Storage Efficiency and Performance of Commvault® on Cisco® UCS S3260 Storage Server

Contents

Introduction

Unprecedented data growth presents new challenges for enterprise organizations worldwide. Current solutions with outdated architectures are failing under the weight of these data loads which are too often stored in disparate, disconnected silos. The result is spiraling costs, the risks associated with the complete inability to manage and protect large data pools, and lack of visibility and access to data stored in data silos. This trend is unsustainable.

To address the challenges highlighted above, two industry leaders, Commvault and Cisco, have come together to provide the next generation of scale-out storage, protection and management. This required a fresh approach to data management with a revolutionary architecture. This is the only solution available with enterprise-class data management, storage and protection services that takes full advantage of an industry-standard Commvault scale-out infrastructure together with Cisco Unified Computing System.

Organizations today are facing an increasingly chaotic technology landscape, and with the rapid growth rates seen in not just data size and types, but also the proliferation across location and services, it's important to gain perspective on where technology can be a benefit. It's a hybrid cloud or multi-cloud world and choosing a solution that intelligently optimizes more than storage is more than a table stakes conversation, it's a fundamental requirement for business agility. The combination of the Cisco Unified Computing System with Commvault Complete™ Backup and Recovery running on the Commvault® Data Platform provides the ultimate solution to effectively address all these critical issues for today's enterprise.

Purpose of this document

This document gives an overview about technologies and options available from Commvault running on Cisco UCS S3260 Storage Server for optimized storage efficiency and performance.

Cisco Unified Computing System

Cisco® UCS is a state-of-the-art data center platform that unites computing, network, storage access, and virtualization resources into a single cohesive system.

Cisco UCS consists of these main resources:

- Computing: The system is based on an entirely new class of computing system that incorporates rack-mount and blade servers using Intel® Xeon® processor CPUs. The Cisco UCS combines industry- standard, x86-architecture servers with networking and storage access into a single unified system. UCS brings increased productivity, reduced total cost of ownership, and scalability into your data center.
- Network: The system is integrated onto a low-latency, lossless, 10/25/40/100 -Gbps Ethernet and 8/16/32-Gbps Fibre Channelunified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing (HPC) networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables, and by decreasing the power and cooling requirements.
- Virtualization: The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- Storage access: The system provides consolidated access to both SAN storage and network-attached storage (NAS) over the unified fabric. By unifying the storage access layer, Cisco UCS can access storage over Ethernet (with Network File System [NFS] or Small Computer System Interface over IP [iSCSI]), Fibre Channel, and Fibre Channel over Ethernet (FCoE). This approach provides customers with choice for storage access and investment protection. In addition, server administrators can pre-assign storage-access policies for system connectivity to storage resources, simplifying storage connectivity and management for increased productivity.

Figure 1. Cisco UCS Manager

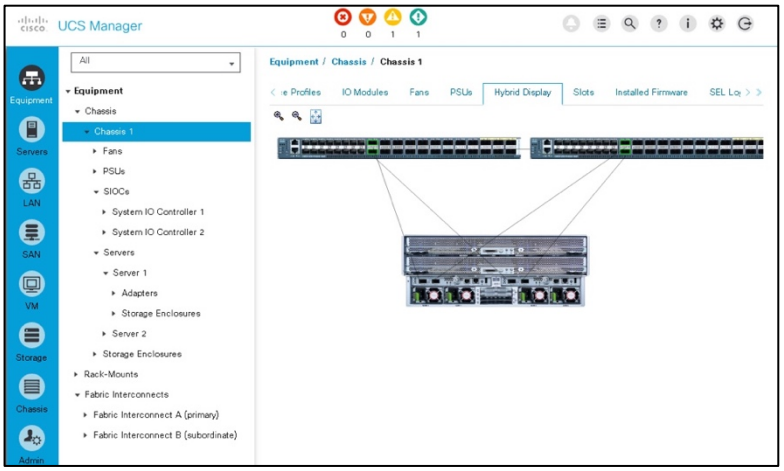

Cisco UCS consists of the following components:

- Cisco UCS Manager provides unified, embedded management of all software and hardware components in the Cisco Unified Computing System (Figure 2).
- Cisco UCS 6000 Series Fabric Interconnects are line-rate, low-latency, lossless, 10- or 40-Gbps Ethernet and FCoE interconnect switches that provide the management and communication backbone for Cisco UCS.
- Cisco UCS 5100 Series Blade Server Chassis supports up to eight blade servers and up to two fabric extenders in a 6-rack-unit (6RU) enclosure.
- Cisco UCS B-Series Blade Servers are Intel-based blade servers that increase performance, efficiency, versatility, and productivity.
- Cisco UCS C-Series Rack Servers deliver unified computing in an industry-standard form factor to reduce total cost of ownership (TCO) and increase agility.
- Cisco UCS S-Series Storage Servers deliver unified computing in an industry-standard form factor to address data-intensive workloads with reduced TCO and increased agility.
- Cisco UCS adapters with wire-once architecture offer a range of options to converge the fabric, optimize virtualization, and simplify management.

Cisco UCS is designed to deliver:

- Reduced TCO and increased business agility
- Increased IT staff productivity through just-in-time provisioning and mobility support
- A cohesive, integrated system that unifies the technology in the data center
- Industry standards supported by a partner ecosystem of industry leaders
- Unified, embedded management for easy-to-scale infrastructure

Cisco UCS S3260 Storage Server

The Cisco UCS S3260 Storage Server (Figure 2) is a modular, high-density, high-availability dual-node rack server well suited for service providers, enterprises, and industry-specific environments. It addresses the need for dense, cost-effective storage for the ever-growing amounts of data. Designed for a new class of cloud-scale applications, it is simple to deploy and excellent for data protection applications such as Commvault Data Platform, software-defined storage environments and other unstructured data repositories, big data applications, media streaming, and content distribution.

Figure 2. Cisco UCS S3260 Storage Server

The Cisco UCS S3260 helps you achieve the highest levels of performance and capacity. With dual-node capability that is based on the Intel® Xeon® Scalable processors or Intel® Xeon® processor E5-2600 v4 series, it features up to 720 TB of local storage in a compact 4-rack-unit (4 RU) form factor. The drives can be configured with enterprise-class Redundant Array of Independent Disks (RAID) redundancy or as JBOD in pass-through mode.

This high-density rack server easily fits in a standard 32-inch-depth rack, such as the Cisco R42610 Rack.

Cisco UCS S-Series Storage Servers can be deployed as standalone servers or as part of a Cisco UCS managed environment to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' Total Cost of Ownership (TCO) and increase their business agility.

The Cisco UCS S3260 uses a modular server architecture that, using Cisco's blade technology expertise, allows you to upgrade the computing or network nodes in the system without the need to migrate data from one system to another. It delivers:

- Dual server nodes
- Up to 36 computing cores per server node
- Up to 60 drives, mixing a large form factor (LFF) with up to 28 solid-state disk (SSD) drives plus 2 SSD SATA boot drives per server node
- Up to 512 GB of memory per server node (1 TB total)
- Support for 12-Gbps serial-attached SCSI (SAS) drives

- A system I/O controller with a Cisco UCS Virtual Interface Card (VIC) 1300 platform embedded chip supporting dual-port 40-Gbps connectivity
- High reliability, availability, and serviceability (RAS) features with tool-free server nodes, system I/O controller, easy-to-use latching lid, and hot-swappable and hot-pluggable components

For updated information on Cisco UCS S3260 Server Platform, go to: https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-s3260-storage-server/index.html

Commvault

Commvault Data Platform with Commvault Complete Backup and Recovery

The Commvault Data Platform is a single platform for automated global protection, retention, and recovery. Commvault enterprise data protection and recovery software automates global data protection, speeds recovery, reduces costs, and simplifies operations. Commvault integrates application awareness with hardware snapshots, indexing, global deduplication, replication, search, and reporting. The Commvault Data Platform converges all the needs of a modern data management solution in one place to seamlessly integrate protection, management, and access in one solution that features Commvault Complete Backup and Recovery and protects and manages data regardless of where it lives—on premises, in private cloud or in public cloud (including SaaS and PaaS options).

