



# **A CommVault White Paper: Singular Information Management™**

*Blueprint for Effective Recovery, Protection and  
Archive in a World of Exploding Data Volumes*

**CommVault Corporate Headquarters  
2 Crescent Place  
Oceanport, New Jersey 07757-0900 USA  
Telephone: 888.746.3849 or 732.870.4000**

# Singular Information Management™: Blueprint for Effective Recovery, Protection and Archive in a World of Exploding Data Volumes

## Table of Contents

Executive Summary: How the Newly-Available Benefits of Recovery Management are Reshaping Data Protection.....	1
<i>The Value of a Unified Approach to Data Management</i> .....	2
<i>Recommended Next Steps</i> .....	3
A Detailed Look at How Recovery Management Solves Challenges with Having Too Much Data: The Size of the Problem.....	4
Some Reasons Why Traditional Backup Software Hasn't Been Able to Cope with Big Data Sets.....	5
Gaining Ground on Data Growth by Managing Disk as Disk.....	6
<i>Direct Access Restore (DAR)</i> .....	6
<i>Single Pass Restore</i> .....	6
<i>Failover Resume and Alternate Data Paths</i> .....	6
<i>Data Mobility – D2D and D2D2T</i> .....	7
Gaining Ground with Active-Archiving.....	8
<i>Critical Capability: Data Mobility for Archives</i> .....	8
The Emergence of Recovery Management.....	9
<i>More About RPO's and RTO's</i> .....	10
<i>Recovery Using Snapshots</i> .....	10
<i>Recovery Using Continuous-Capture Replication</i> .....	11
<i>De-Mystifying CDP</i> .....	11
Singular Information Management™ in Action: Combining Recovery, Protection and Archive Capabilities.....	12
<i>Data Management without Unification</i> .....	12
<i>Singular Information Management™</i> .....	13
Conclusion: Unified Recovery, Protection and Archive Management.....	15
<i>Final Word</i> .....	16

The problem is still with us. Despite decades of technology evolution and investment, data management is still an unsolved source of expanding cost, anxiety and wasted time. With increased government regulations, business requirements and organizational scrutiny, the stakes are high – and traditional approaches to data management are breaking down.

In this blueprint, you will find an outline of how snapshot, replication and archive technologies combined with tiered storage systems are evolving data management. This is a step beyond backup-to-disk strategies, into improved RPO/RTO (*Recovery Point Objective/Recovery Time Objective*) service levels. This is about being able to use your data. This is a glimpse of what's possible when you add Recovery Management to Data Protection.

---

## Executive Summary: How the Newly-Available Benefits of Recovery Management are Reshaping Data Protection

We all understand the problem: too much data, too many applications and not enough time to back it up and restore it all. Recovery Management offers a new approach to solving the problem that has the potential to protect as much data as you need as often as you'd like, and to recover that data virtually instantaneously when you need it. No matter how much data you have.

Think for a moment what this might mean for your email applications and databases: data corruption and virus attack doesn't have to take your business operation down for hours and days. Think what this might mean for your business: continuous access to the data and information needed for better analysis, decision-making, action and competitiveness.

Recovery Management may sound like the same old story: that your backup is only as good as your recovery. But it's not the same old story: Recovery Management is about creating and managing online replicas of production data. When a replica is online it is immediately available; it does not have to go through a lengthy restore process before you can use it.

This is a dramatic change from backup – even backup-to-disk – in which the backup copy is not immediately usable, but must first go through a restore process. Depending on the size of the database, it's common for a restore process for Exchange or Oracle to take days of time. Days during which your end-users cannot use email or process orders. Days during which your business operation is slowed or stopped.

Technologies for creating online replicas – including snapshot and replication technologies – have been with us for some time. So, what's new? First, what's new is that these technologies are at last becoming more widely available. In the past, these capabilities were available on only the largest and most expensive storage devices. Now, storage manufacturers are introducing mid-tier storage devices with these capabilities. These manufacturers include HDS, NetApp, EMC, HP, Engenio, IBM, LeftHand, Xitech and many more. Some of these include API's to enable software management of their embedded snapshot and replication capabilities. Others support the Microsoft Volume Shadow Copy Service (VSS) standard, for the same purpose. Also, distributors such as Dell offer and make available a wide range of such devices including entry-level devices.

Equally important, software for managing these capabilities and the creation and use of online replicas is also becoming available. CommVault now makes it easy and efficient to create online replicas and put them into harness supporting your business. CommVault Quick Recovery integrates with storage-embedded snapshot capabilities to make it easy and simple to schedule the creation of multiple point-in-time snapshot copies of production data. CommVault QR Recovery Volumes built from these snapshots make it possible to reconnect applications rapidly following data corruption and virus attack.

QR integration with Microsoft VSS enables consistent management of all VSS hardware providers using a consistent, simple-to-use interface.

