Education

Trends in Data Protection and Restoration Technologies

Jason Iehl, Network Appliance
Michael Rowan, Consultant

SNIA Legal Notices

- The material contained in this tutorial is copyrighted by the SNIA
- Member companies and individuals may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced without modification
  - The SNIA must be acknowledged as 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

Storage Networking Industry Association

- SNIA is the trade group for storage networks
  - "ensuring that storage networks become complete and trusted solutions across the IT community" - http://www.snia.org
  - SNIA "Dictionary of Storage Networking Terminology" is an excellent resource http://www.snia.org/dictionary
- SNIA Tutorials are on-line:
  - http://www.snia.org/education/tutorials

About SNIA and the DMF

About the Storage Networking Industry Association (SNIA)

- SNIA's primary goal is to ensure that storage networks become complete and trusted solutions across the IT community
- For additional information about SNIA see www.snia.org
- SNIA's "Dictionary of Storage Networking Terminology" is online at www.snia.org/dictionary

About the SNIA Data Management Forum (DMF)

- The DMF is a sub-group of SNIA acting as the worldwide authority on Data Management, Data Protection and ILM www.snia-dmf.org
- The DMF is a collaborative storage industry resource available to anyone responsible for the accessibility and integrity of their organization's information.

DMF		
Data Protection Initiative (DPI)	**Information Lifecycle Management Initiative (ILMI)**	**Long term Archive and Compliance Storage Initiative (LTACSI)**
Defining new approaches and best practices for data protection and recovery	**Developing, teaching and promoting ILM practices, implementation methods, and benefits**	**Addressing challenges in developing, securing, and retaining long-term digital archives**

Trends in Data Protection and Restoration Technologies

- We will cover the following technologies:
  - Application Consistency
  - Snapshots
  - Backup to disk
  - VTL
  - CDP
  - Replication
- What each is, its value, what is good and not so good, integration into infrastructure.

Outline

- Traditional data protection challenges
- Application Consistency
- Snapshots
- Backup to disk
- VTL
- CDP
- Replication
- Evolving data protection
- SNIA resources

Traditional Data Protection Challenges

- Backup window
- Failed, inconsistent, unreliable recovery
- Recovery time too long (poor RTO)
- Negative production impact
- Protection gaps (poor RPO)
- Disaster recovery
- Regulatory compliance and corporate governance
- Costly and inefficient
- Disruptive to change

Protection Based on Recovery

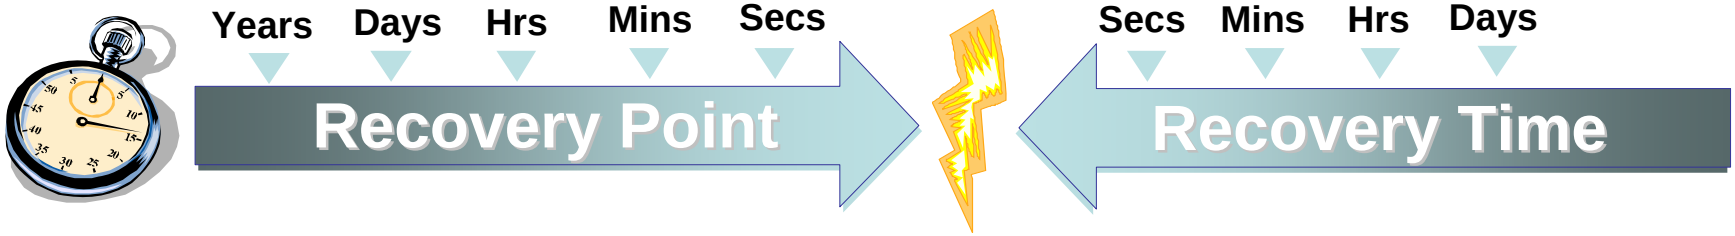

Protection Methods

Tape Backups
Vaults
Archival
Capture on Write
Disk Backups
Snapshots
Synthetic Backup
Real Time Replication

Recovery Methods

Instant Recovery
Disk Restores
Tape Restores
Point-in-Time
Roll Back
Search & Retrieve

Enabling Technologies

- Tape & Automation
- Virtual Tape Library
- De-duplication
- Continuous Data Protection
- Snapshots

Disk-based Data Protection

- Data Protection is about **Data Availability**
  - **Backup is NOT a business requirement -- Availability is.**
  - Data protection as an extension of data availability - Restoration focused
- Disk-Assisted and Disk-based protection methods
  - Speed
    - Backup Windows
    - Recovery Time Objectives
  - Random Access
    - Data set, application object or sub-object recovery - quickly
  - ILM
    - More gradations in lifecycle (B&W -> shades of grey)

Application Consistency

When an application is running during the "copy" process, various techniques are available to ensure data consistency

Much like the "open files" issue when backing up a file system that is in use, applications (like databases, messaging systems, etc) allow for different approaches to capturing a holistic picture of the applications data during a copy process (such as a snapshot, a mirror-split, or CDP protection).

It is important to understand the consistency semantics of your application so that your data protection copies are recoverable.

Cold Protective Copy

- Old School approach:
  - Bring application down
  - Create copy (or mark timeline as a “cold moment”)
  - Bring application back online
- Requires downtime
  - How much is dependent on technology, implementation and environment
    - Can be from seconds to hours
  - This is a true “backup window”
    - Actual downtime during backup copy operation

To Quiesce or Not?

