

WHITE PAPER



Beyond Deduplication Rethinking Data Protection

Syncsort, a global leader in data integration
acceleration and data protection software solutions

syncsort RETHINK THE ECONOMICS OF DATA ▶

Beyond Deduplication

Rethinking Data Protection

I. The State of Back-up	3
A. Enter Deduplication	3
B. The Challenge of Data Lift	3
II. Rethinking Deduplication	4
A. Data Reduction vs Data Deduplication	5
B. Back-up of Virtual Machines	6
C. The NetApp FAS Hardware Target Device	6
III. Getting Data Back: Recovery without Compromise	7
A. Blocks are Better	7
B. The Need for Multi-Tiered Recovery	8
C. NSB: Complete, Multi-Tiered Recovery	10
IV. Conclusion	11

I. The State of Back-up

In August of 2010, Gartner Inc released a significant research paper that referred to the “Broken state of back-up.”¹ Through thousands of end user interactions, the Gartner analysts concluded that enterprises both large and small were facing enormous challenges in protecting data.

How did we get here? Haven’t we been protecting data for decades?

We have, and that is a large part of the problem. Most users are still deploying technology that was crafted decades ago when a large server hard drive was 100 MB (to put it in context, that’s 95% less storage than an iPod Shuffle!). The core technology remains the same: copy files from a source to a target. It’s never been a very efficient model, but we’ve managed to get by because our IT infrastructures have increased greatly in capacity. Servers have vastly more computing power. Disk drives and tape drives are much faster. Network capacity has increased by orders of magnitude – the 10 megabit LAN of the past is now a gigabit LAN, with a hundred-fold increase in carrying capacity.

So why aren’t things working better? The core of the problem is data growth. Even while we were upgrading our infrastructure, the files kept coming. Structured data grew as more and more business processes were computerised. End user productivity applications like word processing and spreadsheets reached every desktop and unstructured data files grew like weeds. E-mail replaced phone calls as the primary means of communicating. And the files kept coming. While tape drives got faster, their sequential nature limited the amount of data we could pump at them. And the files kept coming.

We switched to disk back-ups, faster and more reliable. It helped, for a while, **but the files kept coming.**

A. Enter Deduplication

Disk back-up seemed to solve a lot of problems at first, but data quantities made it cost prohibitive. Even though hard drive prices trended downward, disk was still far more costly per gigabyte than tape. The nature of file back-up exacerbated the problem. Put bluntly, file back-up isn’t smart. It copies the same content over and over

again and stores great quantities of redundant data.

Why not just get rid of the duplicate data? A simple and rather ingenious idea, and it worked exceedingly well once disk-based deduplication technology became commercially viable. Data storage needs collapsed on the order of 90% or more. What used to take 100 TBs to store now took only 10 or 15 TBs.

Deduplication let users reduce or even eliminate the use of tape. Because you could now store weeks and even months of back-ups on disk economically, tape became expendable for anyone without long term retention needs. Replication transferred data off-site to a second device, eliminating the need for mobile media and truck transport. Savings were realised all around.

So with deduplication technology such a success, why is back-up still broken? **Because what deduplication did was fix a problem without fixing the process.**

And the files kept coming.

B. The Challenge of Data Lift

Disk-based deduplication fixed a problem that back-up created, but it was like applying sunscreen at the end of the beach party. Many issues remained unresolved and users found they still had a lot of pains to deal with.

Think about deduplication in a physical context. Imagine you are given a pile of one thousand bricks in a mix of ten different colours. Your task is to move bricks across a room such that at the far side you end up with only one brick of each colour.

You have two options. You can move all one thousand bricks across the room and then sort out the colours. Or you can sort out the colours first, and move only ten unique bricks. Clearly, the second approach makes far more sense and will be both easier and faster. But deduplication takes the first approach. It moves the thousand bricks and then sorts them out. Not so smart after all.

Of course, we’re not talking about bricks, we’re talking about data blocks. Millions and millions of data blocks. And we’re asking our servers to move those blocks, over and over again, even while they are running our applications and enabling our users.

This is the problem that we can summarise as Data Lift: the entire, end-to-end impact of the back-up process

¹See Best Practices for Addressing the Broken State of Back-up, Publication Date: 27 August 2010, ID Number: G00206028.

created by the need to lift and move data from source servers to a target device. Data Lift impacts everything in the back-up path: CPU, memory, disk and network I/O.