CommVault also makes it possible for its Galaxy Backup and Recovery customers to make backup copies from the QR-managed snapshot images, with its Recovery Director product. In this way, CommVault enhances traditional backup by virtually eliminating the backup window, particularly useful for very large sets of data. CommVault provides a one-step restore process back to the original client system protected in this way, to preserve ease-of-management and object-level recovery benefits.

### *The Value of a Unified Approach to Data Management*

What distinguishes CommVault's approach to addressing Recovery Management is not only the storage-agnostic capability to leverage virtually any storage device, but also the unified platform of common services upon which CommVault's capabilities are built and delivered.

CommVault provides a Common Technology Engine (CTE) with integrated capabilities across all three tiers of data management including Recovery, Data Protection and Archive. Integration provides cost-savings and ease-of-management advantages which are unique to CommVault.

This power of unification is apparent in the single GUI management console, single code base to maintain, single index, single systems infrastructure for data movement, single policy used to manage a data set across management tiers, and with a one-step recovery process regardless of which tier has the best available data copy. Device sharing across tiers is key to lowering total-cost-of-ownership investment levels. All of which delivers a uniquely beneficial solution set, which can be expanded incrementally according to your needs over time.

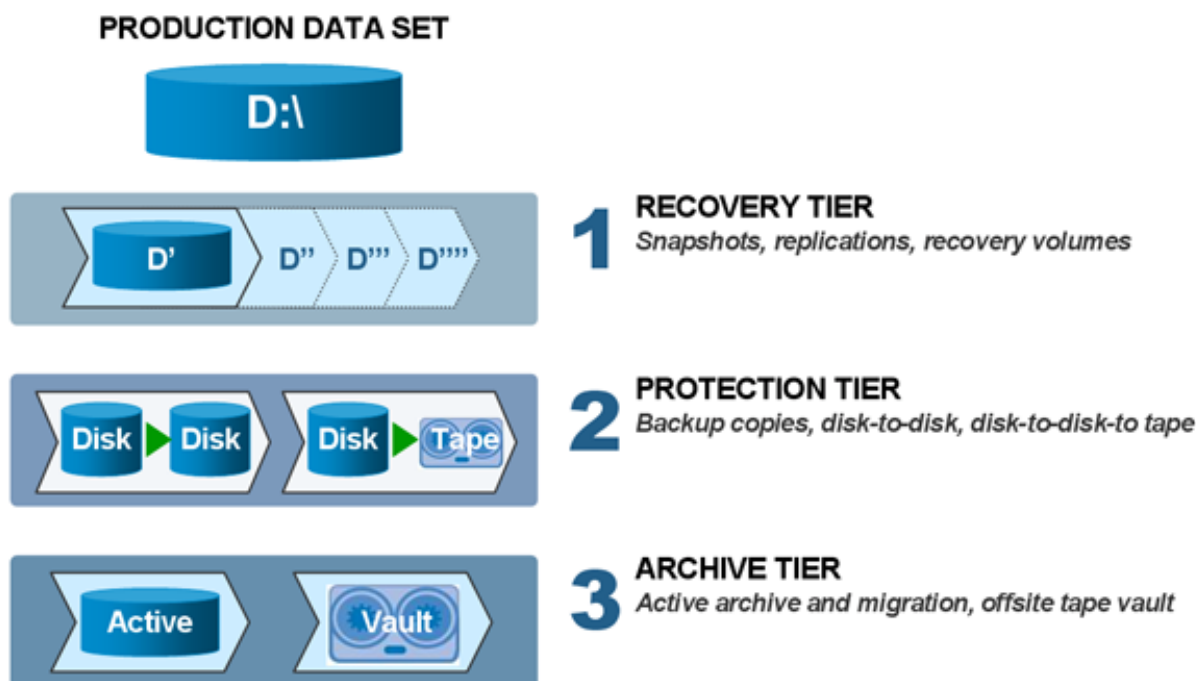


Figure 1. Singular Information Management Integrating Recovery, Protection and Archive for Ultimate TCO Benefits

Note that CommVault Singular Information Management™ provides these unique benefits as illustrated in Figure 1:

- One touch of the production data enables all subsequent data management and protection, to maximize applications performance and end-user productivity while also providing the highest recovery service levels
- One policy to automate how and when snapshots are created, backup copies are made, and auxiliary copies are made using disk-to-disk and disk-to-disk-to-tape device tiers
- Rapid mount of online Recovery Volume when the production data set is corrupted
- One-step restore of data, transparently from the best-available data tier to the production or Recovery Volume

#### *Recommended Next Steps*

Recovery Management is a new class of data management which offers dramatic benefits in data protection with improved RPO/RTO service levels. CommVault Quick Recovery, Recovery Volume and Recovery Director capabilities work in tandem with mid-range storage device capabilities to make these benefits more readily attainable.

Consider what Recovery Management can add to your data management solution, and work with CommVault to explore how to best apply Recovery Management for your business.