- Cold Snapshot
  - Less complex, but backup window is downtime
- Application Consistent Snapshot
  - Application intervention
  - Application dependent
  - “Hot Backup” or “Online Backup”
- Atomic or Crash Consistent Snapshot
  - Ability to take snapshot for entire dataset at exactly same moment
  - Can be done in multiple ways
  - Recovery domain same as high availability systems

Application Assisted Protective Copy

- A “hot” or “online” backup or copy - is done while application is online and active
- Application coordination is invoked to facilitate copy
  - Often done in stages across data set or using a transaction log for shorter durations
  - Performance of application often degrades during “hot” window
- Also known as a “backup window”
  - Application isn’t offline, but isn’t performing at optimal levels
  - Application performance often continues degrading as the window expands or transaction flow increases
- Warnings
  - Not all applications support this operative mode
  - Source of complexity -- high failure rates (test often and thoroughly)

Crash Consistent

- Crash Consistent or Atomic Protective Copy
- An atomic instant picture of the data
  - Application must be capable of crash recovery
  - Can it be clustered?
- No application interaction
  - Application isn't involved in the snapshot or CDP operation
  - No (snap/operation) related performance issues
  - Backup window is eliminated completely
  - Simplifies operation run-book and backups
- Warnings
  - Not all applications can do crash recovery
  - Not all disk protections can do atomic operations across a wide set of disks

Snapshots

A disk based "instant copy" that captures the original data at a specific point in time. Snapshots can be read-only or read-write.

"A fully usable copy of a defined collection of data that contains an image of the data as it appeared at the **point in time** at which the copy was initiated. A snapshot may be either a ***duplicate*** or a ***replicate*** of the data it represents.

www.snia.org/dictionary"

Snapshot of Networked Storage

- **Terminology:** Snapshot, Checkpoint, Point-in-Time, Stable Image = Any technology that presents a consistent point-in-time view of changing data. *Many implementations exist.*
- **Why?** Allows for complete backup or restore, with application downtime measured in minutes (or less)
- Most vendors: Image only = (entire Volume)
- Backup/Restore of individual files is possible
  - If conventional backup is done from snapshot
  - Or, if file-map is stored with Image backup

Full Copy Snapshots

- Keep multiple RAID-1 sets of data
- Cease RAID-1 operations to one set of the disks
- Freezes the contents of that set at a specific point in time
- Can be done by a variety of sources
  - Host based logical volume manager
  - Fabric based logical volume manager or virtualizer
  - Array based functions
- Frozen copy is not dependent on current copy for content
  - Better tolerance for failure
  - Be careful: If alternate set is dependent (cage, power, director), this “independence” is compromised

Differential Snapshots

- Keeps track of “changes” to the primary copy
- Uses a combination of the “change” set and the primary disks to save/present snapshot
- Different approaches optimize for different use cases
  - Copy On Write (CoW)
  - Redirect On Write (RoW)
  - Write Anywhere (WA)

Copy On Write

- Primary disks remain current
- Whenever a Write operation arrives, it is held:
  - First the current contents of the write-destination are read in
  - The old-contents from the primary disk is saved off somewhere and indexed
  - The new write is now allowed to pass through
- Read path of current disks remains optimized
- Write path of current disks is potentially impacted
- Read/Write path of “snapshot” disks impacted

Redirect on Write

- Primary disks are frozen
- New write operations to primary disks are stored in a journal (and indexed)
  - To read current copy, the journal is checked first
  - To read the snapshot copy, the primary disks are used
  - When snapshot is “dissolved”, write journal must be applied to primary disks to “catch up”
- Read path of snapshot is optimized
- Write path of current disks is optimized (no copy)
- Read path of current disks is potentially impacted

Write Anywhere

- All disk blocks are virtualized
  - Current disk is represented by a map to real blocks -- not directly mapped
  - Disk storage is larger than maps present
  - New writes, instead of “overwritting” blocks, are directed to free blocks
  - Maps are kept for “now” and potentially for multiple “snapshots”
  - Reference counts are kept for blocks “in-use”
- Performance doesn’t generally change primary/snapshot
- Performance can be impacted by fragmentation depending implimentation

Snapshot Comparison

	Full Copy Snapshot	Differential Copy Snapshot
Upsides	• No cost during "snapshot" process • Can be used for DR - independent copy	• Less storage consumption - typically 10-20% ◦ Depends on churn • Typically can take advantage of cheaper disk

Downsides	• Massive storage cost
	  ◦ 1x of storage per RPO
	  ◦ Like disk - expensive
	• Often in the same disk frame
	  ◦ Loss of DR component
	• Consider re-sync time in schedules
		• Performance impacts while snapshot exists
		  ◦ Multiple implementations to optimize performance impact
		  ◦ Most vendors don't offer multiple implementations - pick at onset
		• Leverages main copy - not DR capable

Applications	• Disaster Recovery
	• Near zero backup window
	  ◦ 24x7 operations
	• Faster restore
	  ◦ Can do no-copy restore
	  ◦ Most run-books require copy
	• Can help with data repurposing
		• Backup source
		• Near zero backup window
		  ◦ 24x7 operations
		• Fast restore
		  ◦ copy based by definition
		• Can help with data repurposing
		  ◦ Beware performance impact

Backup to Disk

- What and Why?
  - Part of the Backup Application itself
    - Easy to implement
    - Vendors may have low cost or no cost option
  - Can be directed to any type of disk or connection
    - FC, SATA, DAS, NAS, etc
  - Allows backup application to attain better tape utilization – streaming
    - Especially when backing up slower clients
  - Single/Volume Restores are faster than tape
- What to watch out for
  - Can be faster than tape, slower than VTL (generally)
  - Current versions of enterprise backup applications support it
    - May charge more for advanced functionality
  - Backup window issues may still exist - not “instant”

Introduction to VTL

What:

- Originally designed for mainframe
- Moved into open systems market
  - Fits within the backup environment
  - Easy to deploy and integrate
  - Takes advantage of current processes
  - Reduces tape media handling

