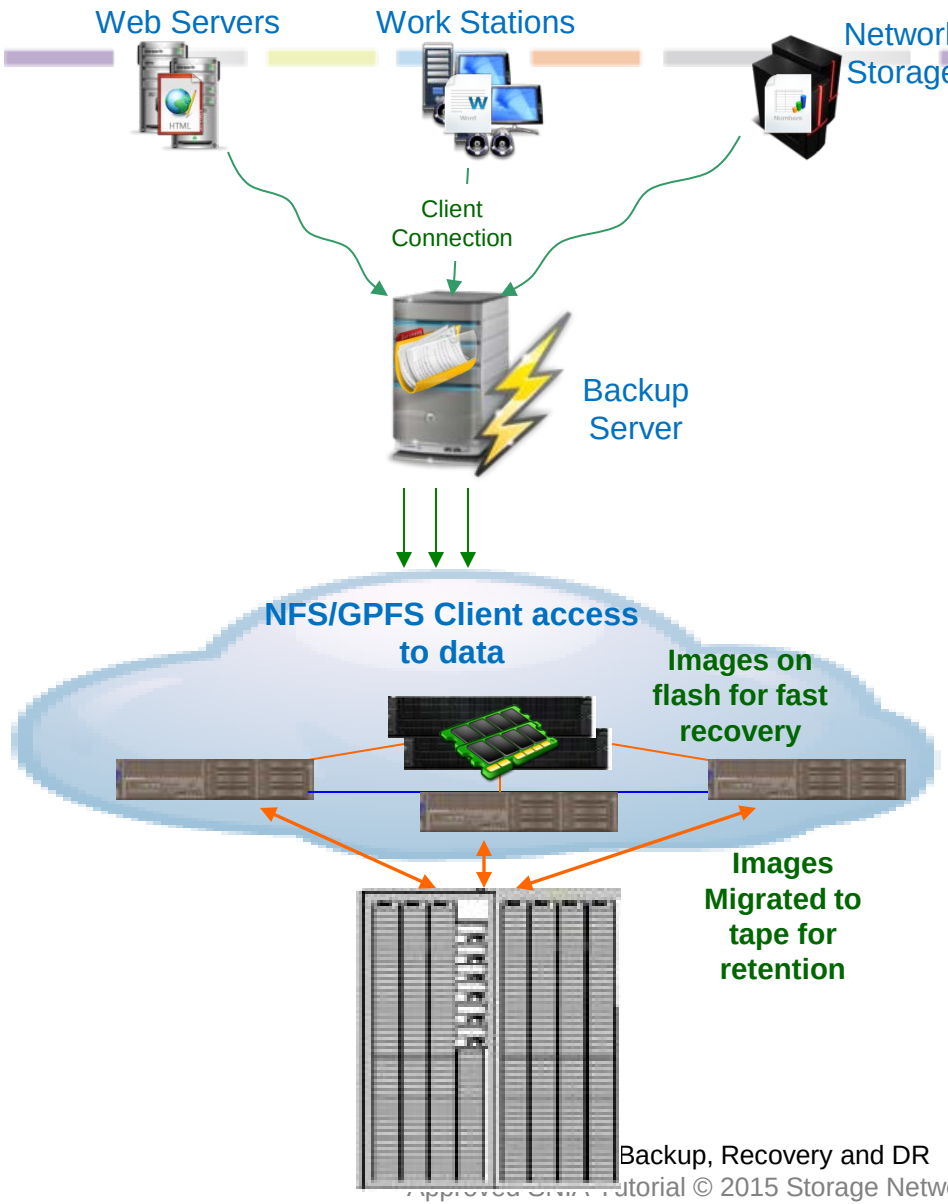

- Configuration
  - 200TB Flash Storage
  - BaR
  - 3 PB Long Term Retention
- Full Back-ups retained on disk until incremental
- Incremental backups remain for 1 week on Flash
- Full and Incrementals retained on tape by Policy
- Near line access to all of the data.
- Up to 75% improvement in Back-up time
- As low as 1/20th the cost of storing back-up images

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

What is a Virtual Tape Library (VTL)?