Work with CommVault to evaluate Data Protection and Archive capabilities, and build a cost-effective and comprehensive environment for unified data management according to your current needs.

## A Detailed Look at How Recovery Management Solves Challenges with Having Too Much Data: The Size of the Problem

The root cause that continues to cause problems with protecting data is the same as it has been almost since the beginning of computing: the continued explosive growth of data.

In a survey conducted by the Enterprise Strategy Group (ESG) analyst firm in 2005, respondents cited their top five problems with their current backup and recovery “solutions.” The winners were all about time: “backups take too long” and “recoveries take too long” were cited 66% and 49% respectively. In short, we still cannot get the majority of our data to fit into a reasonable “backup window”; and, when we need it most, we cannot get our data restored from the backup copy and into a usable version efficiently enough to suit most of us.

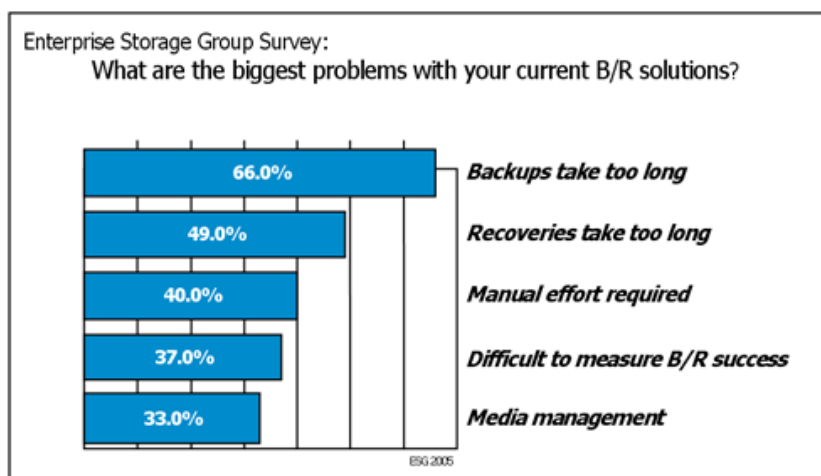


Figure 2. Survey Results Reported by the Enterprise Storage Group in 2005, Responding to the Question “What are the biggest problems with your current B/R solutions?”

Where is all of this data coming from? The InfoPro team published a study in the Spring of 2006, which cites databases and content repositories, files shares and e-mail systems. In short: no surprises. The data comes from every user in every business. It comes from our growing need to keep multiple copies of the same data: copies made by multiple backup jobs; copies in archives for compliance and discovery; copies of copies, in point of fact.

We also keep these data copies for far longer than we used to: it is now normal to keep most data for seven year retention periods, and at least some data for infinite retention. Why? Because the laws are still being written: we do not know which data we can safely discard, and which will cost us millions of dollars if we fail to produce it in a court of law.<sup>1</sup>

So, we keep the data. We index it, so that we can have a chance to find it when we need it – maybe six or seven years from now. We keep the indexes. We pay service companies to scan old media when we lose the indexes. We are amazingly inventive with disk and network technologies to try to get enormous amounts of data protected overnight – when hopefully no one needs to use it – or on weekends.

And, with some exceptions, very little of the backup software currently deployed in the world works very well to protect large amounts of data.

<sup>1</sup> As the cost of disk continues to drop, some speculate that we will reach the point that all data will be kept with no deletion. These futurists refer to the concept of the *personal pedabyte* as being the unit of measure that each individual will consume in a lifetime, and that its cost will be very low.

## Some Reasons Why Traditional Backup Software Hasn't Been Able to Cope with Big Data Sets

Most backup software being used in the world today was designed a long time ago – back in the days when tape or optical platters were the only options for storing backup data copies. Those were the days when an entire backup set could fit onto about 100 tapes, when tape rotation often meant putting tapes on top of the tape drives, and when tape tracking was done on a wall calendar or spreadsheet. Simpler times, indeed.

Interesting fact about tape: it is a sequential technology. As the cost of disk has come down – dramatically over the last 30 years – the advantages of using a random-access technology which is faster, more efficient, and more resilient than tape has transformed data backup. Backup-to-disk is now deployed as part of a comprehensive data protection environment in more than half of the world's enterprises, and growing.<sup>2</sup>

The theory is that disk efficiencies will help us to cope with the expanding data that we need to manage. Also, since we still need to take data off-site, tape remains an essential part of the picture; more on this later.

As the world has changed from sequential to random-access media, most backup software has not changed. Most software backing up data to disk still uses the disk as if it were a tape: data is written sequentially, indexed sequentially, and when it comes to restore time the backup software must read through all of the data on the disk – front-to-back – as if it were on a tape.

Think tape cassette versus CD: it's as if your CD player made you listen sequentially to every track, before you could hear your favorite song.