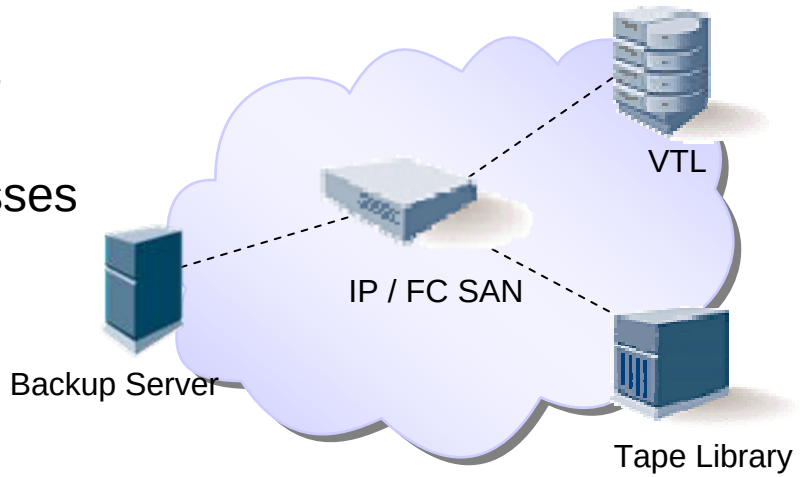

Why:

- Improved speed and reliability
  - Speed of higher end tape without the downside
  - Ability to achieve high performance with or without the use of multiplexing
  - Enables faster access to data for single file/folder restores
  - No detached tape leaders or mechanical failures

The VTL Difference

- ***Easy to manage in traditional backup software environment:***
  - Works like normal tape library
    - Fits into existing backup and restore processes
  - Viewed as open systems cartridges, robot, tape drives, and in some cases even a mail slot
  - Standard tape copy, cloning, or vaulting functions apply for off-site copies
    - Used to replicate data to physical tape for long term retention

- ***Cost effective solution***
  - Leverages lower cost disk
  - Can extend the life of current physical tape investment
    - Used as a front-end to the backup process
    - Tape is still used for longer term retention

© 2007 Storage Networking Industry Association. All Rights Reserved.