A comprehensive data protection and management strategy offers seamless and efficient backup, archiving, storage, and recovery of data in your enterprise from any operating system, database, and application. To protect and manage data in your environment, the Commvault software must be distributed to systems that you want to protect. The CommServe®, MediaAgent and protected systems constitute a CommCell® environment while each protected system is referred to as a client.

Figure 3. Commvault Data Platform overview

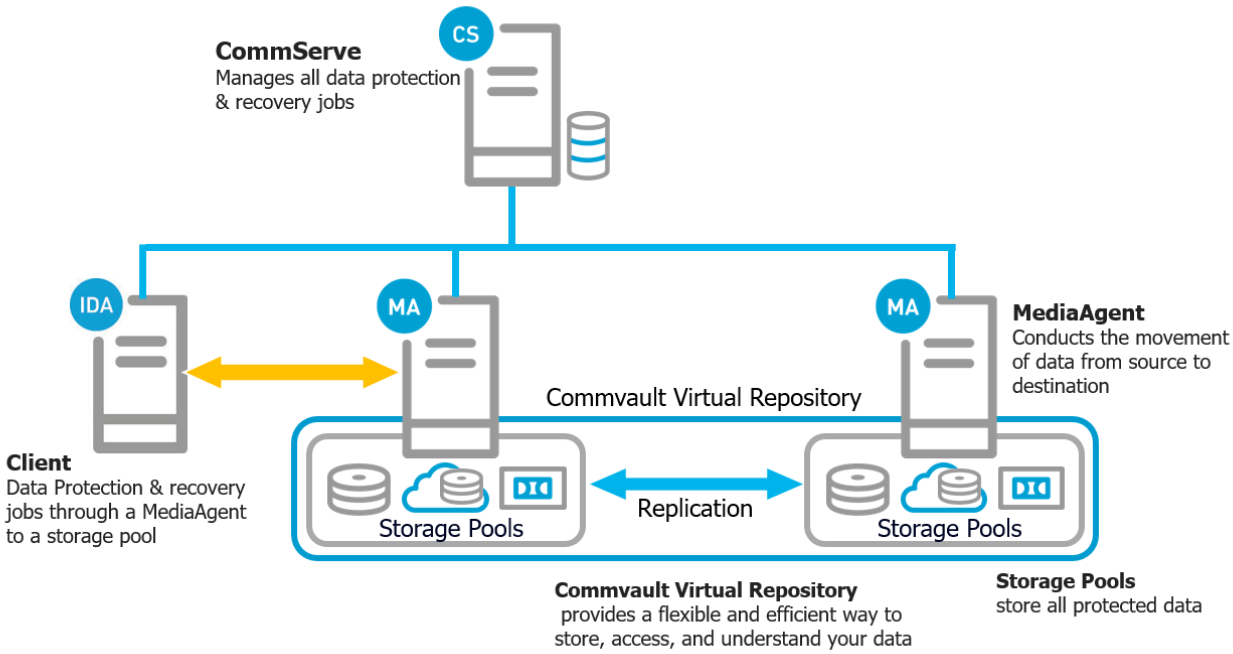

The CommServe server is the command and control center of the CommCell architecture. It coordinates and executes all CommCell operations, maintaining Microsoft SQL Server databases that contain all configuration, security, and operational history for the CommCell environment. There can be only one CommServe host in a CommCell environment. The CommServe software can be installed in physical, virtual, and clustered environments.

The MediaAgent is the data transmission manager providing high performance data movement and manages the data storage pools. When installed on a client system it also manages the Commvault IntelliSnap snapshot integration with the underlying storage.

The Client is any system within a CommCell environment to be protected. iDataAgents are software modules that are installed on computers to access and protect data. The backup and recovery system uses agents to interface with file systems, applications, and databases to facilitate the protection of data on production systems. By default, a filesystem iDataAgent is installed when adding the Commvault software to a system. If the client hosts specific applications or databases, the installation of additional iDataAgents are required.

These three Commvault components combined result in the most comprehensive and flexible data protection solution on the market today.

Deduplication

Deduplication is in use in several technologies today, however the depth and benefits of it can differ between use case and implementation. In this whitepaper we will discuss general deduplication technology in broad strokes, this not meant to generalize the technology, but to highlight that there are a significant number of the variations in use today.

The simplest definition for deduplication for this whitepaper is as follows:

> **deduplication**
> dēˌd(y)o͞oplə'kāSH(ə)n/
>
> *Noun*
> noun: **deduplication**; plural noun: **deduplications**; noun: **de-duplication**; plural noun: **de-duplications**
>
> the elimination of duplicate or redundant information, especially in computer data.
> "deduplication removes the repetitive information before storing it"

Deduplication looks to eliminate redundant information. On the surface this seems simple enough, but it is typically represented in deduplication ratios and savings in storage. Today's datasets are vast and with so many different places where deduplication can reduce data it can be daunting to understand what is important. Eliminating redundant information is important because it can reduce the amount of data stored at the storage layer, but that only represents part of the challenge.

Regardless of the technology used, for redundant data to be eliminated it first must be found. With a plethora of choices available for enterprises today how a vendor implements their deduplication technology fundamentally dictates how and where the savings and efficiency are determined in a given environment.

Compression vs. Deduplication

Compression is a form of deduplication, it's not the focus of this whitepaper, but it is important to understand what it's for. Compression technology reduces the size of a file by removing the redundant data, within the file only. This provides space savings for each individual file that is compressed.

A simple example would be to compress a 1MiB text file. Compressions will look inside that file any redundant data, for instance there is a lot of white space or the same character pattern that repeats and can therefore be replaced with a pointer that is stored instead of the full footprint. This could reduce a 1MiB document to a fraction of its size, when it comes to text files its normal to see a 50% reduction in space, or 1MiB down 512KiB.

Deduplication is a topic that is taken for granted today, and why not? In some form or another it's been around for decades. Before it became a common place technology, forms of deduplication are almost as old as computing in general. Once it was determined that the budgets associated with storing the data were not infinite, manual operations were looked at to reduce the amount of data stored. This was generally done by humans looking for duplicate information and then either removing or combining those duplicate items. It wasn't automated, nor would it be considered technology like we have today, but it was the same concept. Today deduplication is treated as a checkbox in a technology conversation, something to ensure that is included, but are all forms of deduplication equal, and do they achieve the same result?

Deduplication provides a great number of benefits in solutions and technology, and for that very reason this whitepaper will outline some of the various forms of deduplication and highlight that not all technology is created equal. Deduplication technology doesn’t provide the same benefits in all situations, nor should it be the only technology that should be relied upon.