Why this inefficiency? Because, for the most part, broadly deployed backup software cannot change too much. If it changes too much, then it loses a little something called "backward compatibility." Backward compatibility is critical for backup software manufacturer's to maintain – because otherwise upgrading their customer's environments would be too difficult to accomplish. So, the world has been frozen – at or about the level of mid-1990's backup software technology.

---

<sup>2</sup> Enterprise Strategy Group (ESG) Report 2004

## Gaining Ground on Data Growth by Managing Disk as Disk

There are backup software products which build in essential disk-management capabilities. This section reviews these key capabilities which are essential to getting the most value from disk backup and recovery protection. These are also the capabilities which are essential if we are to begin to gain ground on managing data growth.

### *Direct Access Restore (DAR)*

Direct Access Restore (DAR) provides most efficient use of disk, by going to the location required for restore without having to read through all of the data on the disk sequentially – as if it were a tape. DAR is a key capability for enabling rapid access and restore of data from disk, to deliver better value from networked storage investments and to speed recovery times for large amounts of data.

### *Single Pass Restore*

A typical backup cycle is comprised of a “full” backup performed infrequently, typically once every week or two weeks, with “incremental” or “differential” backups performed to catch and protect changed data in between the “fulls”. This makes for a more efficient backup process.

However, most software that performs “fulls” and “incrementals” does a poor and inefficient job during restore: instead of restoring the most recent incrementals and then building the full image, most software begins with the full and reads the incrementals in reverse order. This means that obsolete incrementals are included in the restore process. This means that more data than necessary is restored – typically increasing the time required and the disk-space required by an order of magnitude as compared to an efficient restore process that recovers each piece of data only one time.

In contrast, a single-pass restore process reads the most recent incrementals and differentials first – and then skips past obsolete data copies. This process is only possible with a strong indexing model, which takes advantage of disk random-access capability.

### *Failover Resume and Alternate Data Paths*

When you make an investment in networked storage devices, you should expect to gain some return on that investment. That is reasonable – that’s how businesses operate.

Unfortunately, if you continue to use backup software designed in the mid-1990’s – before networked storage was invented – then you are likely not getting the type of return on your networked storage that you should.

Because it is on the network, part of the beauty of networked storage is that each device is accessible to many server systems. For data protection, this infrastructure can offer dramatic benefits – if the backup software is constructed appropriately to use the infrastructure. Look for two key capabilities:

- Failover resume
- Alternate data pathing

Failover resume is the capability to resume data management operations *from the point of the failure* when backup servers or networks fail. Note that this is far different from failover *restart* which merely restarts the job – from the beginning of the process. Resume is critical in particular for large backup or restore jobs, and in particular if they fail towards the end of the job.

Alternate data pathing is the capability of being able to automatically resume a data management job using an alternative network path to the storage device. This is particular useful when NICs fail on IP networks, and when fiber cards fail in SAN deployments.

Note that these capabilities are critically beneficial for tape protection as well as disk-based protection, and can considerably enhance disk-to-disk-to-tape data management systems.

*Data Mobility – D2D and D2D2T*

Moving data among storage tiers and from disk to tape is part of the challenge of managing stratified, fully mature tiered storage environments. You need the mobility to move between classes and brands of storage, how and when you need it moved, and without creating “dead-end” archives that cannot be moved without being entirely restored or retrieved.

Such data mobility is provided in mature data protection solutions, as disk-to-disk (D2D) and disk-to-disk-to-tape (D2D2T) capabilities. The type of copy capabilities to look for include:

- Inline copy, which enables to you make more than one backup copy *at the same time*
- Selective copy, which enables to you make a second copy onto a different class of storage *just before the original data copy is pruned from the original storage location*
- Cascading copy, which enables to you make a second copy from a copy

This type of flexibility in creating data copies is necessary, for you to get the value of having different classes of storage. So that you can retain some data on higher-costing storage – like disk – for rapid recovery, while moving other data onto lower-costing storage – like tape.

Flexibility in creating secondary data copies does not need to compromise restore capabilities: mature data protection solutions offer one-step recovery directly back to the original client system from any secondary backup copy. Better systems also keep track and find the most readily available copy of the data, without requiring administrators to manage the copies.

## Gaining Ground with Active-Archiving

Adding active-archiving to disk-based data protection provides more help in our quest to manage expanding data. Active-archiving does this by reducing the amount of data that needs to be managed – but without sacrificing end-user access to the data.

On a given file and print or application server, as little as only 15% is typically in active use. Yet, this data takes space on the disk – and forces new disk purchasing. It also lengthens backup time, and takes extra disk and tape storage for the backup copies.

CommVault Archive software can solve this, by moving the inactive data to secondary disk or tape storage. With this type of approach, only the data part of the file is actually moved; the file entry is left in the original file system location. This “stub”, as it is typically called, takes very little space as compared to the original file but enables end-users to continue to see and select the file as if the data were present. When an end-user selects and opens the file, the software retrieves and restores the data from the secondary storage location – which alleviates administrators from needing to assist in the recovery process.