Deduplication addresses only the last step on the path, the final destination of the data. The problems that deduplication leaves unsolved are numerous and critical.

- **Back-ups take too long.** While disk-based dedupe targets provide marginal speed increases over tape, they have not solved the number one issue of data protection, which is meeting the back-up window in the face of relentless data growth. Often, target speed increases do not help because servers can't send data as fast as the target can receive it.
- **Restores take too long.** Deduplication doesn't transform the old model. We still restore by streaming files back to a target. This is far too slow for today's recovery needs.
- **Client impact.** Back-ups create too much impact on our servers and applications. That's why we have the notion of the "back-up window," a period of time when we can get away with back-up impact. But that window keeps getting smaller as applications are ever more ubiquitous. Really, when does e-mail ever go to bed?
- **Virtual Machine impact.** As 2010 drew to a close, six out of ten workloads were virtualised, and that number continues to rise. Hypervisor platforms share physical resources among multiple virtual machines — ten VMs per server on average — leaving much less capacity for back-up processing. Want to try backing up ten applications at the same time? Good luck.
- **Network impact.** Back-ups are sending too much data across the network. A target-side solution doesn't address this.
- **Media Server impact.** Users need to continually upgrade or deploy additional media servers as data growth overwhelms current resources.
- **Too much data loss.** Deduplication doesn't change the fact that we back-up once a day, leaving us exposed to as much as 24 hours of data loss. How much data can you afford to lose?

And the files keep coming.

II. Rethinking Deduplication

In the final analysis, deduplication is necessary *but not sufficient*. We need to rethink deduplication to understand that we must move beyond deduplication alone to solve the problems of Data Lift, which are the root cause of the broken state of back-up. As long as we're thinking things through, let's envision an ideal data protection scenario and then see how close we can get to it.

To start, we know there's too much data. We also know that it is physically impossible to protect data without making copies of it. This would seem a Gordian knot we can't unravel, unless we seize upon the very composition of computer data. The traditional unit of back-up – the file – is by no means the smallest unit available to us. Each file is composed of many blocks, typically in increments of 4 or 8 KB. (For the sake of our discussion, we will use a typical Windows 4 KB block size as our standard.) So a 2 MB PowerPoint file consists of 512 individual 4 KB blocks. A user who opens the file to make modifications rarely changes the entire 2 MB of content. Instead, some amount of data is added, removed or both. Then the file is saved again. Let's assume that about 10% of the file was actually modified, or 50 blocks worth (200 KB total).

During a file back-up, the back-up application would copy the entire file to the target media, even though 90% of the contents are the same as the last time it was backed up. The process of lifting the data would impact every significant hardware resource: CPU, memory, disk and network. Yet 90% of the effort is, essentially, wasted: it is doing work that's already been done. Even if another process deletes the excess data on the target side, you've still suffered the impact across the entire scope of your back-up process. Now multiply this by thousands or millions of files, across dozens or hundreds of servers, and you realise how much needless Data Lift burden your infrastructure carries.

Wouldn't it be better to only move the changed blocks? Yes – better, faster, and smarter.

A. Data Reduction vs Data Deduplication

NetApp Syncsort Integrated Back-up (NSB) uses unique technology to detect and copy only changed data blocks, eliminating unnecessary overhead during the back-up process. Other technologies in the marketplace also attempt to reduce the amount of data sent from the host, but there are significant differences in approach. Let's look at the two main methods.