When looking at secondary workloads and datasets, which account for the vast majority of storage today, it's important to address the requirements that will achieve optimal results in today's hybrid world. Regardless of where deduplication is utilized it's important to understand that there are more components involved than just storage reduction. Deduplication technology can have a significant impact on:

- Service Level Agreements – Both RPO and RTO
- Infrastructure Costs – LAN, WAN, Infrastructure, and Operational Efficient
- Impact on both primary and secondary infrastructure

Achieving breakthrough performance typically focuses on how fast the clients and the solution can process the data for protection. If production infrastructure is designed to focus on the serving production workloads, should the definition of modern data protection still be bound by how fast the client can send data? Taking more resources for secondary workloads can have an adverse effect on the production workloads.

When looking at the SLAs, clients speed is typically what drives the SLA in terms of what Recovery Point Objective is possible. Wouldn’t it be beneficial if that wasn’t the determining factor?

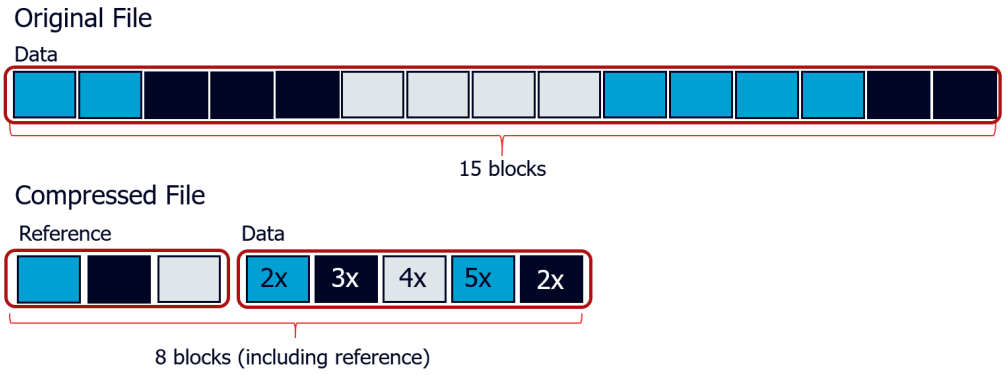

Compression does not look at data across multiple different files, just within a single file itself. Therefore, if the text file from the previous example was copied multiple times each new copy would consume space even though the data is identical from one copy to the next.

Compression can be applied to a great number of file types, and this technology is built-in to several production applications and filesystems today. With compression active in several areas that are in production today it can alter the amount of savings that can be found when deduplication for backup technology is used. No two organizations data footprint is identical it's important to understand that not every conversation about data reduction will apply in each use case.

One of the main differences between compression and deduplication is the ability to identify redundant patterns between datasets, not just within the same dataset. Compression is just one element of data reduction that can change the outcome of deduplication technology, because it is in fact a deduplication technology. The impact is that sizing is typically based on guidelines of what the general industry will see, and mileage may vary in a specific dataset or use case.

Transport Layer

Data needs to be processed for protection, however where will the deduplication processing be done?

Where deduplication processing occurs is an equally important factor, if not more important, than the space savings for some companies. There are two main transport methods for deduplication processing, source-side meaning the data is processed before it leaves the client, or target-side done at the MediaAgent or even the storage layer.

Both methods can achieve the space savings and efficiency as discussed previously, which saves on the storage layer, but not all methods will save on transport costs.

Target-side Deduplication

Target-side deduplication has a couple of different variants, which do have some impact on processing but the critical component to understand is that the processing of data for deduplication is done at the target, and not at the client itself.

Traditional target-side deduplication requires all data to be sent from the client across the network to be processed, even if the data has previously been stored and deduplicated. This achieves storage savings, but it has a heavy burden in terms of data transmission. Target-side deduplication is the least efficient in terms of network bandwidth utilization, clients must read their data and send it across the network to look for the unique blocks of data.

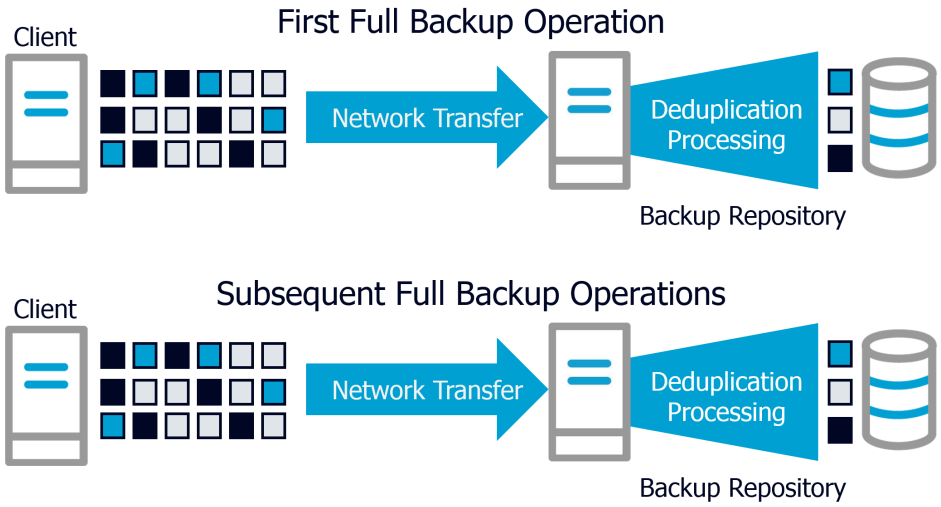

The above examples compares 2 traditional full backup operations, the initial full or baseline, and then a subsequent full backup operation. In each case, and with every other traditional full backup operation the client and the network must send all the data to be processed on the target side.

Target-side deduplication typically runs into a simple constraint, available network bandwidth. Backup windows are always shrinking, and companies are looking to improve their SLAs surrounding data protection all the time.

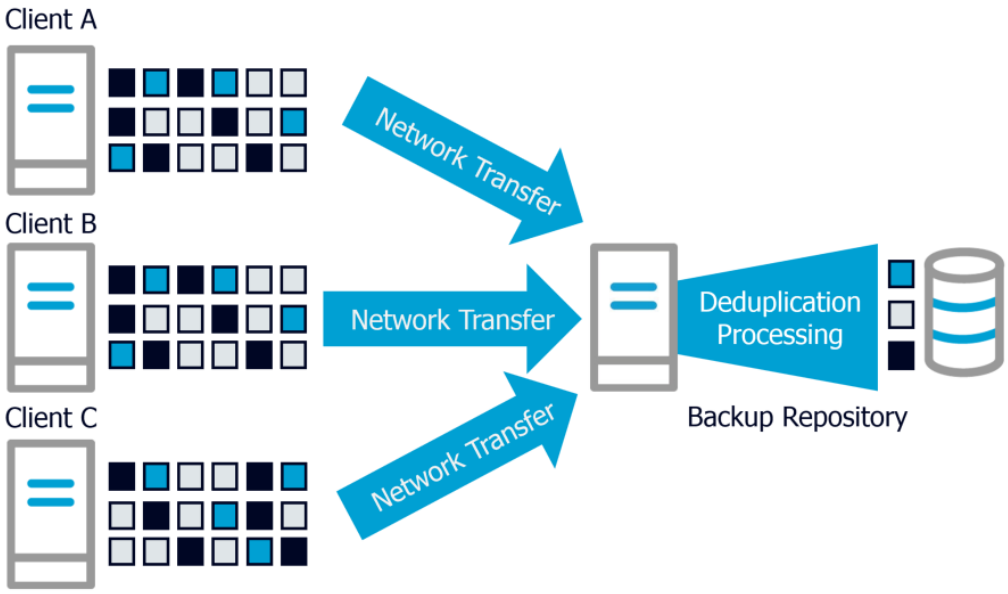

In the above example, let's say that those 3 clients represented 100TB of data. All that data for the first full, and each subsequent full operation must be moved across the network for processing. Several companies have standardized on 10Gbps infrastructure for both clients and backups, but even at that speed the theoretical maximum of a single 10Gbps network card is 4.5TB/hr. of data

movement. Typically, in the real world the maximums aren't typically achieve due to several other factors, such as production application traffic, but even at this maximum speed it would take just over 22 hours to move 100TB of data.