VTL Deployment

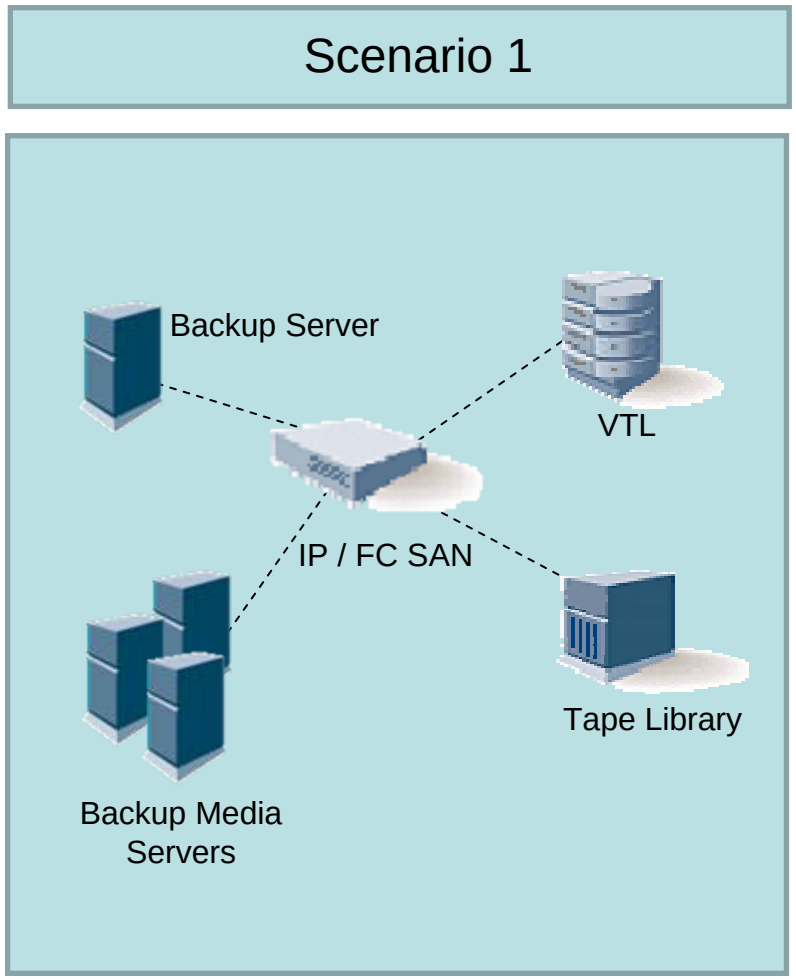

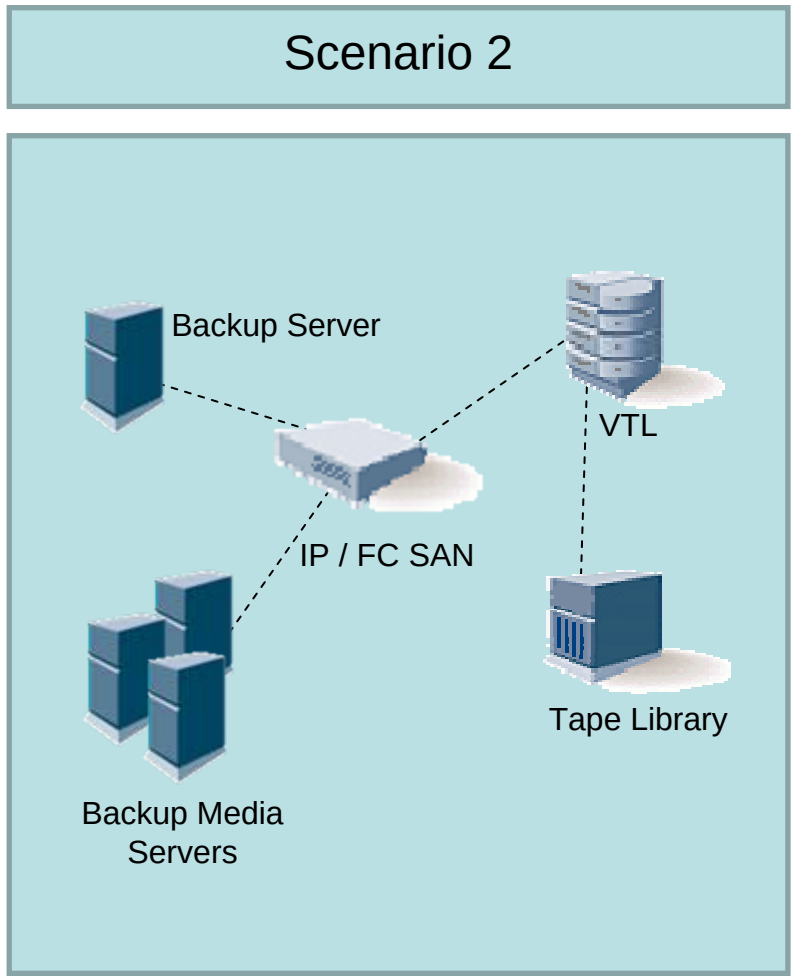

VTL Best Practices

- Use as primary backup target to reduce backup window
  - Still take advantage of tape as backup target where appropriate
- Add storage to enable additional restore time objectives
- Rules the same rules for libraries and tape when physically connecting
  - Virtual drives per connection
- Tape redeployment
  - Eject process, controlled by the backup software
- Offsite requirements
  - Bandwidth, connectivity, time to complete tape copies

© 2007 Storage Networking Industry Association. All Rights Reserved.

Introduction to CDP

What:

- Capture every change as it occurs
- Protected copy in a secondary locatior
- Recover to any point in time

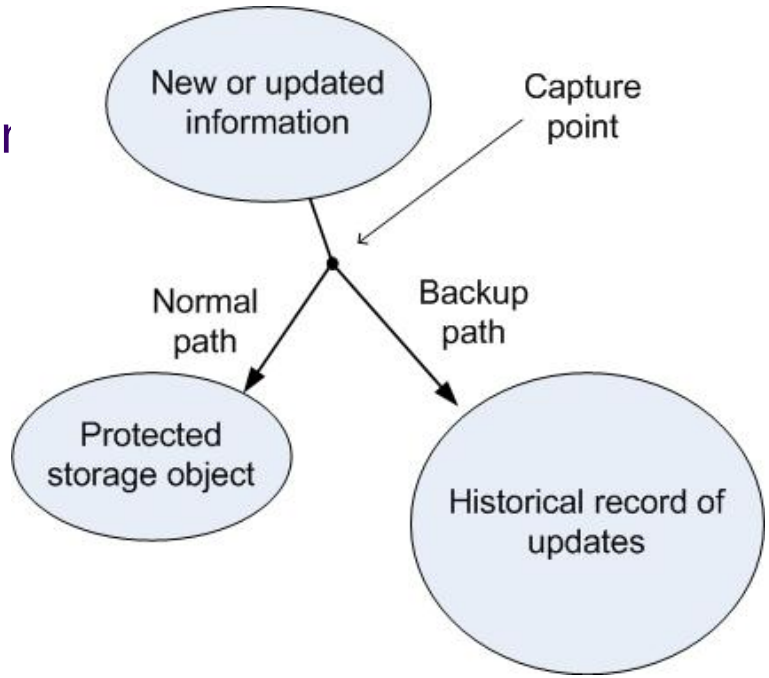

How:

- Block-based
- File-based
- Application-based

Why:

Implementations of true CDP today are delivering zero data loss, zero backup window and simple recovery. CDP customers can protect all data at all times and recover directly to any point in time.

© 2007 Storage Networking Industry Association. All Rights Reserved.

The CDP Difference

- ***Replication is not CDP:***
  - Maintains only a current copy of the data
  - May be combined with some snapshot capabilities
- ***Snapshots are not CDP:***
  - Snapshots are scheduled events
    - Data loss possible if crash or corruption happens between snaps
    - Snapshots frequently to same system as primary
    - Lack continuous index with embedded knowledge of relationship of data to files, folders, application and server
- ***Scheduled events are not CDP:***
  - Scheduled backup processes
  - Log collection for database style applications, rolling transactions forwards or backwards

© 2007 Storage Networking Industry Association. All Rights Reserved.