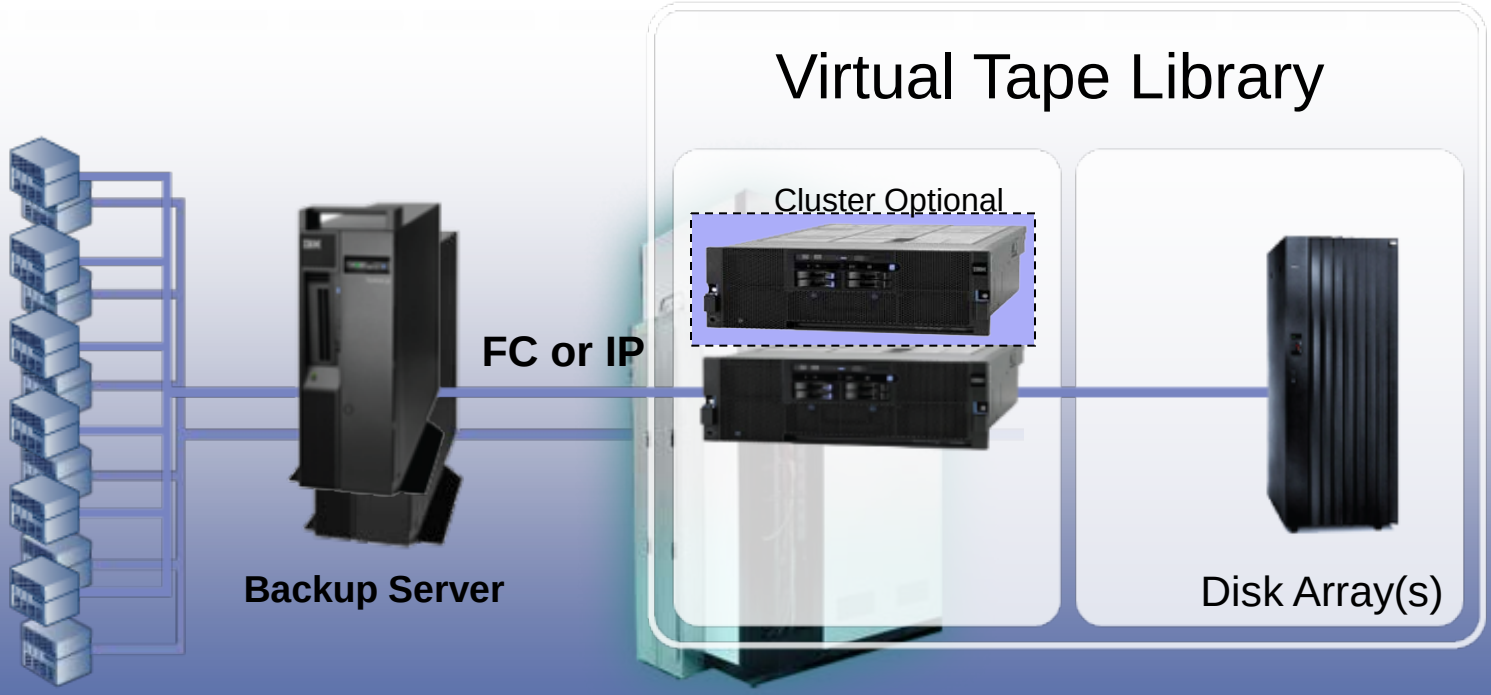

- **Software solution that typically resides on standard server.**
- **Performs VTL, NAS, or OST, Compression, and Deduplication and replication function.**
- **Uses disk array as the backup medium.**
- ***Licensed on the physical disk capacity (TB tiers) and features (cluster and replication).***

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Deduplication is Very Useful for Backup and Recovery….

