



comtec

infrastructure | datacentre | communications



# Backup Software Is Dead.

Long Live Data Management!

## BACKUP SOFTWARE IS BORN

Backup software as we know it first emerged in the 1970s, with the advent of disk drives as our primary media. At the time, spinning media was cost prohibitive to store anything but production data. Tape was the cost-effective alternative for storing backup copies.

The origins of backup software were humble enough, with functionality limited to reading data from a server, optimising it for tape and cataloguing its location. It wasn't long before the volume of data being backed up exceeded the capacity of a single tape, so tape libraries were introduced.

Backup software added library and media management functionality; including robotics management for tape selection and the re-use or retirement of media.

### THE GENIUS OF BACKUP SOFTWARE DESIGN

For something that is so prevalent within IT, backup software has an ingenious design. Software developers not only had to deal with technology issues, they also had to factor-in business operational and workflow constraints.

When designing backup software, a number of assumptions were made:

- Storage is an integral part of the Operating System and the only way to access data is to read it from the OS. Hence the need for backup agents.
- Business applications can handle downtime, overnight or at weekends, during periods of reduced demand
- Tape media allows for rapid streaming of data in and out, as long as access is sequential
- Disks are too expensive
- In case of disaster recovery, data needs to be restored quickly

### THERE IS A DISTURBANCE IN THE MATRIX

Deployment of backup software in IT infrastructure led to the creation of a whole new set of processes and behaviours.

Backup software designers had to reconcile downtime that a business could reasonably afford during backup (the Backup Window) against downtime resulting from the need to restore or recover lost data (the Restore Window).

Backup software traditionally assumes that no other activity takes place during the Backup Window. As businesses moved toward an always-on, 24/7 operational model, the Backup Window became smaller and smaller. At the same time, the volume of data required for backup was increasing, requiring a bigger window. The results were inevitable; if business operations weren't to be impacted, backups were incomplete, leaving business critical data vulnerable. The time needed to recover data depended largely on the time it took to load the appropriate backup tape and extract the data. Whilst a daily backup of all data ensured quicker recovery, performing a full back-up every night was both expensive and IT resource heavy.

To balance the competing needs of business uptime, Backup Window and Restore Window, IT departments created the WFDI schedule (weekly full, daily incremental). The assumption was that business activity was low during evenings and weekends so having backups running during these times would have the least impact.

The WFDI schedule also addressed the need for a reasonable recovery time. In a worst-case scenario, recovery would constitute the weekend tape plus any subsequent daily, incremental tapes. WFDI was an effective solution, until the very principles on which it was built changed. As it became more of a liability, it led to the development of a new approach to backup – Data De-Duplication.

The origins of backup software were humble enough, with functionality limited to reading data from a server, optimising it for tape and cataloguing its location.

## MEANWHILE, THE WORLD BECAME FLAT

During the nineties and noughties data became huge; literally. The need for greater storage capacity and reliability drove the development of external storage devices and server virtualisation and the increase in workforce mobility drove the need for networked storage.

As data grew, the cost of disk-based media came down and businesses could now choose from a number of options, varying in price, performance and capacity. De-duplication technology fixed the problems inherent in the WFDI schedule and eliminated redundant data on disks, significantly reducing the data footprint.

Disks were now competitively priced and supplanted tape as the media of choice. They were easier to handle and provided quicker access. The world was now open for business 24/7 and there was no such thing as downtime.

### BACKUP FOR THE FLAT, VIRTUALISED WORLD OF CLOUD COMPUTING

The IT world is a dynamic one; and it is very good at adapting quickly to changing business requirements and customer demand. As the business environment evolved, technology advanced a new set of assumptions:

- Business needs IT support 24/7
- Server virtualisation delivers unparalleled agility, mobility and simplicity
- Data protection involves moving data from disk to disk
- Disk storage has become commoditised
- IT architecture is evolving to become hosted, hybrid, SaaS
- Traditional backup software is incapable of providing effective, efficient data protection

These guidelines inform the new-look data protection process. One with a number of definable characteristics:

- Data can now be captured instantly and stored on backup disk media without the need for reformatting – eliminating the Backup Window.
- Data recovery is now a simple process of mounting and binding data to the appropriate application. There is no longer a need to restore data to the production storage system – eliminating the Restore Window.
- Data can be maintained in multiple formats, including native application format for instant recovery or optimised for long-term storage.
- Network utility can be optimised with purpose-built protocols designed for the efficient movement of data to and from remote sites for disaster recovery or business continuity.
- Data can be moved seamlessly across sites; transparent to the backup process and across vendor-neutral storage media, either physical or cloud-based.