Host-based Data Deduplication

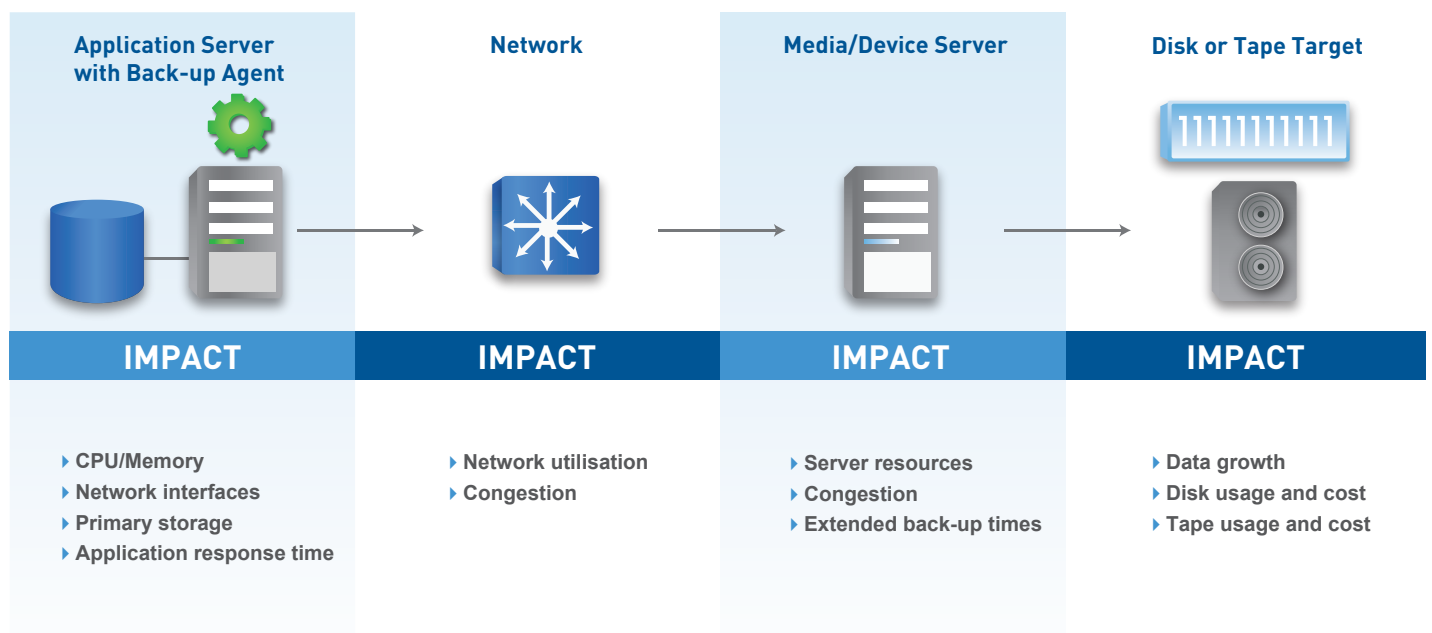
Data deduplication finds duplicate data blocks within a server data set, and then skips them during the back-up process. There are various technical ways of doing this, but all methods require that the data be scanned before back-up. That is how new blocks are detected. The process is both CPU and disk intensive, and tends to generate a significant utilisation spike on the host. Because the total back-up time is reduced, the overall

impact may be less than that of a lengthy file-based back-up, but the impact is still significant. This is one reason why most host-based dedupe solutions suggest only one back-up per day.

Host-based Data Reduction

Data reduction eliminates the need to scan data by tracking block changes as they happen. As NSB has implemented the technology, an agent on the host monitors changes at the disk layer – below the file system, to avoid file system overhead – and when the back-up time arrives, the new data blocks are already known. They are then copied to the target. This methodology eliminates the need to scan the file system, thereby drastically reducing impact. With impact nearly eliminated, data reduction lends itself to multiple back-ups per day, even as often as several times an hour. This lets you create far more data protection points, without accumulating extra data at the target.

Data Lift and Traditional File Back-up



“Data Lift” is defined as the entire, end-to-end impact of the back-up process created by the need to lift and move data from source servers to a target device. In a traditional file back-up environment, Data Lift creates significant impact at every step of the way, causing degraded performance and resulting in the need to continually expand resources, driving up costs.

While actual rates of deduplication may vary somewhat between the two approaches, it is clear that a *data reduction* model goes beyond deduplication to create far less impact on systems and lets your organisation easily meet all their back-up window needs.

Another advantage of host-based data reduction is that it eliminates the need to run repeated full back-ups. The traditional back-up model relies on daily incremental and weekly full back-ups (though it should be noted that due to restore complexities, many users do daily full back-ups of structured data such as databases and e-mail). The weekly full, in particular, is responsible for enormous amounts of needless Data Lift. Data reduction uses only one full back-up: the first, base level back-up. After that, only block level updates are required, yet every back-up job can be viewed as a full set of data. In other words, you have the operational advantages of a full back-up every time, without the Data Lift overhead.