With most backup windows much smaller than 22 hours, copying 100TB of data every night and processing it at the destination likely wouldn't meet the requirements of most companies SLAs today. This would provide a customer a RPO objective of 22 hours, however realistically it would be 1 day, due to scheduling and backing windows.

This example of course uses a LAN connection. WAN links further complicate the situation. Today's distributed environments can include hybrid cloud scenarios, and many remote sites. Moving data across a WAN introduces new costs and efficiency problems.

Target-side deduplication can be found in many appliances and storage arrays today, and while they can provide storage savings all the data needs to be transported to them which can be costly and time-consuming. This type of technology doesn't address the constant client utilization and the network traffic utilized for all operations.

Source-side Deduplication

Source-side deduplication shifts the processing of deduplication to the source, the client (or a proxy to the client that can access the data directly). By shifting the processing to the source, the client will read and search for unique data patterns that match against a central repository. This method ensures that only the unique data is sent across the network. It provides the same savings in terms of storage efficiency as target-side deduplication, but it also is highly efficient in terms of network bandwidth utilization.

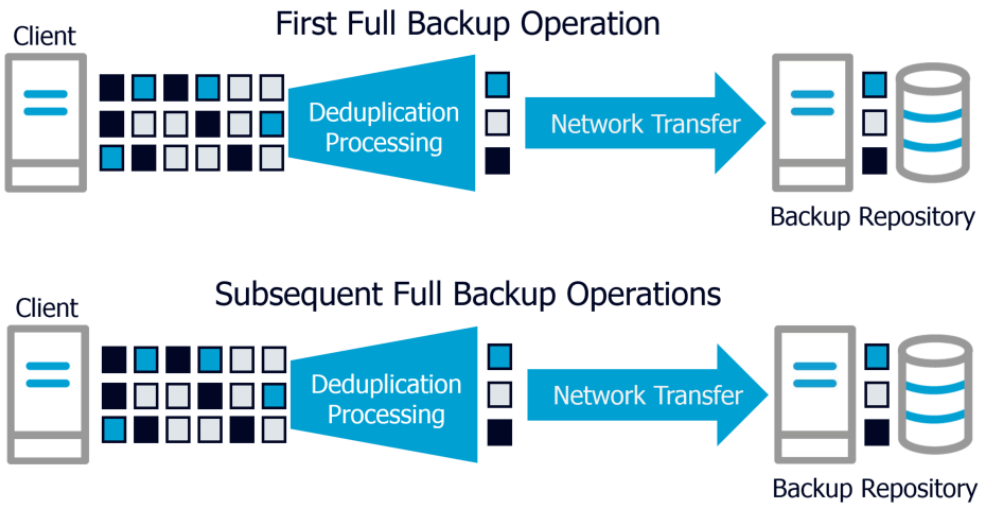

Another benefit of source-side deduplication is that the processing for deduplication can be spread across a larger number of clients instead of centralizing all the work into dedicated processing hosts.

Source-side deduplication moves the processing to the source, which maximizes the network efficiency of the backup for each individual client. Although the processing requirements for target-side deduplication can be heavy, the client requirements are light because the footprints for the individual source workloads are smaller than the overall footprint for target-side processing. The data is read, and the deduplication checksums are calculated, and compared against a central repository, if that block is unique it will send the block across the network, otherwise the data payload isn't sent. This has a massive impact in terms of both backup window reduction and network traffic reduction for both LAN and WAN configurations. MediaAgents can process significantly more data as they are no longer constrained by network infrastructure and the client infrastructure doesn't get saturated sending data to process.

All things being equal source-sided deduplication is clearly the choice for maximum bandwidth and storage savings. There are a couple of use cases in which a traditional source-side deduplication scenario isn't utilized, one of which is when the client itself is completely overloaded. In these scenarios some solutions will opt to move back to a target-side deduplication as any client impact, however negligible is unacceptable.

There are additional features available as part of the Commvault Data Platform that can alleviate the impact to a client, an example one of which is IntelliSnap Technology which directly integrates with the storage layer and orchestrates snapshots that can then be mounted to an external proxy that has its own processing power that wouldn't impact the production compute layer. In effect that proxy becomes the source, without the client being impacted. This is just one example of additional processing that can automatically be utilized as part of the Commvault Data Platform, however listing them all is beyond the scope of this whitepaper.

Service Level Agreements

Data protection is not built on the principle of a single operation. As data in modern enterprises change the backup and recovery system needs to keep up with those changes. Service Level Agreements (SLAs) normally include components about Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs).

RPOs focus on the amount of data loss is tolerable for an enterprise, the more critical the data set the less amount of data loss is acceptable. If a data protection operation happened just once a day, that would mean there is a single daily recovery point, and it would be a maximum of 24 hours. The maximum RPO in that example would be 24 hours (there can be a couple of extra variables on that, but in general terms). Moving the frequency up to every 4 hours would provide an RPO of a maximum of 4 hours.

Of course, there are scenarios where no data loss is acceptable, and there are several ways to achieve that, however that typically takes a layered approach and is a larger conversation around business continuance and disaster recovery. Replication can provide a replica copy of the data for instance, but if the data was corrupted at the source the corruption typically replicates along with it. IntelliSnap technology is another way that data can be protected in plays with the integration of storage arrays. There are several additional solution components that Commvault and Cisco can utilize in a modern data protection solution. The focus of this whitepaper however is on deduplication and streaming technology.

RPOs combined with retention determine how many recovery points are available for recovery. Use cases for recoveries differ as well, for instance an accidental file deletion sometime in the past month versus a database corruption or failure. The accidental deletion might not be noticed for a few days, or even weeks, therefore retaining multiple backup copies can show the historic view of the information and recover that single file from days or even weeks ago. A database problem however is typically something that requires the most immediately available file to ensure that any data loss is mitigated.

In traditional processing frequency and retention traditionally dictate how much storage reduction will be achievable. In general terms if the amount of change rate in the backup data was 0 then backups performed daily and stored for 30 days would have the equivalent reduction of 30 backups processed in a single day.

Deduplication Ratios

Deduplication effectiveness is represented in the industry typically as a deduplication ratio, the larger the ratio, the better the savings are. Ratios listed by vendors aren't always realistic, they are favorable ratios or the best that a vendor can achieve under certain conditions. It's a data point, but it is often used as a yardstick for how effective a solution would be. Will this technology provide me with a 20:1 reduction? A 30:1 reduction? More? Less?

Is the total worth of a solution determined by a marketing brochure with a ratio that has an asterisk beside it?

If deduplication ratios were the only thing that mattered, how about a 1,000,000:1 deduplication ratio? Anyone providing that would have the best product on the market!