The assumptions on which backup software design was based have changed. So why are we still using backup software?

## GOODBYE BACKUP SOFTWARE, HELLO DATA MANAGEMENT

With both applications and backup utilising a common storage media, business data can be divided into two simple types. Production data used by business applications in day-to day operations, and copy data, the backed-up versions of production data.

The unification of media to store copy data drives one of the key paradigm shifts in data management: a single, unified approach to managing all copy data with the smallest storage footprint – VIRTUAL DATA MANAGEMENT

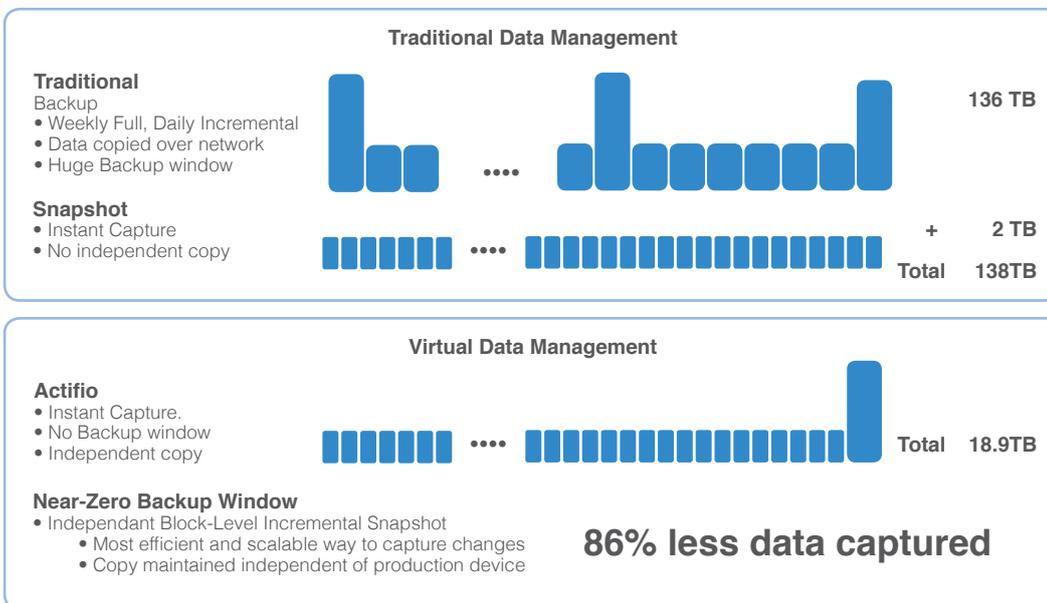
### VIRTUALISATION COMES TO DATA MANAGEMENT

Actifio set about designing a data management solution for disk-based media and pioneered Virtual Data Management (VDM). This technology virtualises the management of independent, duplicate copies of data into a single platform; eliminating redundancy and re-purposing unique data for business applications.

The result is an application-centric, SLA-driven solution that efficiently manages all copy data for multiple business requirements: data protection, disaster recovery, business continuity, compliance etc. working across physical, virtual or cloud infrastructures: AnyIT™.

At the heart of the Actifio technology is an SLA that drives the lifecycle of business application data. AppSnapCopy™ captures only the changed blocks of application data, resulting in a near-zero Recovery Point Objective (RPO). The pipeline architecture of this technology results in up to 80% reduction in the amount of data being copied.

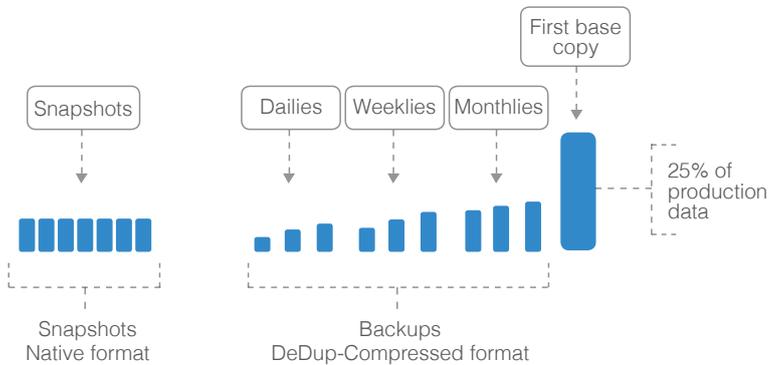
The figure below demonstrates the impact on a 10TB data set with 2% daily change rate.