Traditional Recovery

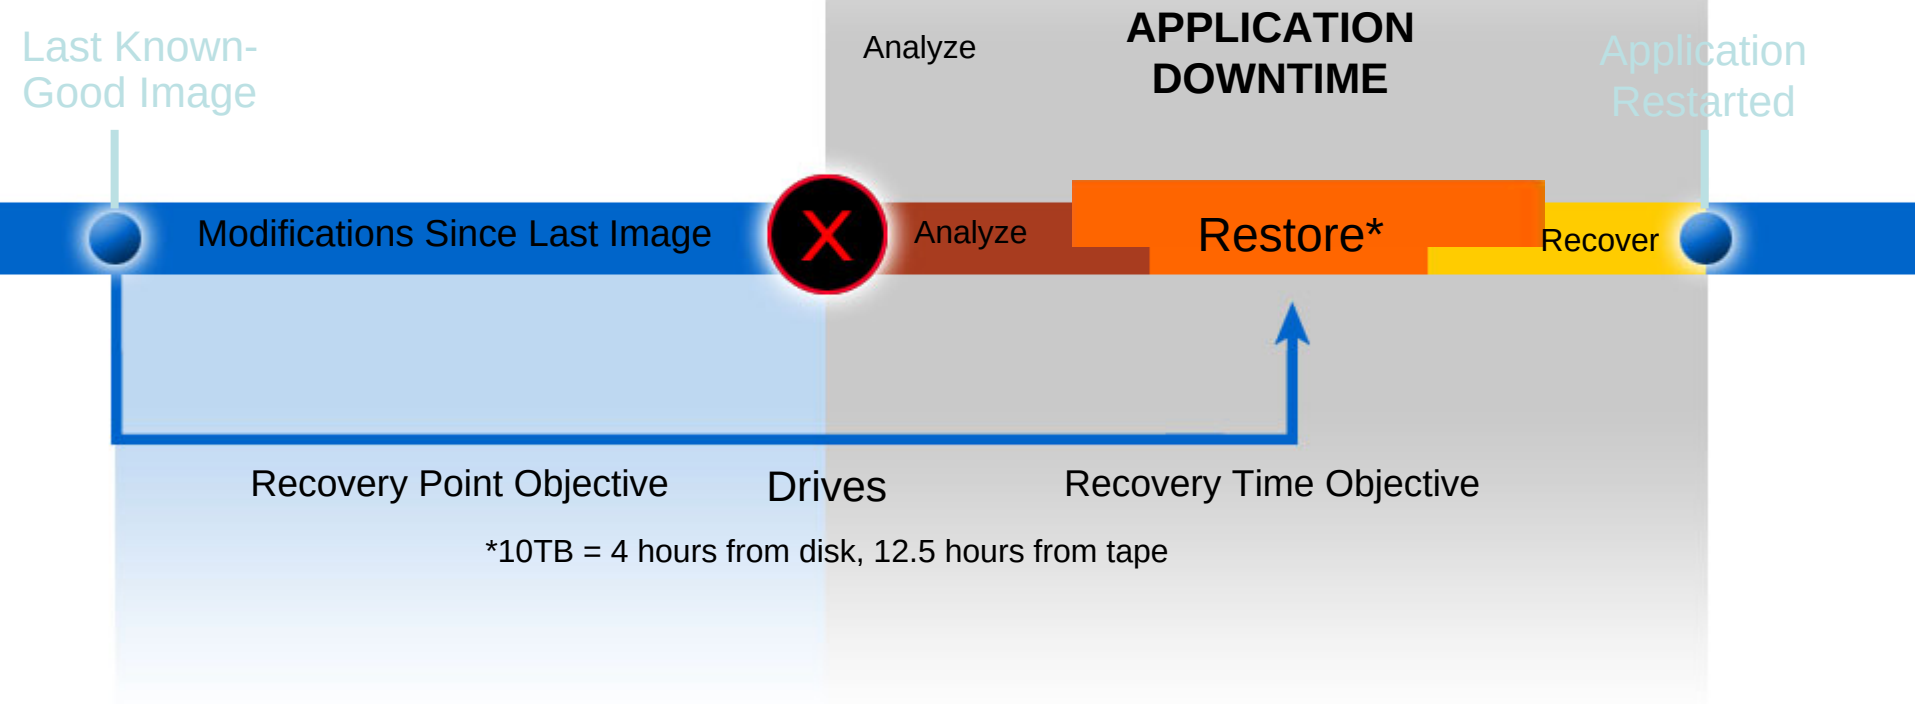

Recovery with CDP

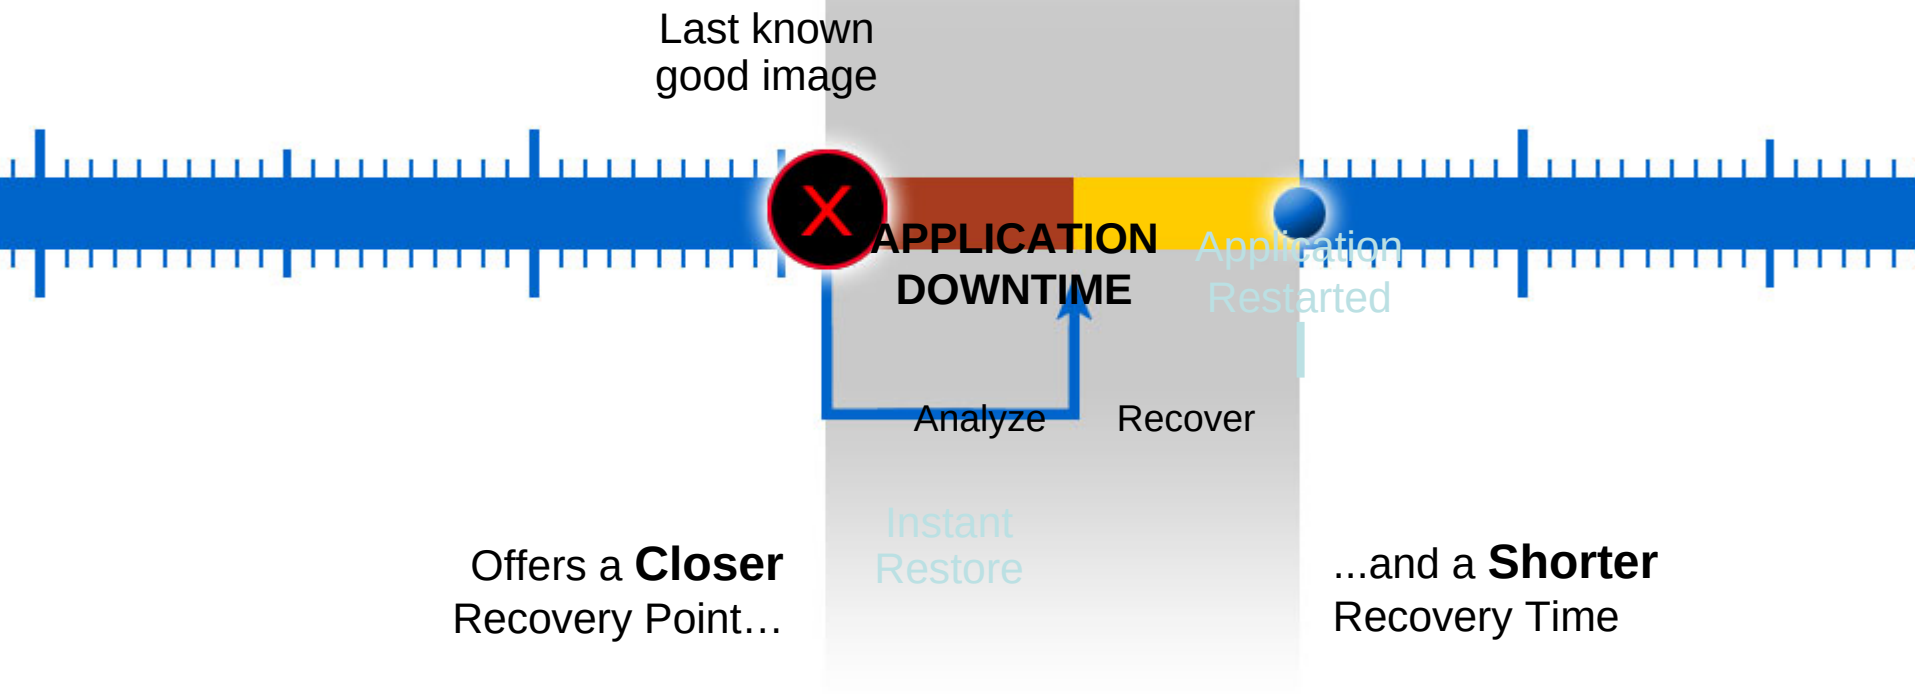

CDP Deployment

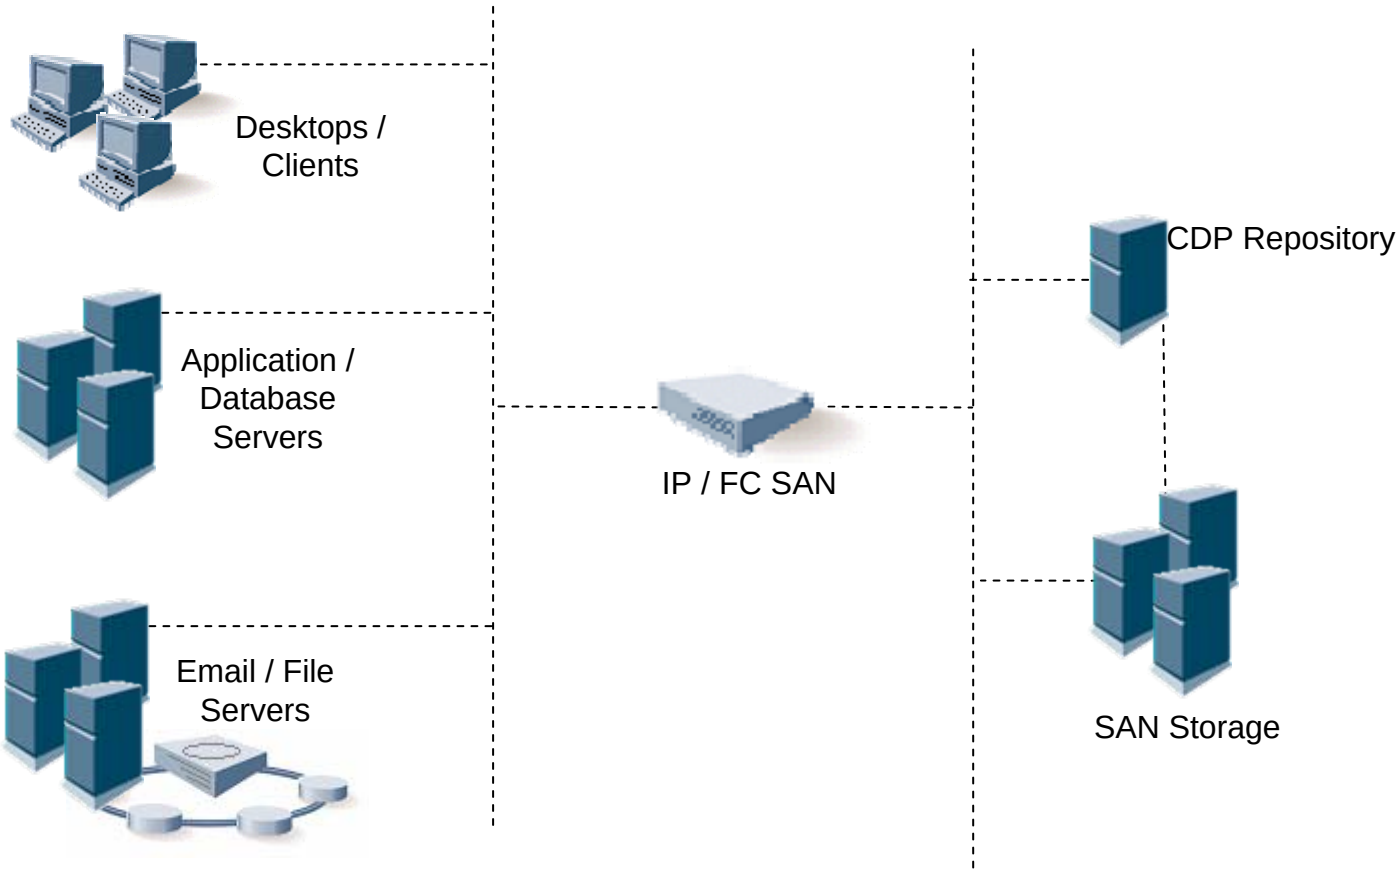

CDP Implementation Models

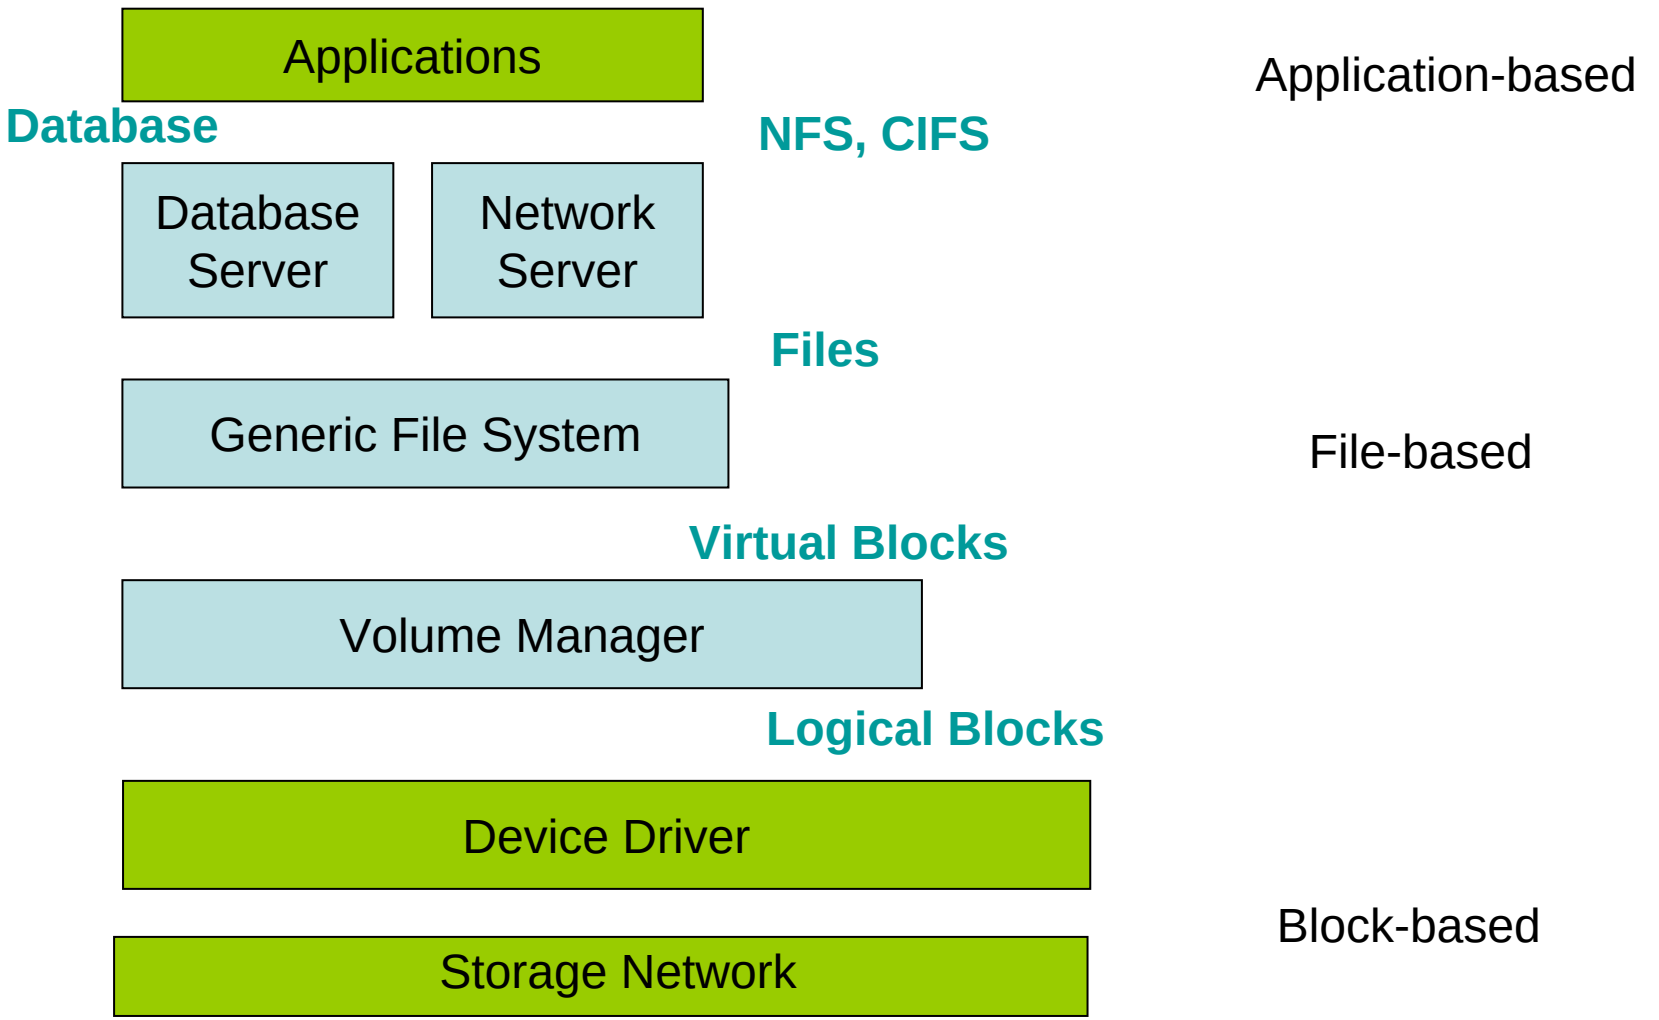

Consistency

- Logical Coherence
- In electronic data management, a set of data is said to be ‘consistent’ when the data can be correctly and unambiguously interpreted by an application.

Data Consistency is THE critical issue in data recovery

Remote Replication

- What and Why?
  - Extension of full snapshots, delta snapshots, DP images or CDP (depending on implementation/vendor)
  - File or block replication across a communications pipe synchronously or asynchronously (remote)
  - Combines DR with corruption protection
  - Extend tier 1 data protection to second and thirds data tier's
  - Can be used to consolidate tape creation to central location
- What to watch out for
  - Speed of light/ bandwidth issues
  - Site-to-site recovery can be awkward -- local corruption issues better dealt with locally