Here's the thing, that deduplication ratio can easily be achieved. Here are a couple examples on how that would work:

- Take a file, copy it 1,000,000 times and back them up. Since each file would be identical to the others the data would be highly redundant also, after the first file is stored all remaining files would therefore be deduplicated against the first.
- Take a VM, assuming no data is changing inside of it, and back it up 1,000,000 times in a row. Again, since the data would be identical in the VM performing the same operation over and over would achieve a great deduplication ratio.

How relevant are either of those scenarios? Not very.

The truth about deduplication ratios is that they are a number, and that number is typically associated in marketing collateral to promote the cost savings that could be achieved. Like purchasing a car, the marketing material associated with the average miles per gallon (MPG) provides the fuel efficiency for the vehicle. Is that real world achievable? Is that really what will happen for every car or is that providing information about what could be achievable, however your mileage may vary.

When two solutions are compared on deduplication ratios alone, it's easy to see how it can be misleading. Take the following ratios for example 10:1 versus 20:1. If the only number that is reviewed is the deduplication ratio, the effectiveness of the 20:1 solution must be double of the 10:1 solution. Is that really the case though?

Deduplication ratios also translate direction to the storage reduction achieved. The following chart represents the storage savings beside the deduplication ratio.

Deduplication Ratio	Deduplication Reduction Percentage
2 : 1	1/2 = 50%
5 : 1	4/5 = 80%
10 : 1	9/10 = 90%
20 : 1	19/20 = 95%
30 : 1	29/30 = 96.6%
100 : 1	99/100 = 99%
200 : 1	199/200 = 99.5%

The real difference between a deduplication ratio of 20:1 and that of 10:1 ratio is 5% additional storage savings. A 10:1 deduplication ratio would provide a 90% savings of storage, while a 20:1 deduplication ratio would provide 95%. The ratios themselves don't state that the savings are double, they are just higher, and not typically explained.

Of course, the higher the ratio, the better the savings, right?

Yes, while technically true, after a certain point there is a point of diminishing return to the achievable results. Just as 10:1 versus 20:1 provide just a 5% difference in storage reduction, 20:1 versus 30:1 the amount of savings is further diminished. The additional storage reduction difference drops to just 1.7% between to two. 20:1 is 95% storage reduction, while 30:1 is 96.7%.

Storage reduction is an important component of the overall solution, after all it represents the overall savings from the reduction in actual storage needed to house the data. After a certain level however, the higher the ratio doesn't save that much more storage, it just becomes an impressive number. The chart below shows the deduplication ratio contrasted against the storage reduction.

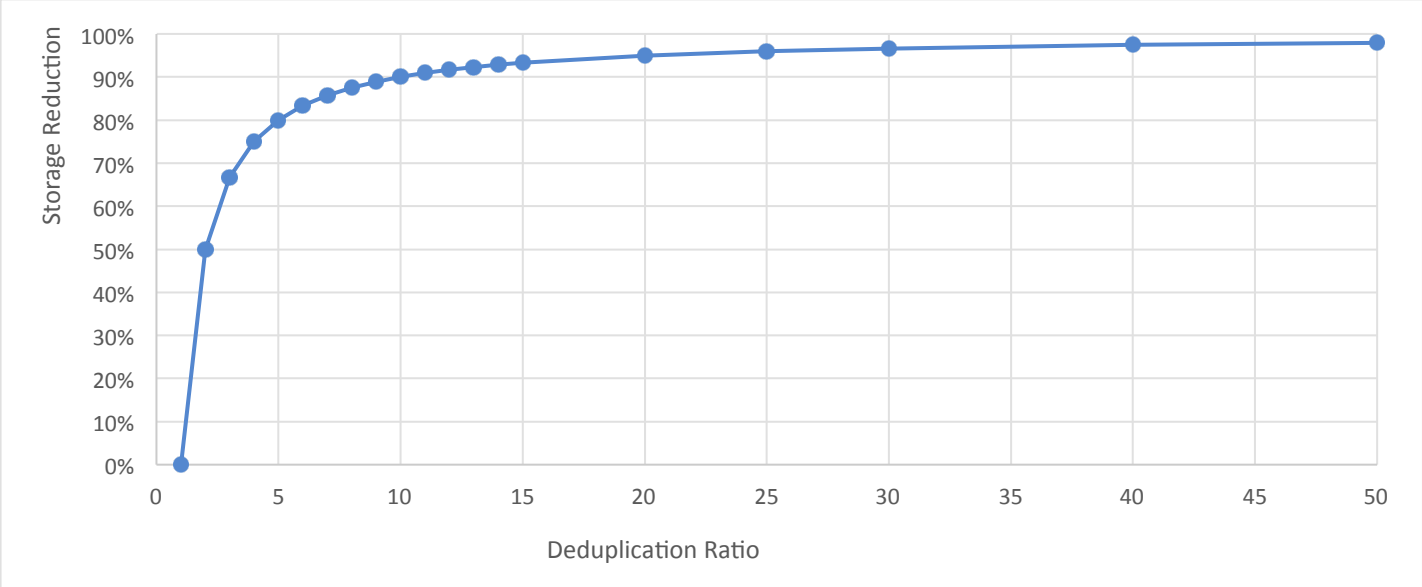

Deduplication and their represented storage savings are part of just about every modern backup solution. After a certain point however, the savings are diminished and are not alone a value add. Additionally, as previously outlined being able to achieve a 1,000,000-deduplication ratio can also be done, though unrealistic because who would continually backup the same data just for a better deduplication ratio?

If we change the 1,000,000-file example from identical files, to completely unique files and back them all up once, what would the deduplication ratio be? Effectively 0% and that's a critical point. The ratios provided for most solutions are based on moving the same amount of data over and over to increase the ratio. Regardless if the data is unique or not, simply performing the same operation will improve the ratio if the data hasn't changed. While this does provide a better ratio, the cost of this isn't free.

The production infrastructure must re-read the same stale data over and over to achieve a better ratio. Looking at the deduplication by itself can be misleading, and leads us to a larger question, if the data hasn't changed, why would it need to be protected again? It wouldn't.

Data Processing

Being able to reduce a storage footprint for data protection by 90-95%+ has transformed data protection solutions and eliminating network traffic, whether it is LAN or WAN, has been equally critical. On the surface this makes great use of technology that provides some elements of a modern data protection solution. It doesn't however have a positive impact on all areas of the infrastructure.

There is an ugly truth about these deduplication ratios when looking for the redundant data, traditional solutions will do the same operations repeatedly to get those savings. The same data must be re-read continuously, and while this has a positive impact in some areas, the production storage infrastructure certainly doesn't benefit from this.

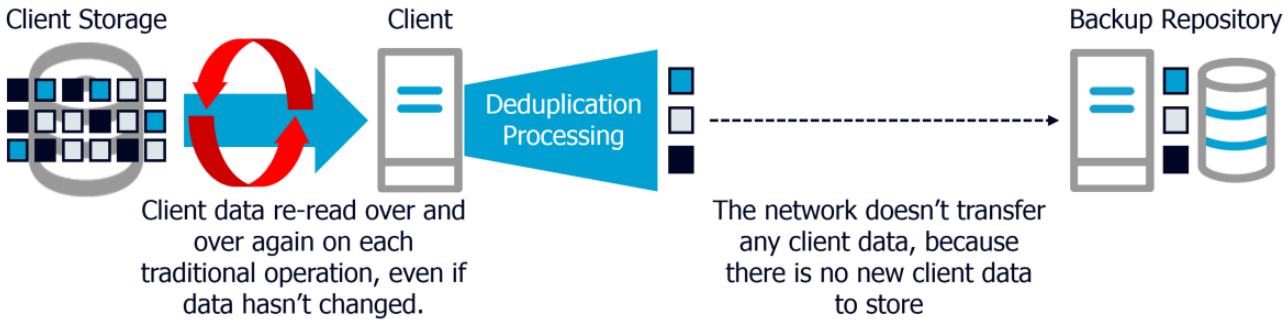

The above example now includes the client storage layer, and traditional solutions will put additional constraints on the production storage to look for new data for backing up. To achieve a 20:1 deduplication ratio the client would have to perform at least 20 full backups, and that's if there were no changes on the client itself. It's highly unlikely that the client data isn't changing at all, so if that is factored in the amount of operations might be closer to 25 to 30 operations to achieve that ratio. Put another way, if that client had 10TB of data, it would have to read at least 200-300TB of data to achieve a 20:1 deduplication ratio. This is a high burden cost for production clients.