Well-designed active-archiving software enables tape as well as disk to be used as a secondary storage location. The retrieval time may be a bit longer than with disk, but the lower-cost of storage can make the trade-off worthwhile for some types of data. In this case, the active-migration software must be able to notify end-users and system administrators in the event that tapes required to retrieve data are not present.

Active-archiving software most often migrates data based on last-use or age of the data. Less often, data classification capabilities can also be used to migrate based on additional factors – such as application, class of user, type of data, and so forth.

Integration with backup software is key for unlocking the value of active-archive software. A software which archives data but which is unsupported with backup and restore processes, or which causes a retrieve of archived data when the backup software “walks” the file system, does not offer full value.

In the ideal cases, archive software is integrated to work with backup and restore operations to offer the highest value. Such integration is offered by Data Archiver in the Singular Information Management™ suite of products. CommVault provides software which integrates to offer a single management GUI, a single policy for data management which includes active-migration and backup, and a single index for easy search and retrieval.

### *Critical Capability: Data Mobility for Archives*

Most archive software in the world today creates what could be termed “dead-end” archives of data: archives which cannot be copied, migrated, nor moved without having to be entirely retrieved from the archive as part of the process. This is an obvious problem and deficiency – and one which you can avoid by selecting archive software which offers data mobility with inline, secondary and selective copy options.

### The Emergence of Recovery Management

Recovery Management has emerged within the last 12 months as a separate and distinct category of data management. Recovery Management is about improving RPO (*Recovery Point Objective*) and RTO (*Recovery Time Objective*) service levels dramatically as compared against traditional data protection approaches. This is achieved by creating and managing online replicas of production data using snapshot and continuous-capture replication technologies. Because the data is online and immediately usable without needing to go through a restore process, RTO times can be virtually instantaneous. Continuous-capture gives virtually continuous RPO through most types of failures. Snapshots are required to ensure protection against virus attacks and other data corruption events, and provide RPO levels at 1 hour or better intervals depending on the type of snapshot being used.

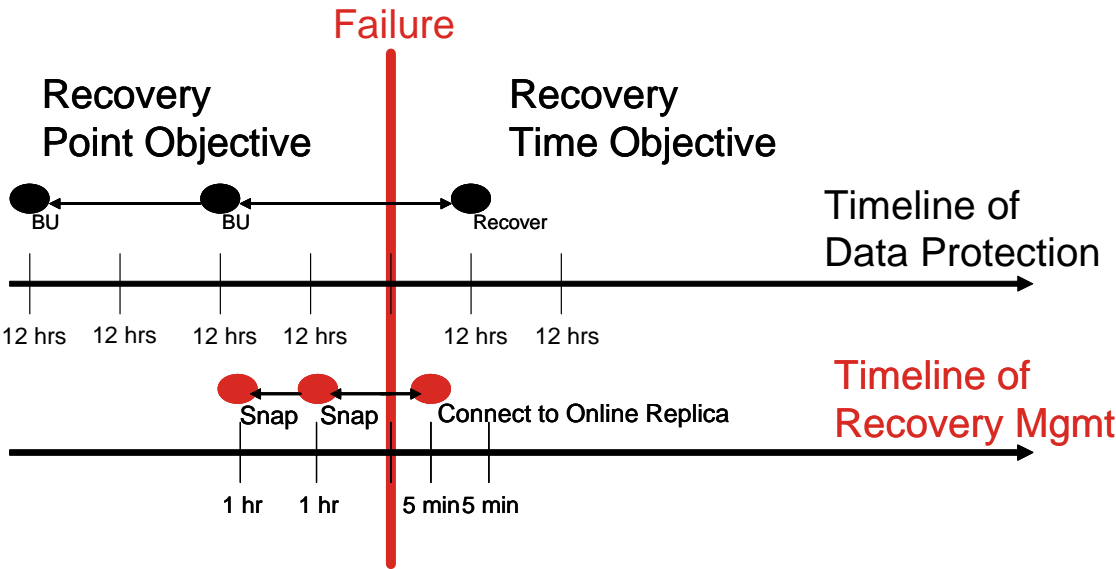


Figure 3. Dramatic Improvements in RPO and RTO Service Levels Possible with Recovery Management as Compared with Traditional Data Protection

The dramatic benefits of Recovery Management are clear. Snapshot and replication technology is not new.

What is new in the last 12 months? What has enabled Recovery Management to come into use – and pushed it along to some extent – is the wider availability of storage-embedded snapshot and replication technologies. These technologies were once available on only the most expensive, top-of-the-line storage devices. Now, manufacturers make these more widely available by offering mid-tier storage systems that embed these capabilities. CommVault enables these embedded technologies to be simple and easy to use, by integrating with them and offering automated policy-management.