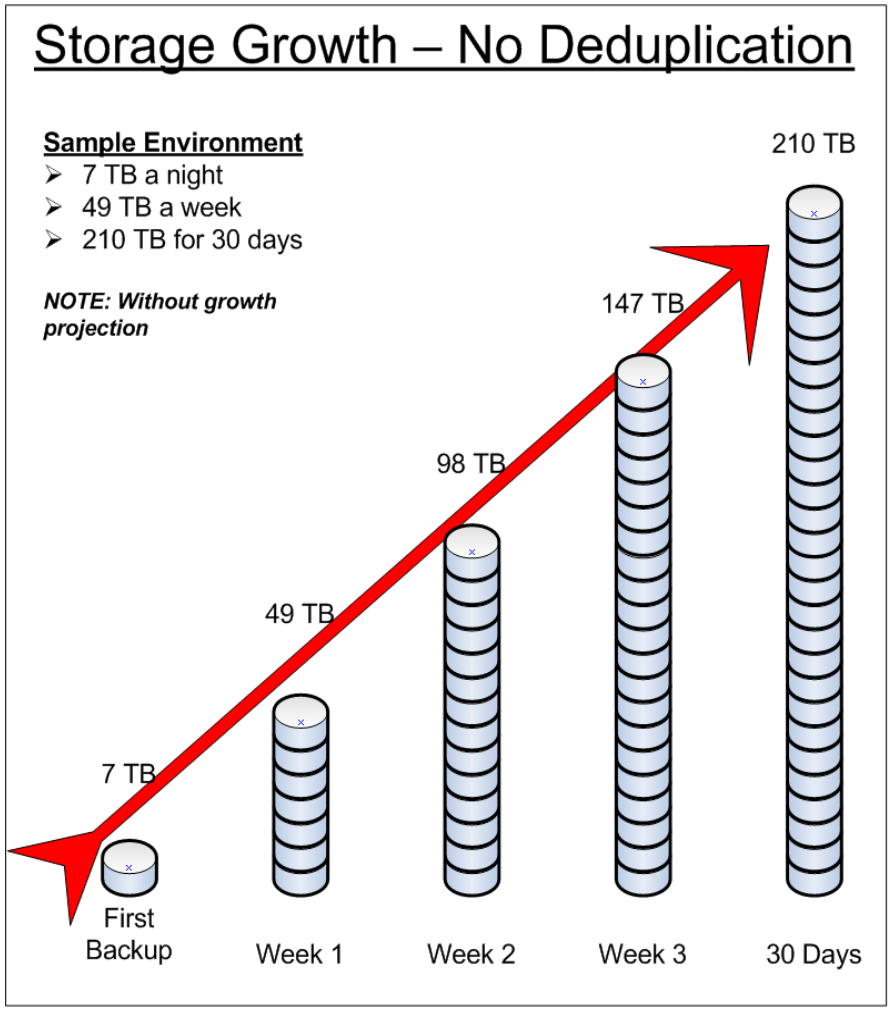

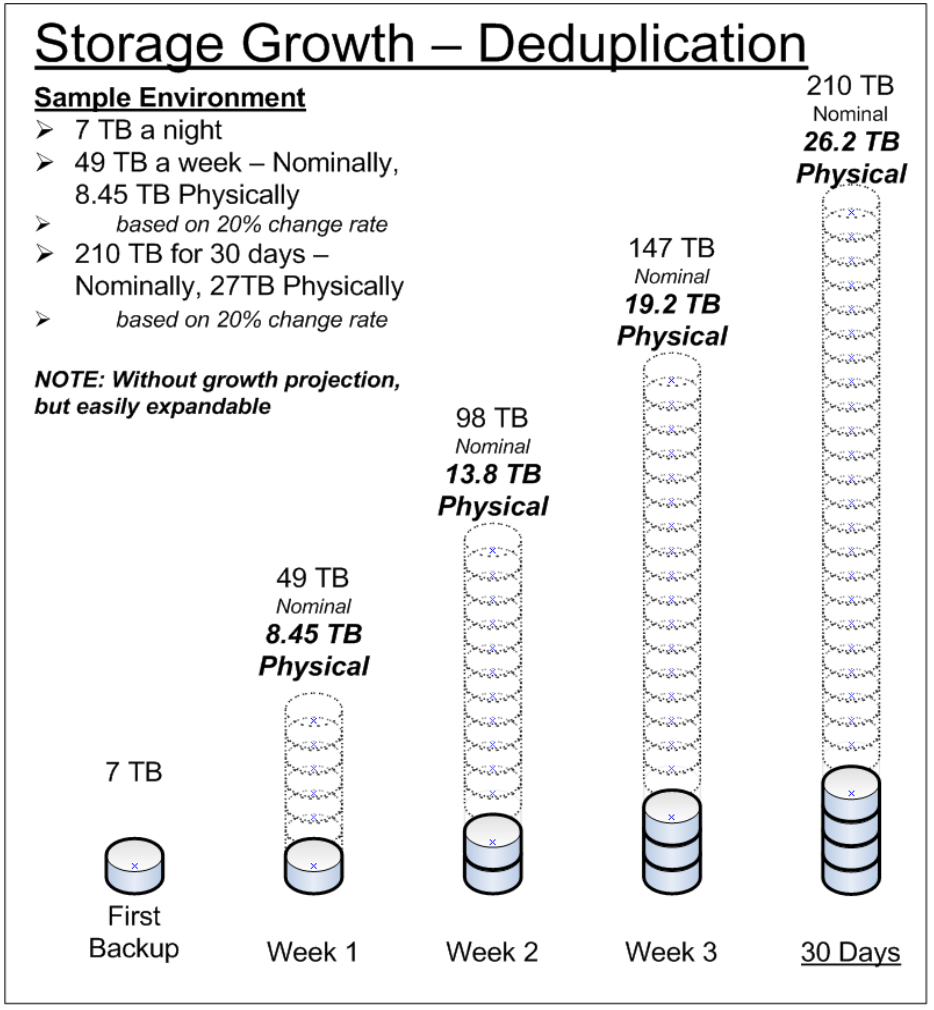

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

DR Environment Accelerated with Flash and VTL

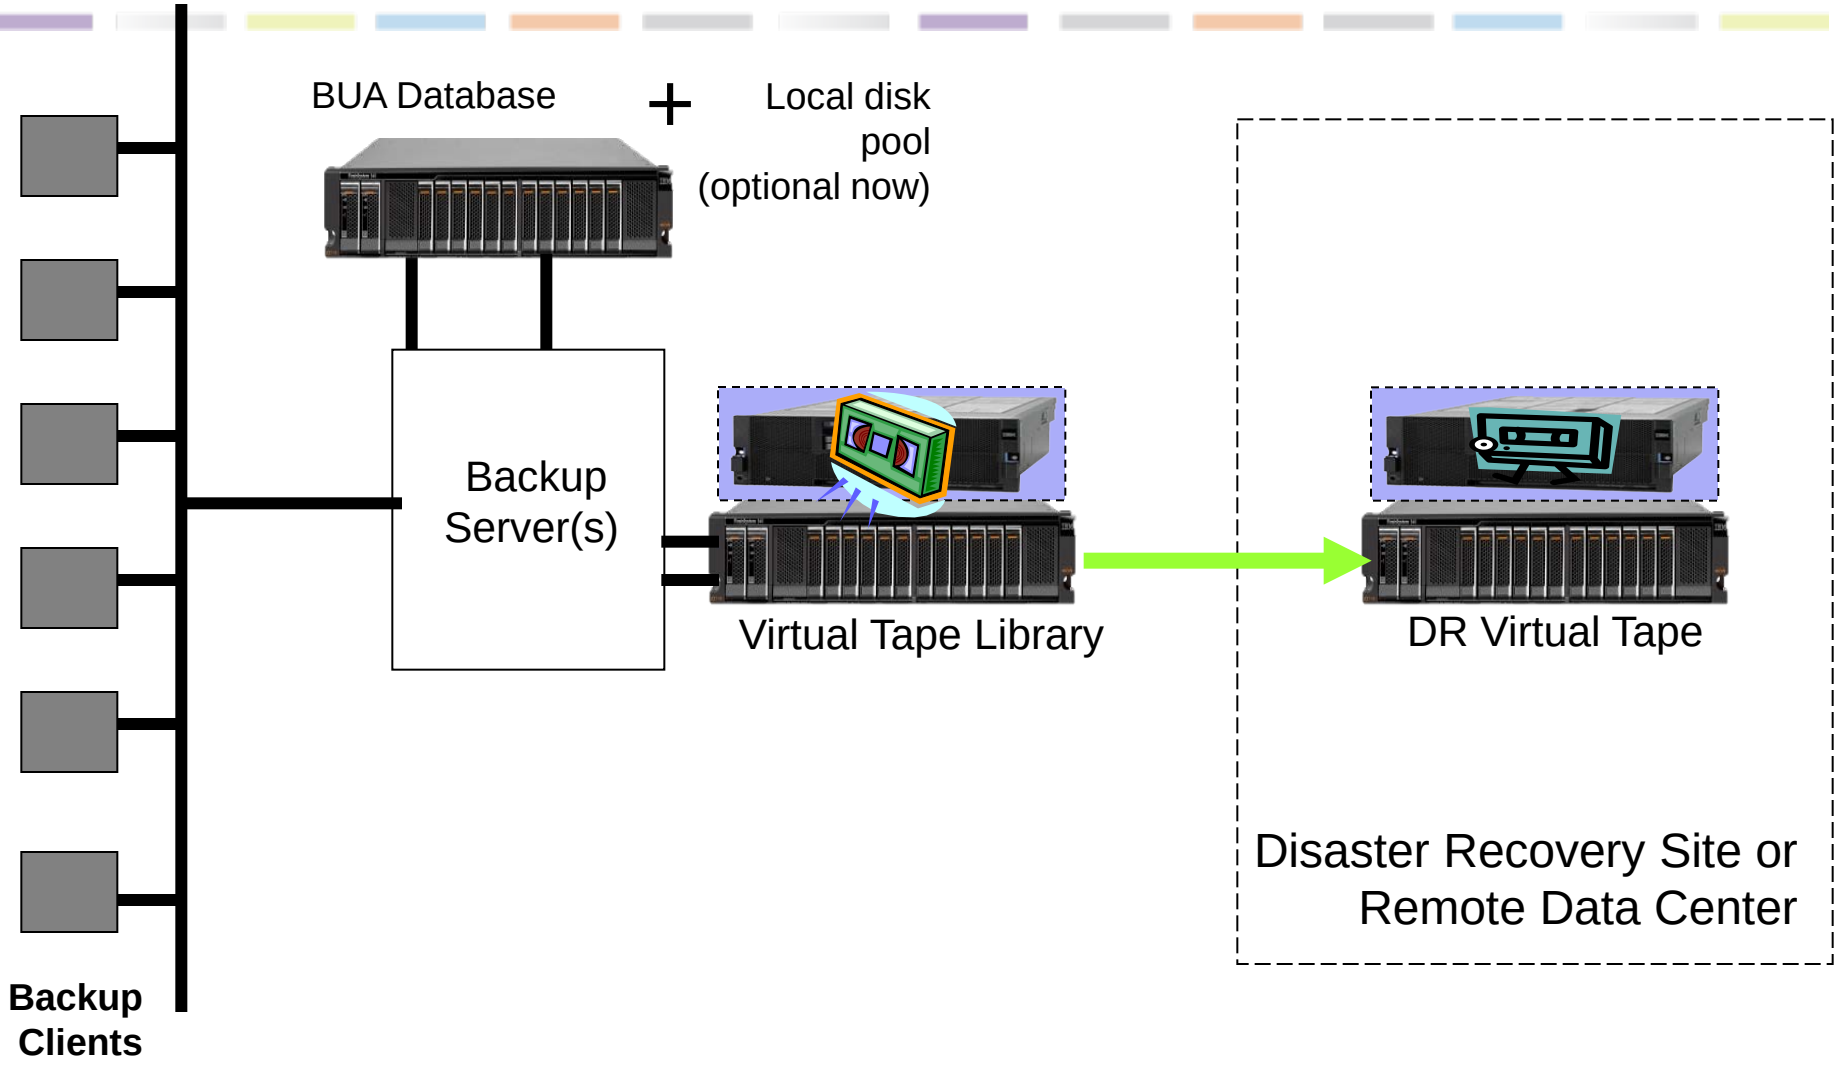

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Don’t Forget the Benefits of Putting the Applications on Flash Too!

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

SQL cluster

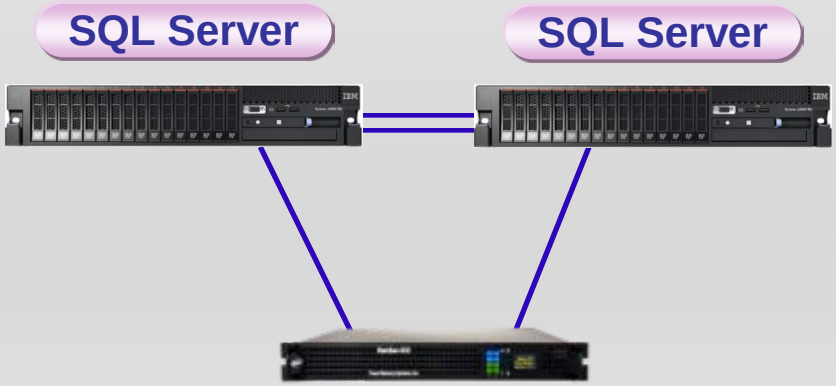

Flash Storage

Problem

- Experiencing pain with JDE BD loads / backups / restores
- Needed better system performance for the end user

Solution

- Installed flash storage into a SQL DB, clustered, running Oracle JDE
- Included Oracle OLAP processes

Benefit

- Backup Time improved from 5 hours to 42 minutes
- Restore Time improved from 6.5 hours to 1.2 hours
- Batch times went from 7:30 hours to 2:37 and 17:47 to 7:07

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

19

Questions?

and

Thank you!

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.

Attribution & Feedback

The SNIA Education Committee thanks the following Individuals for their contributions to this Tutorial.

Authorship History

Ron Herrmann/June 22, 2015

Updates:
Ron Herrmann/July 22, 2015

Additional Contributors

Joseph L. White
Thomas Rivera

Please send any questions or comments regarding this SNIA Tutorial to ***tracktutorials@snia.org***

Approved SNIA Tutorial © 2015 Storage Networking Industry Association. All Rights Reserved.