Solving the problem

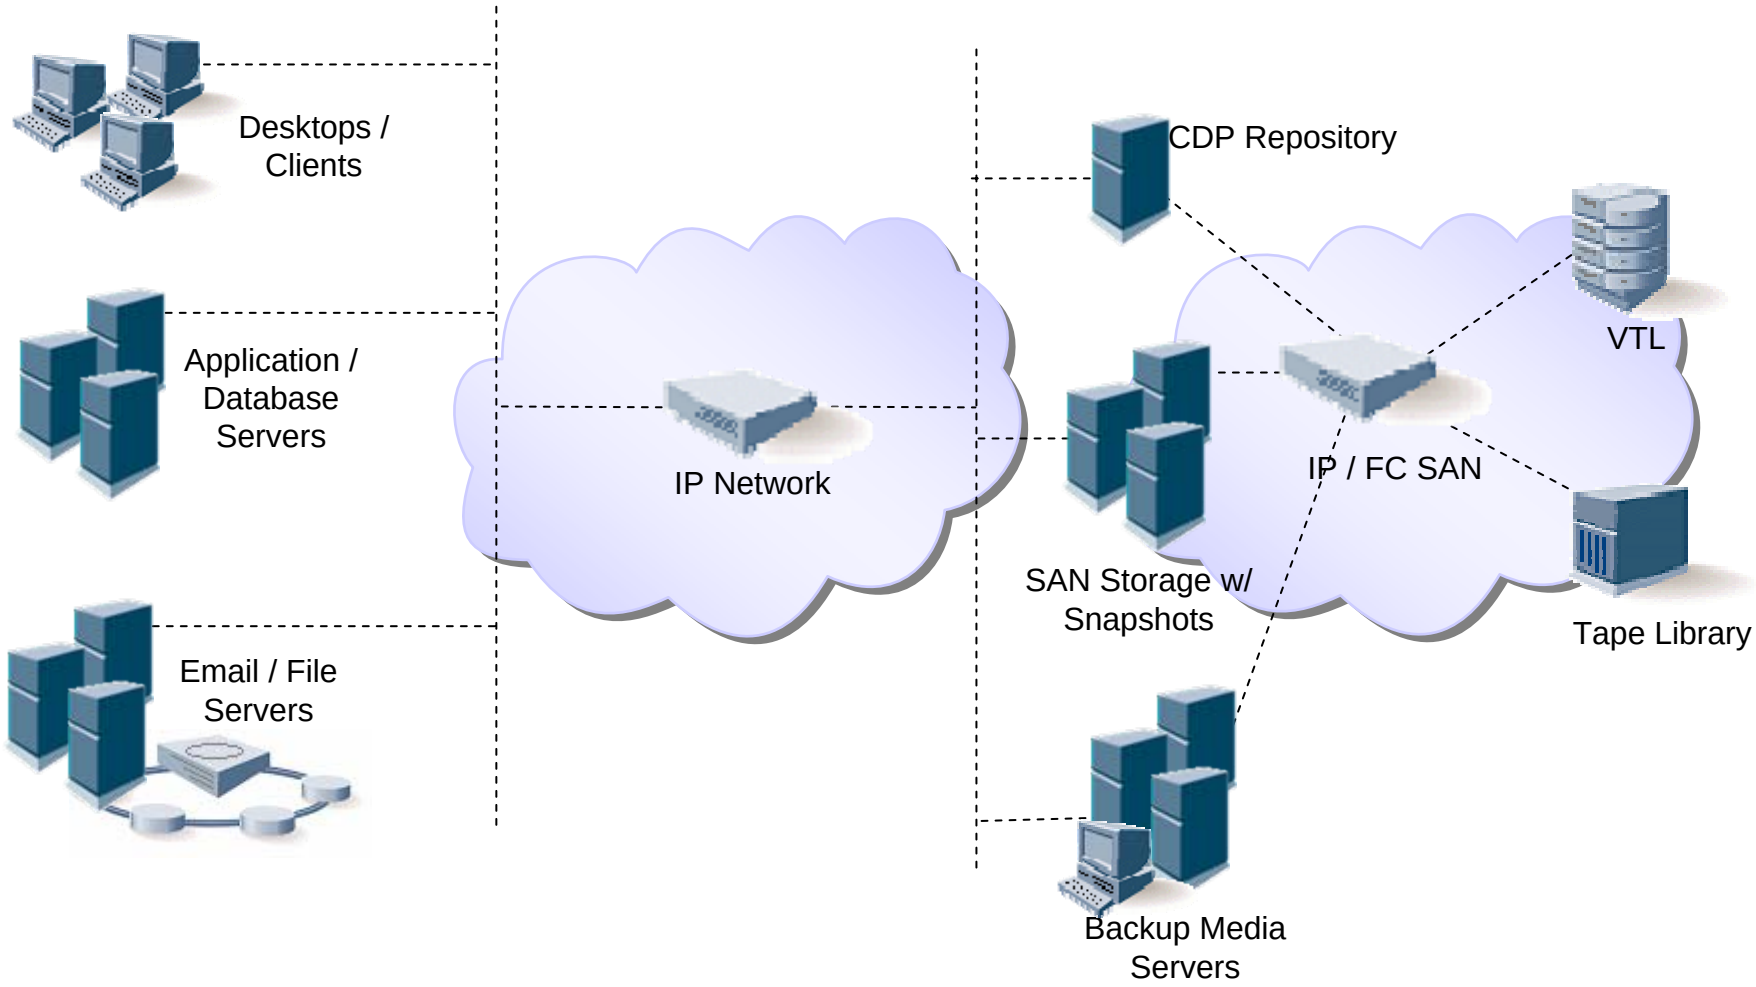

SNIA Resources

- Related tutorials
  - Disk and Tape Backup Mechanisms
  - Disk Based Restoration Technology
- Visit the Data Management Forum website at http://www.snia-dmf.org
- Data Protection Buyers Guide available

Q&A / Feedback

- Please send any questions or comments on this presentation to the SNIA: trackvirtualization@snia.org

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

SNIA Education Committee

Frank Bunn
Curt Kolovson
Ben Kuo
John Logan
Gene Nagle
Rob Peglar
Abbott Schindler
Nik Simpson
Wolfgang Singer
David Thiel

Q&A / Feedback

Please send any comments on this tutorial to SNIA: trackdatamgmt@snia.org

Many thanks to the following individuals for their contributions to this tutorial:

SNIA Data protection Initiative
SNIA "Data Management Forum"
SNIA Tech Council

Nancy Clay
Mike Rowan
SW Worth
Mike Fishman
Jason Iehl

Get Involved !

www.snia-dmf.org

- **Find a passion**
- **Join a committee**
- **Gain knowledge & influence**
- **Make a difference**

Thank you for your feedback

Education SNIA

Questions and Answers