In addition to offering dramatic improvement in RPO/RTO service levels, data replicas can be used to create backup copies. This means that instead of having to work with production data directly – which slows or interrupts application processing – data backups can be taken from the secondary copy of the data. With Recovery Management, the concept of trying to fit all data into a “backup window” is suddenly obsolete.

### *More About RPO's and RTO's*

A simple definition of the Recovery Point Objective (RPO) is that it represents the maximum amount of data loss that a business can tolerate. To understand this better, think about what happens with traditional data protection. A backup is performed typically at best, once a day – probably over night. If a failure occurs the next day, then any data written between when the backup was made and when the failure occurs is irretrievable – it's lost – because it has not been backed up. If the failure occurs towards the end of the day, as much as an entire day's worth of new data could be lost. In this case, therefore, the RPO is 24 hours' worth of data.

To improve the RPO, data must be protected more frequently. Most of the time, it's impractical to increase the backup frequency: 1) it impacts the application performance and cannot be performed during peak use hours – during the day; 2) the process of data backup takes too long.

To improve the RPO, a new approach is required – and that new approach is provided by Recovery Management. Using continuous-capture and snapshot technologies can both dramatically improve RPO's.

A simple definition of Recovery Time Objective (RTO) is that it represents the time that a company can tolerate needing to spend to recover data. With traditional data protection, a backup copy of data is not immediately usable – it must first be restored. Object-level restore capabilities can make this very efficient for individual data objects such as a data file or an email, or even collections of many files and mailboxes. However, restoring entire databases and very large amounts of data takes time. Commonly, this process can take days of time for average-sized systems.

The answer to this is to have online replicas of production data, created using mirror-image snapshots or continuous-capture replication or both. Because the replica is online, it is immediately available for use. It is possible to have very fast access to the data replica – with an RTO of five minutes or less for entire databases and other very large data sets.

### *Recovery Using Snapshots*

There are two types of snapshot capabilities with which to be familiar: copy-on-write (COW) and mirror snapshot technologies.

COW is more efficient than a mirror, because it does not make a copy of the data but instead creates a set of pointers to the data. Creating a set of pointers is a virtually instant process with virtually no system overhead – which makes it possible to create multiple point-in-time snapshot copies of the data at very close intervals of time without compromising system performance. When a new piece of data is written over a byte in the protected data, at that time the COW technology makes a copy of the original data byte to preserve it. COW capability is also efficient on disk because it only makes a copy of the changes in the data as they occur rather than an entire data copy – perhaps 80% of which is not required.

However, recovering from a COW snapshot takes longer than with a mirror image – because the recovery requires recreating the original data set by reading through the pointers and changed data bytes. Mirror-image snapshots recreate an entire copy of the data. This process takes longer than a COW snapshot, but is still dramatically faster than performing a traditional backup. Because the mirror-image of the data is online and immediately usable, it also dramatically shortens recovery times.

### *Recovery Using Continuous-Capture Replication*

Continuous-capture replication technologies are also good options for improving RPO's – and make them virtually continuous. As changes in data are written to the production data on disk, continuous-capture technology duplicates the I/O stream and then replicates it across a network connection to a second system – where that I/O stream is replayed against a second copy of the data. Replication technologies can be either synchronous or asynchronous, and keep a duplicate copy of the production data constantly available. As with a mirror-image snapshot, the duplicate copy of the production data is online – which makes it immediately available for use when failures occur to dramatically shorten recovery times.

However, replication technology has a vulnerability which is important to understand: by design with replication any changes in the production data are sent to the second copy of the data. This includes data corruption from virus attacks or other events. Because of this, it is beneficial to combine replication with snapshots of the online data replica – to ensure that the second copy of data can be recovered from a point in time before the virus attack. This also means that continuous-capture replication cannot deliver continuous RPO's through all types of failure events.

### *De-Mystifying CDP*

Continuous Data Protection, or CDP, is the concept of providing continuous RPO protection. Most offerings which are called CDP, combine continuous-capture replication with snapshot protection – and deliver continuous RPO's through all events except data corruption, and virtually instantaneous RTO's through every failure event. The RPO on data corruption is the last snapshot; which is very frequently made – usually within the hour of the failure event.

These types of CDP solutions – which combine replication with snapshot capabilities – are termed *near* CDP by most, and *kind-of* CDP by some. These technologies for the most part are extremely accessible and practical to deploy, and provide the majority of benefits required.

*True* CDP offerings solve the problem of being only as good as the most recent snapshot, to provide continuous RPO's even through data corruption events. However, these offerings are less widely available, are more difficult to deploy, require over-provisioning of hardware, and come at a premium. These technologies are also not as well-proven nor easy-to-manage as replication and snapshot capabilities.

For most types of systems and data, *near* CDP is the option which balances ease-of-management and cost against RPO and RTO service levels.