Production applications are typically architected for the production workload, and that doesn't typically factor in the workload of secondary operations such as backup, dev/test copies, or other secondary operations. For instance, a critical application that has a high read/write requirement could use Flash/NVMe or some other type of acceleration due to the intensity of the workload. Layering on top of this the requirement for multiple recovery points could lead to an aggressive data protection policy of maybe 6 times a day, or even 12. Reading all the data on the drive 12 times a day can have a significant impact on the production infrastructure.

Eliminating redundant data, achieving high deduplication ratios, and saving LAN/WAN bandwidth makes for excellent datasheet material. If it just shifts the work to the client and that places an additional burden on the production infrastructure, then that's not a completely modern solution.

Full vs. Incremental

Before disk-based solutions were utilized for backup operations, tape was the predominant technology. Tape systems struggled to keep up with the data growth and it became clear that performing full backup operations nightly would overrun the nightly backup windows.

Shifting full backups to the weekend where the backup windows could be significantly large enough to house a full backup while performing incremental backups during the weeknights became the main data protection strategy. This would allow the data to be protected in the desired protection windows.

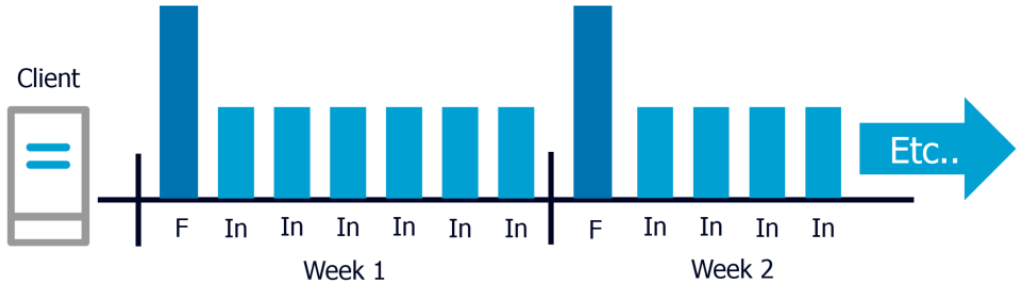

There was a tradeoff with this shift however, in the event of a recovery the data for recovery could be housed across several tapes and therefore recovery times would be extended. A recovery operation could be up to 7 different actual operations because the most updated data that was stored is no longer in the same physical space.

Fast forward to the use of disk basked protection with deduplication and a number of these elements are still in play today. Full backups on the weekend with incremental processing during the week, the difference of course is that all the physical media is together. It can process data that much faster than tape, and with deduplication the savings are lightyears ahead of what tape provided, however the data growth hasn't stopped either and this of course means that processing full backups each night isn't realistic.

A client with 10TB of data and a 4-week retention of data would realistically perform 30 full backups. Let's take a simple example of a client that is changing 1% data each day.

Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
Full – 10TB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB
Full – 10TB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB
Full – 10TB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB
Full – 10TB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB	Incr – 100GB

Processing on this cycle there would be four full backups and 24 incremental backups. If this client was using source side deduplication, the deduplication ratio could be extremely low, perhaps 4:1 or less. Since the system is not processing all the data over an over the deduplication ratio represented wouldn't be high, although the storage space utilized would be about the same. The client has avoided the burden of processing the same data repeatedly.

So, while it might seem counterintuitive a lower deduplication ratio can be a good thing.

Using incremental data protection operations isn't new, and it in traditional data protection operations it doesn't solve all scenarios. For instance, in other traditional incremental approach multiple recovery operations could be required. Take the above example, if there was a full recovery operation needed it might need to perform up to 7 recovery operations to get all the data back, starting with the full, and then recovering each incremental on top.

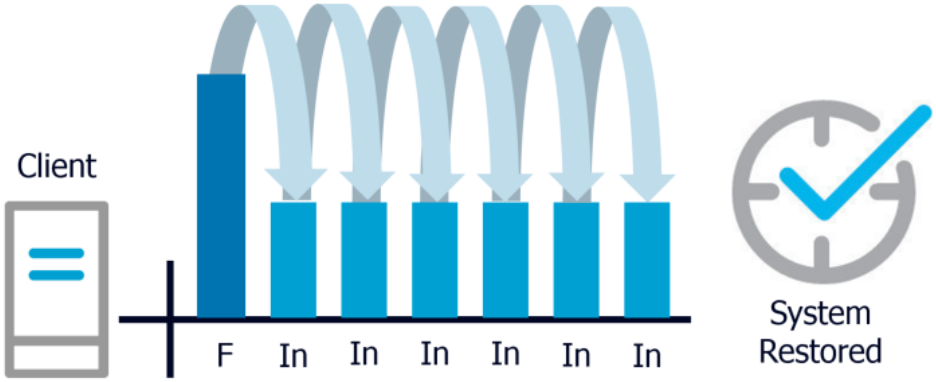

Clearly, that's not efficient. It solves one problem, but it introduces new inefficiencies into the mix.

Synthetic Full

Synthetic Full operations were introduced to eliminate the issue of multiple recoveries based on full and incremental data. As the name implies a Synthetic Full operation will take full and incremental protection operations and create a synthetic full from them. This allows the client to perform more incremental backups and have the data protection solution process a synthetic full without the client required.

Traditional Synthetic Full operations shift the processing of the data from the client to perform a full operation to that of the data protection infrastructure. This of course introducers a new burden or challenge, the processing power of the data protection infrastructure itself.

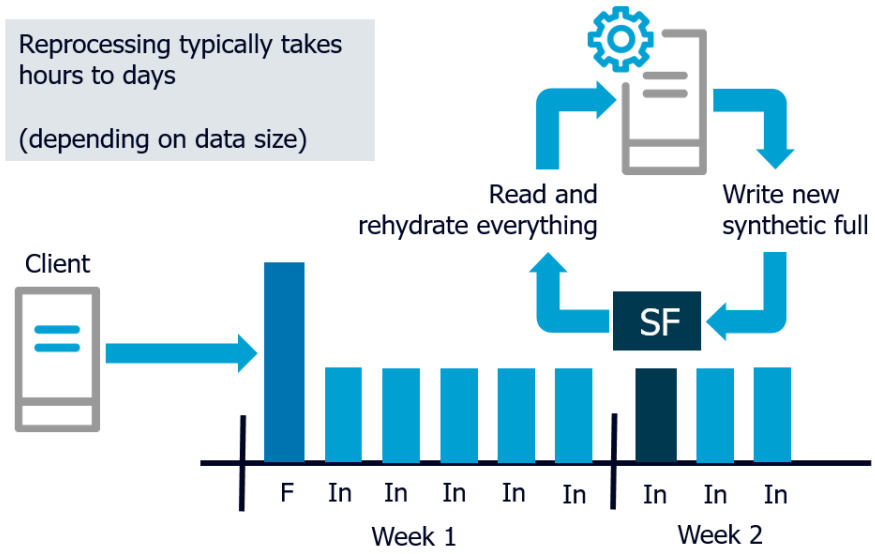

In this example you can see the processing time for the client is reduced after the first full operation and is just sending the incremental data that has changed. This frees up the client to focus on just the changed blocks, however traditional data protection solutions needs to read (and rehydrate) all of the data and write it all down to create a new synthetic full. Data protection solutions are focused on retention and not speed, so therefore this brute force method of processing can be extremely slow. Sometimes it can take more time to process a regular synthetic full than it does to take another full backup.