AppSnapCopy™ captures only the changed blocks of application data, resulting in a near-zero Recovery Point Objective (RPO).

Virtual Data Pipeline File System (VDPFS), an Object file system, stores the delta-copy objects in raw, de-duplicated and compressed formats locally and/or remotely. The net result is a reduction of up to 95% in the storage footprint, which comprises only the unique blocks of copy data.

Unlike the current de-duplication devices, VDPFS implements a unique, context-sensitive de-duplication technology with self-describing objects. These objects comprise application data with temporal, casual and SLA relationship associations, enabling intelligent data management and re-use. Since data is captured at block level, the resource required to de-duplicate on fixed block boundaries are 6x less than the variable length process used by current de-duplication devices.



**DeDuplication + Compression format**

- Base copy reduced by 50% dedup + 50% Compression
- Only unique data kept across all secondary data

**Multiple data formats stored**

- Native for speed and dedup for reduction
- Define by SLA

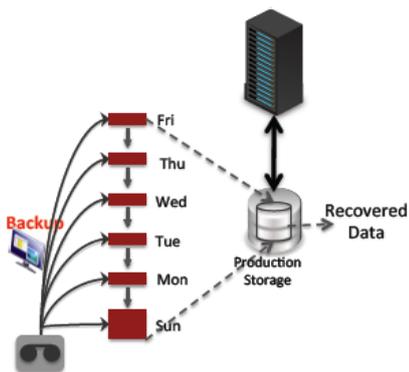
By virtualising vendor-agnostic physical or cloud storage resources into private, public or hybrid cloud storage pools, users are able to leverage commodity storage, or existing storage devices, to reduce storage costs by over 75%.

Data movement is performed by DedupAsync™. This purpose-designed network protocol results in efficient data movement through the optimisation of bandwidth usage by only transferring unique data to the remote storage site. In addition to the global de-duplication of data, DedupAsync™ drives down network usage by as much as 70%.

AppSnapRestore™ parses the VDPFS object file system to recover data in application format, at the point in time of recovery. By directly mounting the recovered data to a server, no data movement is necessary; resulting in a new zero Recovery Time Objective (RTO).

By virtualising resources into private, public or hybrid cloud storage, users are able to reduce costs by over 75%.

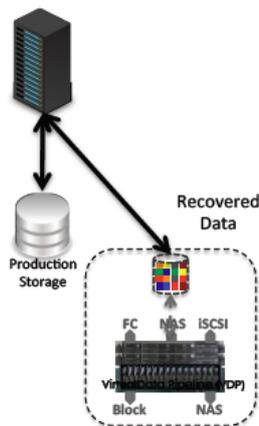
**Traditional Data Management**



**Restore Window**

- Sequential copy of incrementals
  - Overlay Full with incrementals
  - Copy data to production storage

**Virtual Data Management**



**Instant Restore**

- App-Centric Snaprestore
  - Instant restore of any data from any time
  - Restore with "mount". No data movement
  - Restore files, email, mailboxes or volumes

The result is not only a significantly reduced data footprint for recovery but also instant access by applications to data from any point in time, without the need to use production storage, application server or network resources.

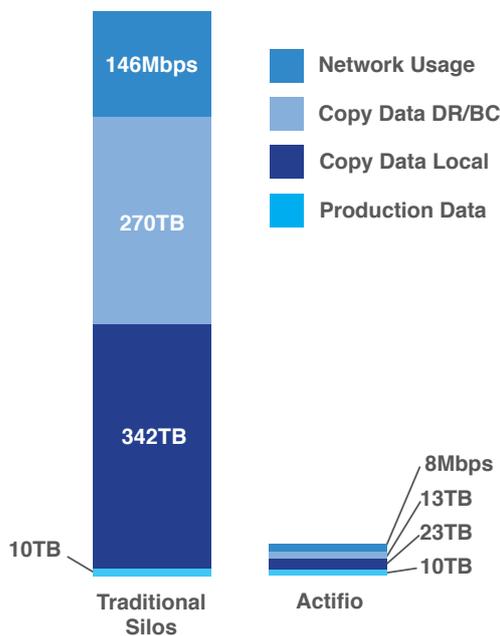
Deployment of the Actifio solution is zero-impact and requires no change to any existing data management applications such as backup, snapshot or replication. Actifio can be deployed as a single solution for all data management requirements or as a solution for a specific management application.