We can quantify the impact of Data Lift with a simple example. Compare the Data Lift impact of a 100 TB primary data set over the course of a week. A target dedupe solution uses traditional file-based back-ups, with a weekly full and daily incremental model. NSB uses host-based data reduction. The data change rates are industry averages based on a mix of structured and unstructured data.

Clients - 100 TB Data Lift		
Target Dedupe	10% daily file change	10 TB x 6 days = 60 TB
	100% weekly full back-up	100 TB
Weekly total		160 TB
NSB	3% daily block change	3 TB x 7 days = 21 TB
Weekly total		21 TB

The comparison is instructive. In the target dedupe case, the Data Lift for a week is 160 TB. That will create significant impact across all key application server components. It is particularly problematic for Virtual Machine environments which are far more constrained in terms of spare system resources.

With the NSB solution, the application servers are processing 139 fewer terabytes per week. This is a phenomenal savings of 7 petabytes of Data Lift over the course of a year. And this is even with the unlikely

assumption that none of the servers receiving file back-ups will require daily full back-ups.

Clearly, the NSB model removes an enormous amount of impact from the application environment based on amount of data moved. This impact is even greater than noted here, because the client-side processing of NSB creates far less impact than an equivalent file-based agent. So it actually results in much lower impact than a 21 TB file back-up would create.

B. Back-up of Virtual Machines

The impact of back-up processing is particularly critical in virtual machine environments. The reasons for this are well understood: in a hypervisor environment of shared physical resources, far less spare capacity exists for back-up processing. When high-impact back-up methods are used, the results are predictable: failed back-ups, missed back-up windows, severe application response problems due to resource constraints. Some users try to work around this with complicated back-up schedules, trying to ensure that no more than two or three back-ups ever take place at the same time. But this adds greatly to management complexity and such schemes tend to be brittle and easily broken by something as simple as one back-up running long.

NetApp Syncsort Integrated Back-up is uniquely suited to VM back-up because of the minimal impact it creates, as in the Data Lift example above, along with very short back-up times. It works *within* the resource constraints of virtual machines, not outside of them.

C. The NetApp FAS Hardware Target Device

Every disk-to-disk back-up needs a target device. That's self evident. However, not all target devices are equal.

NetApp Syncsort Integrated Back-up utilises the architectural and design innovations found in NetApp FAS storage devices. Unlike most disk-to-disk models, NSB doesn't use a media or device server that sits in front of a disk array managing resources, handling snapshots and so on. Instead, NSB agents send data directly to the NetApp FAS, with no other device in the data path. This removes an entire layer from the traditional back-up

architecture, making back-up faster, more efficient and cost effective.

It also makes use of disk array resiliency, offering far better reliability than having a PC-based hardware layer sitting between server and target.

In addition, when NSB stores back-up images as snapshots, it relies on the field-proven, market leading NetApp Snapshot technology, which is well known for its high performance and reliability. When snapshots need to be mounted for instant data recovery, NSB utilises NetApp FlexClone to create minimal impact, zero footprint data images.