Although this helps client workload, network usage, space savings, and the need for frequent full backups, it negatively effects the overall RTO because it takes so long to process the data.

Intelligent Processing

There are several elements that can add efficiency into the ability to protect data, it's not just about a singular item, but the solution. The ability to deduplicate data at the storage layer solves the need to store dozens of independent copies of data, which saves storage. Shifting deduplication to the source mitigates constraints on LAN/WAN traffic by only sending the unique data blocks across the wire. While these two parts are critical, they don't represent mitigating the effects of backup data across the entire infrastructure, the burden on the client still exists in this traditional approach to find the data that is unique.

If the client continually must reread all the data to find those unique blocks than the infrastructure that is running the production workloads must utilize critical resources to find the data that needs protection. This is a waste of resources and time to do that.

Commvault's solution utilizes an intelligent processing layer to ensure that there is an understanding of what data is changing and only process that data that has been changed, it eliminates the brute force approach of looking for data that has changed.

Tracking Changes

Tracking changes from the production infrastructure allow the data that has been change to immediately be tracked and identified instead of wasting resource scanning for changes. By tracking the changes that production environments make it creates a map of data that needs to be protected instead of a brute force approach of looking for changes by sifting through all the data.

Some production environments can track their own changes and provide a method to interface to that, for instance VMware's Change Block Tracking, which identifies which blocks have changed since the last operation. Solutions can make use of this by simply querying that API and provide a list of changed data elements. This is great for VMware environments, but there are numerous other production environments and workloads and not all of them can provide a change list for protection. Commvault can hook into existing tracking mechanisms where applicable but can also track changes in workloads that don't support it to identify the changed data. For instance, Commvault has a block level driver that, which provides the ability to track changes at the block level for protection even if the application or workload doesn't have that capability. Additionally, Commvault's Optimized Scan will track the changes on the filesystem so that any file additions, modifications or deletions are tracked therefore protection operations don't have to waste time scanning infrastructure to look for changes.

By intelligently tracking the changes in production data the time that it takes to identify those changes further reduces the time it takes to protect data.

DASH Full – Intelligent Synthetic Full

Synthetic full operations help mitigate the need for clients to perform full backup operations, however it can greatly impact the RPO in a negative fashion and typically utilizes a great number of resources on the backup infrastructure to create these full copies. Commvault's DASH Full technology adds a layer of intelligence to traditional synthetic full operation. This technology coupled with the intelligent change tracking shifts backup operations from that of brute force to that of intelligence. If all the data is tracked while it is protected why does it have to be completely be reprocessed to perform a full recovery?

DASH Full technology combines multiple components that act in unison. No client interaction required for this to work, it is Commvault's default behavior. Commvault's content aware deduplication understands the source data types from the beginning of processing and stores the relevant blocks for protection, but it also stores the metadata and index information of what has changed. Instead of relying on a brute force approach to creating the next full it makes use of the metadata and index information to create a new full without having to process all the blocks. Quite simply we create a new job index and update all the relevant pointers that have the latest information.

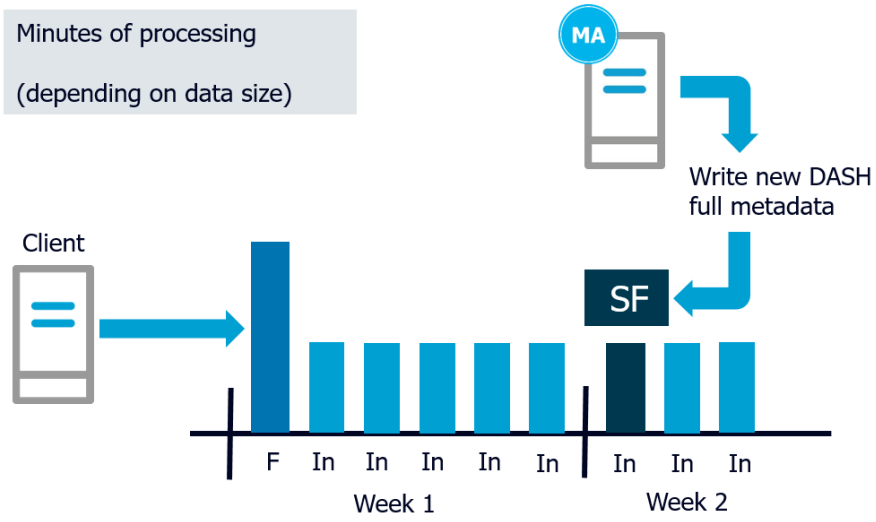

DASH Full technology allows clients to truly shift to an incremental forever approach, after the baseline, and only send the unique blocks of data. The network infrastructure usage is always minimized because the network only must transfer the unique blocks from the client. The space savings is intact due to only storing the unique blocks and although it won't necessarily look like a high deduplication ratio, the storage savings are equivalent to traditional solutions because of the intelligent processing. The processing time is spent on indexes and metadata and not rehydrating and re-deduplicating the same datasets which eliminates the long processing times of traditional synthetic full operations.

This combination of components can drastically shift the protection operations and shrink backup windows significantly because it's not reliant on how fast the client and the network is, but on the intelligence of the solution.

There is one more critical benefit with Commvault's DASH Full technology. Although the testing we undertook with Cisco outlines the performance characteristics of the joint solution there is one more important note to make. Commvault's DASH full technology is automatic, even if the process hasn't been run. Unlike other synthetic full operations that need to be scheduled for full roll-up operations to occur Commvault will automatically do this upon a recovery operation should a DASH full not exist after the last incremental. This stops the need to have DASH full operations run on a daily or even weekly basis and ensures that each incremental copy is essentially always a full copy for recovery.

The shift to intelligent processing has it benefits, on top of these all these benefits it allows us to bend the laws of physics because we aren't bound by regular processing rules.

Backup efficiencies with Deduplication

Where and how the processing is happening plays a large part in the ability to provide efficient use of resources while meeting the requirements of most modern SLAs. Now that we've covered some of the different ways that deduplication technology can increase storage savings, improve network utilization, and drive down client workload, the chart below provides a quick summary.

Processing	Target-Side with Full on Full	Source-Side with Full on Full	Source-Side with Inc. Forever
Deduplication?	At Target	At Client	At Client
Baseline Required?	Yes	Yes	Yes
Ongoing Full?	Yes	Yes	No
Network Impact	Not Efficient	Very Efficient	Very Efficient
Storage Impact	Very Efficient	Very Efficient	Very Efficient
Client Impact	Not Efficient	Not Efficient	Very Efficient
Speed Bound By	Network/Client	Client Processing	Backup Infrastructure

Cisco and Commvault put together a test lab to put a focus on the processing power of our combined solution. Deduplication savings at the storage layer isn't unique but providing an entirely modern approach to data protection that will result in minimal impact across all layers of the impact while providing accelerated performance is a critical component.