## Singular Information Management™ in Action: Combining Recovery, Protection and Archive Capabilities

Each of the capabilities described in this *blueprint* can be evaluated, selected and purchased as individual components offered by individual vendors. This can be a useful method for enhancing an existing data backup deployment, when you cannot afford to replace it. And, indeed, in most cases even when you select products offered by the same vendor, it is most often the case that the products do not integrate – were built by separate software companies, use different GUIs and indexes, and are offered from a single vendor now only by the virtue of having been acquired to fill in the gaps of a software portfolio.

CommVault offers a better approach, with its Simpna™ suite of software all of which is built from a common set of data management services. CommVault provides one GUI from which to manage recovery, protection and archive, with one index, and one policy for data movement across storage tiers.

To better illustrate the benefits of Singular Information Management™, this section compares and contrasts managing a data set with and without an integrated approach.

### *Data Management without Unification*

In this example, let's say that we have a set of Exchange data and a set of Oracle data, both of which require rapid recovery of under an hour in the event of data corruption as well as persistent protection with a seven year retention period.

Using a point-product approach requires all of the following software components:

- Snapshot capability embedded in SAN or network-attached storage
- Scripts to manage the snapshots – or, snapshot management software that works only with a single disk brand
- Backup software
- Library sharing software, for large-scale deployments
- Tape vaulting software

Typically, each of these software components operates independently – requiring separate GUIs, training, and maintenance. Moreover, the software often does not work well together.

For example, backup software may not be supported with the snapshot management software. This means that the backup copy would have to be made from the production data copy rather than from the snapshot copy, effectively disrupting the use of that data during the time of the backup and putting you right back where you started – with too much data to fit comfortably in your backup window.

Even if the backup can be taken from the snapshot, since there is no integration it cannot be easily automated. How do you know when the snapshot is available for backup? How can you schedule the backup when the snapshot might not be coherent? How can you recover back to the production data copy from the backup without disrupting snapshot protection?

Add active-migration, and you get an even bigger management challenge.

So, it all sounds good, but unless it has been designed to work together it can create more problems than it solves. It's exactly these practical realities that have prevented most of us from deploying more than the basics, even though our need has been great.

### *Singular Information Management™*

In contrast, the unified approach offers extreme efficiency and capability. Because the snapshot is integrated with backup, the data is always coherent – and scheduling the backup and the snapshot can be done in a single policy. Because the production data is touched only once – for snapshot – it's extremely efficient. Restore is also efficient, since CommVault's Proxy host capability makes the recovery process the same as if the backup was made from the production data: there is no difference in selecting the client, selecting the data to recover, and then executing the recovery process.

When active-migration of data is added, CommVault adds this as just another integrated capability.

Let's take the Exchange and Oracle examples:

- CommVault Quick Recovery enables the storage-embedded snapshot to take a snapshot of the Exchange and Oracle databases
  - This works across a wide range of storage options available from Dell, including EMC Clariion, Symetrix, and Dell mid-range storage systems
  - CommVault builds Recovery Volumes from the snapped data, which offer protection at multiple points-in-time – for example, once an hour with the most recent six being retained for rapid recovery
- CommVault Recovery Director enables easy use of the snapshots for backup – using CommVault Galaxy
  - Using a single scheduled policy for each set of data
  - Touching the production data only one time, when the snap is made
- CommVault Galaxy enables easy device sharing through its scalable Media Agent architecture, for connection to and management of a full range of tape libraries available from Dell – with no additional software required
- CommVault VaultTracker enables scheduled and automated tape rotation and tracking to offsite locations

CommVault offers all of this with a single GUI, a single index, a single policy-management engine, and total device sharing. If active-migration is added, using CommVault Data Archiver, then the disk and tape storage used for backup can be shared and managed using the same Media Agents – for further cost savings and efficiencies.