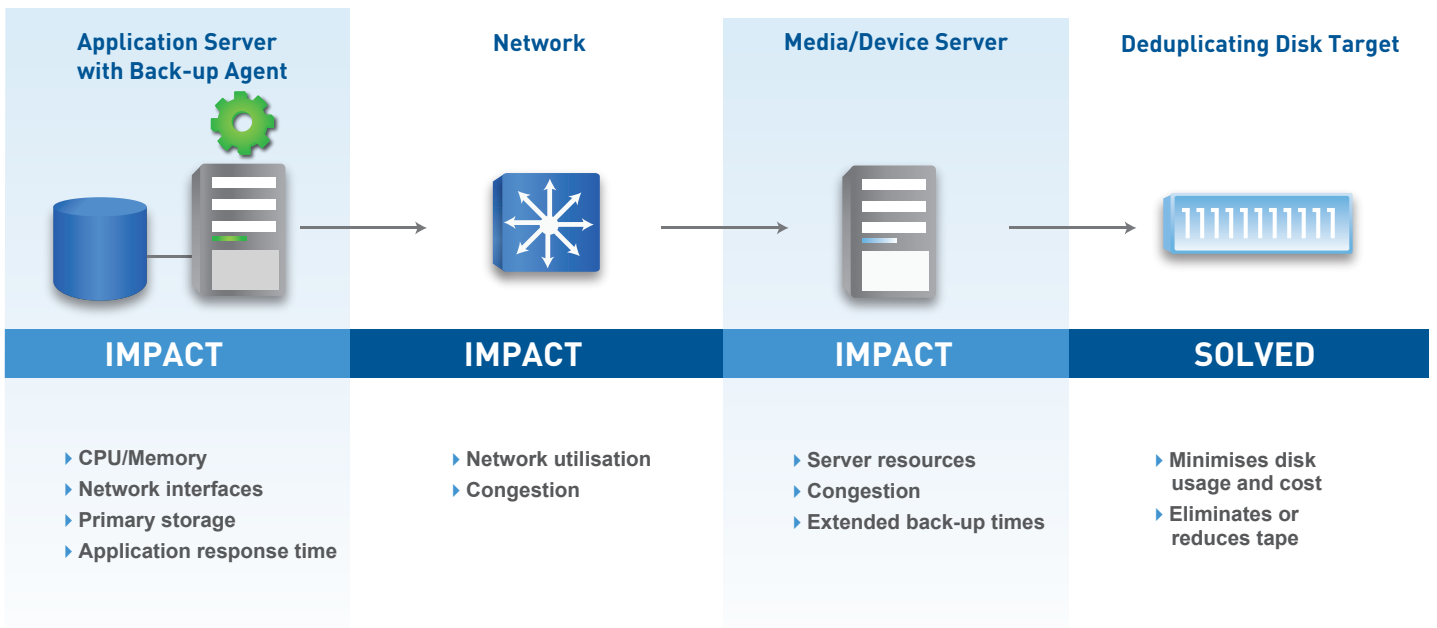
Finally, even though back-up sizes are already reduced by the block-level incremental back-up model, the NetApp FAS also includes data deduplication and compression. This further decreases the total amount of back-up data stored, allowing many more days of recovery to be maintained in a

small storage footprint, resulting in substantial cost savings.

III. Getting Data Back: Recovery without Compromise

It's a truism that we back-up data in order to recover it. But many data protection solutions fall down when the dreaded recovery task comes around. They are often complex and subject to user error, or the recovery process takes too long to meet ever shorter recovery service level agreements. Recovery failures are a significant part of why back-up is broken. When it comes to evaluating data recovery solutions, it is important to keep in mind that what you are recovering isn't "data" in the abstract. What you are recovering is your business. Speed and reliability are paramount.

Data Lift with Deduplication Appliances



Adding a deduplication appliance to a traditional back-up environment solves only one part of the Data Lift problem, that of the target storage capacity. By removing the duplicate data that back-up creates, deduplication makes disk-based back-up storage financially feasible, but it does nothing to remove the remaining impact across the environment.

A. Blocks are Better

Block-level back-up is not just a better and faster way to protect applications. It also enables better and faster recovery.

Traditional file back-up stores data in a non-native format (sometimes jokingly referred to as a “back-up blob”). In order to restore the data, it must be translated back to its original format. This takes time and extends the recovery window, especially when restoring large data sets.

With NSB back-ups, data is always accessible in the native data format. This means that back-ups can be accessed directly, via a simple snapshot mapping process. The benefits of this are discussed further below.

B. The Need for Multi-Tiered Recovery

Data and applications are lost in many different ways: human error, malicious actions, hardware failures, corruption – the list is endless. Failures also take place on many levels: files, volumes, LUNs, applications, servers.

Since all loss and recovery scenarios are not equal, it makes little sense to assume one recovery method is best for all scenarios. Yet that is exactly what many back-up administrators deal with every day because they have access to only one option: file recovery. Back-ups are based on files and restores are based on files, and the time-to-recovery always depends on the same thing: how long will it take to move one or more (or many more!) files from point A to point B. (In essence, this is the Data Lift problem in reverse.)