## 10X BUSINESS IMPACT

Virtual Data Management brings with it the proven benefits of virtualisation: simplicity, service-oriented management, flexibility and enhanced utilisation. Actifio users have been able to drive down Total Cost of Ownership (TCO) by up to 90% whilst simplifying multiple data management operations.

Virtual Data Management reduces costs by:

- Reducing copy data footprint
- Reducing the unit cost of copy data storage
- Reducing network bandwidth for replication
- Eliminating multiple software licenses
- Freeing production data storage capacity
- Eliminating the overhead typically imposed by backup software on application server CPUs, network and production data storage devices



### 10X Financial Impact

- Delay/reduce storage/tape cost by up to 95%
- Reduce bandwidth cost by up to 75%
- Reduce HW/SW licensing cost up to 65%
- Breakthrough the 30% VMware barrier
- Enable commercial/viable use of cloud

### 10X Operational Impact

- Single solution for Backup, Snapshot, dedup, DR, BC, test-dev, analytics and compliance
- Co-Exists with current infrastructure
- Manage physical/virtual environments
- Application driven SLA management

Actifio users have been able to drive down Total Cost of Ownership (TCO) by up to 90%.

## ACTIFIO SOLUTIONS

### Eliminate Backup and Restore Windows

With the rapid growth of production data volumes, backup software tools are unable to protect critical business data or recover data quickly. Actifio AppSnapCopy™ and AppSnapRestore™ are designed to provide instant protection and recovery of large quantities of data; eliminating the windows of vulnerability. Actifio also eliminates backup agents, freeing business application servers from backup and restore operations. Integrated data de-duplication, compression and encrypted replication reduce the storage footprint and network utilisation and eliminate software licence costs for point tools.

### Optimise Replication for Disaster Recovery and Business Continuity

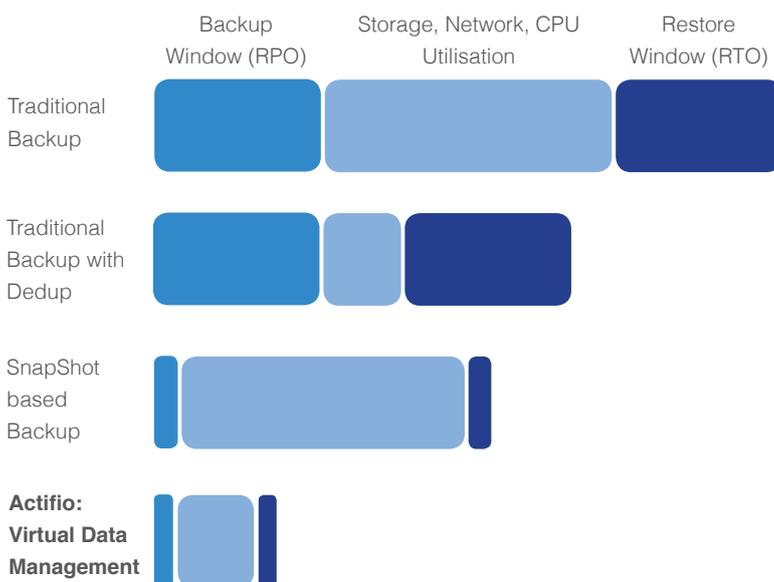
DeDupAsync™ eliminates the movement of duplicate data, reducing network bandwidth utilisation by 70%. In addition, the integrated storage virtualisation technology enables heterogeneous data replication, further reducing the unit cost of storage by 75%.

### Instant Copies of Data for Test, Development or Analytics Applications

AppSnapRestore™ allows copies of data to be instantly available, for any application, with no additional storage requirements. Clones of data can be made for I/O intensive applications. AppSnapCopy™ allows users to efficiently capture application-consistent data, multiple times per day, with no impact on host or storage systems.

## A PARADIGM SHIFT IN DATA MANAGEMENT

Actifio not only delivers a unified solution to address multiple copy data management applications, but also revolutionises the core backup function. The new architecture eliminates the need for backup software and all the complexity, inflexibility and inherent limitations associated with it.



DeDupAsync™  
eliminates the  
movement of  
duplicate data,  
reducing network  
bandwidth utilisation  
by 70%.