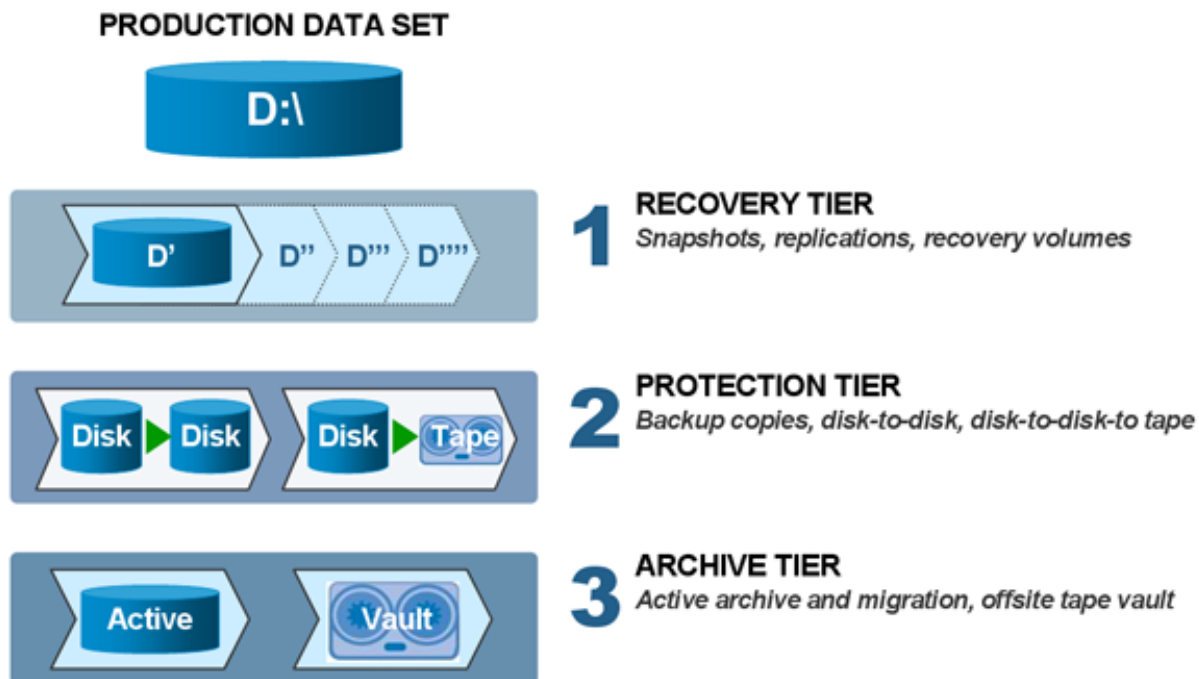


Figure 4. CommVault Singular Information Management™ Spans Data Tiers for Efficient Recovery, Protection and Archive

Note that CommVault Singular Information Management™ provides these unique benefits as illustrated in Figure 4:

- One touch of the production data enables all subsequent data management and protection, to maximize applications performance and end-user productivity while also providing the highest recovery service levels
- One policy to automate how and when snapshots are created, backup copies are made, and auxiliary copies are made using disk-to-disk and disk-to-disk-to-tape device tiers
- Rapid mount of online Recovery Volume when the production data set is corrupted
- One-step restore of data, transparently from the best-available data tier to the production or Recovery Volume

In this way, CommVault Singular Information Management™ delivers the power of one. Singular Information Management™ offers unique and compelling benefits for comprehensive data recovery, protection and archive.

## Conclusion: Unified Recovery, Protection and Archive Management

As we have discussed in this *blueprint*, we have a broader range of tiered storage running in our environments than ever before – and we're deploying more mid-tier storage this year. The mid-tier storage has more and better embedded capabilities for snapshot and replication – capabilities which previously have been available only in the largest storage devices.

Most of this storage is forming a new tier in our environments, called the Recovery Tier. This tier complements and enhances the Protection tier, which has been present in some for for a long while, and the Archive tier which has more recently emerged. Each of these tiers offers different capabilities and service levels for data management, to enable management of large amounts of data:

- Key Recovery Tier technologies
  - Snapshot
  - Continuous-capture replication
  - Data classification
- Key Protection Tier technologies
  - Disk-based backup, disk-to-disk and disk-to-disk-to-tape
  - Data mobility with inline, secondary and cascading data copies
  - Direct Access Restore (DAR)
  - Single Instance Storage (SIS)
  - Object-level restore
- Key Archive Tier technologies
  - Active-migration
  - Active-archive
  - Archive mobility with inline, secondary and cascading archive copies
  - Automated vault tracking of removable media

These capabilities can be used individually or in combination. Use of snapshot and replication technologies offers dramatic improvement in RPO/RTO service levels. Centralizing data with remote replication offers cost-savings advantages by alleviating the need to manage data protection in isolated locations with low-experience staff. Creating online data images also enables data backup to be performed on secondary copies of the production data – making it possible to perform backup when required with no impact on production application services. Adding active-migration of data also dramatically reduces the amount of production data that needs to be managed, saving time and resources.

A unified approach in which these capabilities are commonly managed, scheduled, indexed, and recovered is the most cost-effective and efficient approach. CommVault software uniquely offers a Singular Information Management™ suite that spans these tiers of storage and data management, with one GUI, one index, one policy-management engine, one restore process, and shared devices and infrastructure.

*Final Word*

CommVault invites you to consider what Recovery Management can add to your data management solution. CommVault experts are available to work with you to explore how to best apply Recovery Management to your current data management deployment, and working with the storage systems of your choice. Ask for a CommVault Recovery Management consultation today and get started on a better road to recovery.

CommVault® provides Singular Information Management™ solutions for high- performance data protection, universal availability and simplified management of data on complex storage networks. The CommVault® Simpana™ platform, based on CommVault's Common Technology Engine, integrates Galaxy backup and recovery, snapshot management and recovery, remote replication, active data migration and archiving, e-mail compliance, enterprise service level management and reporting and storage resource management software solutions. The Simpana™ unified approach allows customers to add/integrate Simpana™ components, at a fraction of the time, effort and money required by separate point products.