In the case of Tier 1 applications, there may be primary storage snapshots maintained by a SAN disk array. While these can be quite useful, they tend to be of limited duration due to the associated storage costs and impact. Yet snapshot-based restores are the best method known for restoring large data sets quickly. What better way to recover a corrupt database than by simply mounting an earlier instance of the data in a few minutes? Yet mounting a snapshot just to hunt for and find a single file

is overkill, not to mention operationally less efficient than a file-based restore via catalogue.

This is where the concept of *tiered-recovery* comes into play. Different types of recoveries are best served by different recovery methods. There are five common recovery scenarios to consider:

- File Level Recovery
- Single Item Recovery
- Volume/LUN/Object Recovery
- VM Recovery
- Physical Server Recovery

Let’s look at each of these in some detail to understand what is involved.

File Level Recovery

By far the most common recovery task, file level recovery is well understood by back-up administrators. The admin searches for the file or files in the back-up catalogue, and the restore is targeted back to the original source or an alternate location. The only hitch is when an older file is needed and the corresponding tape must be retrieved from an off-site vault. This can add hours or even days to the recovery.

Deduplication has largely alleviated this problem by allowing far more days worth of data to be kept on local disk. The vast majority of file recoveries are from data a few weeks old or less, so if 30 or 60 days of back-ups can be kept in place on disk, it solves 99% of your file recovery issues.

One caveat to keep in mind: In the case of host-based deduplication, many products do not offer a back-up catalogue (this bit of information is rarely volunteered by vendors, so make sure you ask). No catalogue means that a file recovery involves mapping to a series of snapshots and individually examining the file systems to find the file you are looking for. Not too bad if you happen to know exactly where the file resides. But if you don’t have precise information – an all too common problem – you can end up on a needle-in-a-haystack search that adds precious minutes or even hours to your recovery time.

Single Item Recovery

Single item recovery (SIR) is different than file recovery in that it refers to an item within a larger database file. The most common SIR request is for individual e-mail items (messages, contacts, calendar entries, etc), though SharePoint document recovery requests are growing rapidly.

Operationally, SIR presents a challenge because all too often the first step in recovering an item is restoring the entire database or storage group. This means that you need to restore a very large item (easily 500 GB or more) in order to restore an item that may not even be 1 MB in size. This is one of the ultimate Data Lift nightmares, where 99% of the time and effort is wasted.

A better approach is to utilise snapshots. SIR tools can quickly scan through the contents to find and recover one or more items. This is often how primary storage snapshots are used, if they are available. Ideally, you would have long-term back-ups that can be accessed

as snapshots to facilitate speedy single item recovery. Target-based deduplication solutions do not provide fast access to snapshots because they rely on traditional file back-up, and many host-side products also suffer similar limitations.

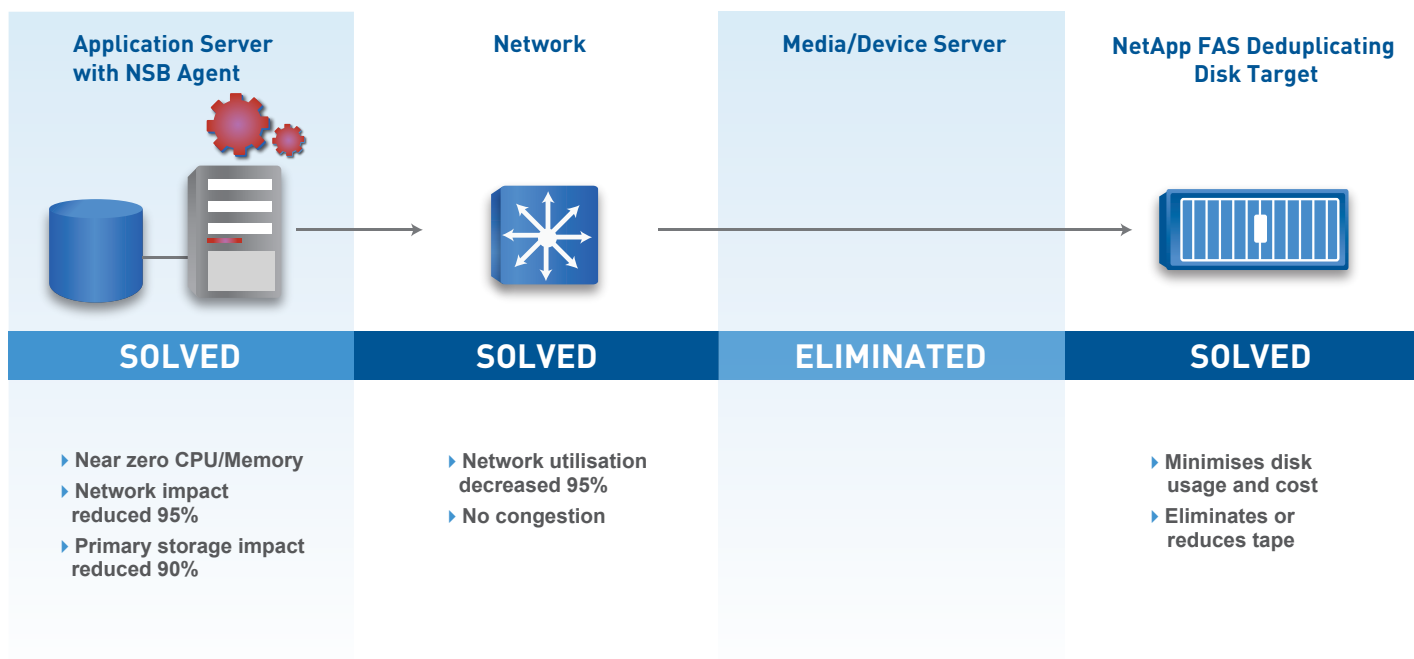
Volume/LUN/Object Recovery

While volumes, LUNs and Objects (eg, a database) are not the same thing, from a restore perspective they are similar enough to discuss together. All tend to be far larger than a typical unstructured data file (such as PowerPoint slides).

The source of data loss can be different. A volume may be lost through user error, a LUN through hardware failure, and an object through software corruption, but these lines cross often enough and are often contained one within the other like Russian dolls (eg, a failed LUN may have contained a volume in which an object resided).

The bottom line is that file level recovery is not suited to any of these cases. The amount of data is too large and

NSB and the Data Lift Problem



The NSB solution tackles the challenge of Data Lift by eliminating the problem at the source. Starting at the application server, NSB moves far less data than traditional solutions. This eliminates nearly all impact on applications. The benefits extend through the system by reducing both network traffic and storage footprint. In addition, NSB's unique architecture entirely eliminates the need for the media or device server, simplifying deployments, reducing potential points of failure, and significantly reducing costs.

file-restores are too slow. This is another case where snapshot-based data access can restore applications in a timely fashion.

Virtual Machine Recovery

With virtual machine workloads now surpassing physical machine workloads, VM recovery is something required of any data protection solution. VMs simplify recovery because of the way the entire server operating system is encapsulated into a few files. Hardware compatibilities are generally insignificant as they are handled by the hypervisor.

Still, it is important to consider two things with VM recovery: speed and operational ease. Because the VM is encapsulated in a single large “blob” file (ignoring for now any small configuration files), you need to restore the entire VM at once. As always, a file-based solution will recover only as quickly as the data can stream back to the target. Snapshots can greatly speed up the recovery process by making VM data available in minutes.

From an operational perspective, the recovery process must be well-integrated with the underlying hypervisor. Most products today have at least some level of integration and automation, but they do vary in terms of ease-of-use and sophistication. For example, with some solutions, you may need to first restore a full VMDK file, and then overlay that with additional file-level restores to get back to the most current data state. This both slows down the recovery process and introduces additional risk and complexity. VM restores should always be one-step.

Physical Server Recovery (Bare Metal Recovery)

Server recovery adds a wrinkle in that it's more than just recovering data. You have to first restore hardware-specific operating system information (drivers, registry settings, etc) that may or may not match the new, replacement hardware you are using for restore. Because the restore by definition uses new hardware, this is commonly referred to as Bare Metal Restore (BMR).

A key success factor in server recovery is the level of operational ease. It starts with the back-up. Some

products require an entirely separate back-up process for Bare Metal Recovery. Usually, this means only periodic BMR back-ups, spaced out between more frequent file back-ups. This greatly complicates restore, because you may have to overlay the BMR restore with one or more file recoveries.

During recovery, further complications arise due to hardware incompatibilities. When a server fails, it's usually because it's getting older. That means the same hardware model may no longer be available, and the replacement server has very different components that need different drivers. Some BMR tools put the burden of gathering the correct drivers on the user. Ideally, you always maintain up to date sets of drivers for all your server hardware that are readily available when needed. Realistically, very few IT departments do this with any consistency. More likely, when the disaster strikes you'll find yourself in a scramble to get everything you need together. Look for tools that provide you with up-to-date drivers while also being able to inject the new hardware drivers into the restore process.

With all these different recovery scenarios, many IT organisations find themselves deploying multiple tools to cover all the possibilities. Target-based deduplication devices do little to help with any of them, as they rely on the features of the file-based back-up software that drives them. As we have seen, snapshot-based restores are the key to fast and effective data recovery, but snapshots aren't a component of deduplication. Again, this is why the need to go *beyond deduplication* is paramount to achieving a truly effective data protection solution.

C. NSB: Complete, Multi-Tiered Recovery

- Rapid and efficient data recovery was a key design feature of NetApp Syncsort Integrated Backup from its inception. By utilising the robust efficiency of NetApp Snapshots, NSB shrinks recovery times from hours to minutes, no matter how much data is involved. Let's quickly look at our five recovery scenarios and how NSB handles them.

- **File Level Recovery:** NSB provides a searchable file catalogue that makes it easy to locate specific files to recover. They can be restored back to the original location, or to an alternate location.
- **Single Item Recovery:** NSB is application aware and understands discrete data objects, such as a database or a SharePoint site. If time is not critical, data objects can be restored in the traditional way by copying the object back to the target. However, when time is pressing, data objects can be accessed in the form of snapshots within minutes.
- **Volume/LUN/Object Recovery:** Any size data volume can be accessed within minutes through NSB snapshot mapping. By integrating the client agent with the NetApp FAS device, a FlexClone read/write image can be created and mapped to a server, all driven from within the NSB console.
- **Virtual Machine Recovery:** NSB automates virtual machine recovery by integrating operations with VMware. Only a few simple steps through an NSB recovery wizard are needed (these can be pre-configured) in order to recover any server back-up, whether from a physical or virtual source, as a new virtual machine. By utilising snapshots, the entire process takes only minutes. If preferred, a full VMDK file can be created by migrating data in the background after the VM has booted from the snapshot.
- **Physical Server Recovery (BMR):** NSB simplifies this often complex process by providing an easy to use BMR boot image that includes up to date drivers and allows for restore across different hardware. Fast, block level copy of data reduces data copy time to a minimum. For even faster recovery, BMR can be combined with snapshots. For example, BMR can be used to restore the boot drive, while snapshots are then used to connect to data drives, eliminating the

need to wait for the large volumes to copy.

Combining the operational benefits of a file catalogue with the speed and efficiency of NetApp Snapshots, NSB delivers multi-tiered recovery from a single solution and – very importantly – a single back-up. Additional point products or extra back-up steps are never required. The solution is comprehensive, unified and easily deployed across both physical and virtual server environments.

IV. Conclusion

Gartner is indeed correct in their estimation that back-up is broken. For too long, IT departments have been relying on decades old back-up and restore models that can no longer keep up with rapid data growth. Data deduplication, while important, is little more than a band-aid and doesn't resolve the majority of problems created by the challenges of Data Lift.

NetApp Syncsort Integrated Back-up takes you *beyond deduplication* to provide a simple, integrated and proven solution that solves all the challenges IT faces today from rapid data growth, accelerating server virtualisation and the need for ever-more-rapid recovery times.

About Syncsort

Syncsort is a global software company that helps the world's most successful organisations rethink the economics of data. Syncsort provides extreme data performance and rapid time to value through easy to use data integration and data protection solutions. With over 15,000 deployments, Syncsort has transformed decision making and delivered more profitable results to thousands of customers worldwide. Syncsort Incorporated, headquartered at 50 Tice Boulevard, Woodcliff Lake, NJ 07677, USA. www.syncsort.com

syncsort RETHINK THE ECONOMICS OF DATA[®]

Syncsort Ltd., UK, Tel +44 (0)1732 849000

www.syncsort.co.uk

Syncsort S.A., France, Tel +33 (0)1 46 67 46 10

www.syncsort.fr

Syncsort GmbH, Germany, Tel +49 (0)2102-9546-0

www.syncsort.de