Generally, speed is a byproduct of the efficiency in processing the data, and the infrastructure that it is run upon. It's important to understand that not all speed comparisons are the same, and that while there are several performance studies already out there they can sometimes be comparing apples to oranges.

When using a target-side solution the bottleneck is using a combination of the network and client infrastructure sending all the data to the target for processing, typically the constraint is not the target deduplication engine. Shifting to source-side deduplication alleviates the networking challenges and keeps the same efficiency on the storage layer, however several approaches still rely on full on full technology or forcing the client to read additional information because the solution lacks the intelligence to track all the change blocks. This makes the clients the binding factor in those solutions. In terms of bottlenecks, a source-side deduplication solution with full on full technology is still going to be bound by how fast the clients can read the data, which puts additional wear and tear on the production infrastructure.

Lastly with the intelligent data movement that provides incremental forever technology and with intelligent incremental controls the least amount of data is read from the client, the least amount of data is processed, and the networking is utilized in the most efficient fashion. Additionally, since the client and the network aren't the bottlenecks after the baseline operation the speed becomes bound by the infrastructure associated with doing the data protection.

Overview of Lab

The overall testing utilized a server environment design to simulate a real-world dataset for a small-scale test, with one notable exception, the data generate for the tests would all be unique with no data being shared between clients. This would ensure that each client would have to process a unique footprint of data to ensure that any resulting metrics wouldn't benefit from the other systems in the test.

Equipment Utilized

The lab environment was an isolated lab environment based on the following hardware:

- CommServe (Not shown in diagram as it is not in the data-path)
- Cisco C240
- Management point for all jobs.
- MediaAgent
- Cisco UCS S3260 M4
- Dual 40Gbps uplinks
- Clients
- 34 UCS C240 M5 Clients
- 10 Gbps Uplinks per client
- Cisco Nexus 9372PX Switch
- Providing 10Gbps connectivity to the client infrastructure, and 40Gbps uplink capability to other switch segments
- Cisco Nexus 9332PX Switches
- Providing 40Gbps uplink to the S3260 MediaAgent

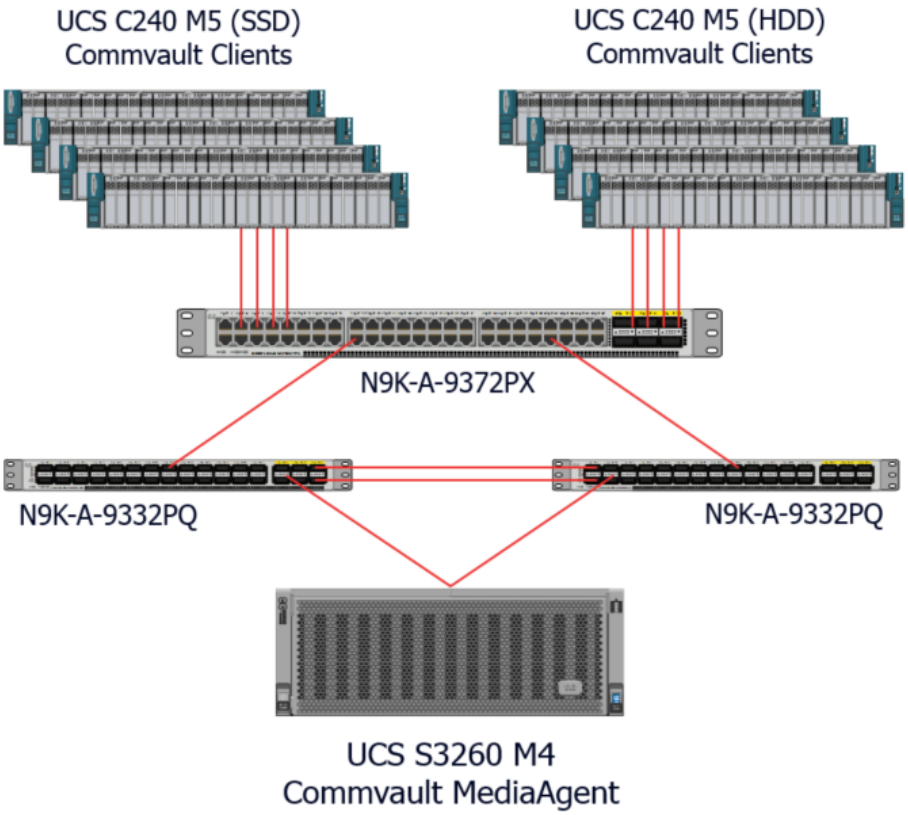

Table 1. Commvault MediaAgent Hardware Configuration

Commvault – Cisco Extra Large MediaAgent	
Cisco UCS Server Type	Cisco UCS S3260
Chipset	Intel C610 Series
CPU Type	2 x Intel Xeon E5-2650v4 Processors
Memory	(128GB) 8 x 16 DDR4 2400MHz RDIMM
Storage Controller	UCS C3000 12Gbps SAS RAID Controller with 4GB Cache
Storage - Boot	2 x 480GB SSD 6Gbps - Raid 1 (480GB Total)*
Storage - Metadata	1 x 6400GB PCIe/NVMe Storage*
Storage – Storage Pool	48 x 6TB NL-SAS -RAID 60 (234TB Total with 3 Hot Spares)*
Network Options	2 x 40Gbps connections

Testing Methodology

Normal customer environments have multiple different workloads that will be protected, these workloads differ greatly from customer to customer as will their data footprint. These variances link directly to the amount of data that will be deduplicated in each environment. Deduplication ratios on their own can be misleading and in some ways are an antiqued view to gauge the success of a modern data protection solution. Being efficient across the entire infrastructure not just at the storage layer and ensuring that the processing of data is handled in an intelligent way instead of using brute force methods.

Production environments will also share data between clients, whether that is via operating system files that have a lot of binaries in common, or copies of data in production environments such as file shares or even database copies for different needs (e.g. dev/test, upgrade testing, etc.). In this test environment though, the focus wasn't on gaining the maximum amount of deduplication ratio, a series of tests could have easily been created to maximize that number artificially.

In a production environment matching data from client to client would be the desired result to further maximize efficiency, not in the test lab.

	Day 1	Days 2-6	Day 7
Week 1	Baseline (Full) + Full on Full	Daily Incremental 1%	Incremental 1% + DASH Full
Week 2	Incremental (1%)	Daily Incremental 1%	Incremental 1% + DASH Full
Week 3	Incremental (1%)	Daily Incremental 1%	Incremental 1% + DASH Full
Week 4	Incremental (1%)	Daily Incremental 1%	Incremental 1% + DASH Full

The above schedule simulates a live customer environment with a 1% daily change rate. The 1% change rate is not related directly to the amount of source data in total, but the amount of data that has changed. In a production environment it could look like there is a much higher daily change rate because of the macro elements of change rates and not the data changing itself. An example of this would be a 10MB file, if a single block changes inside of the file than from a data change rate there is only 1 block that needs to be protected if the rest of the document is the same, however at a macro level that 10MB file has changed. Data generation would change information inside of the files that were generated previously.

The data in this lab was generated via a series of automated workflows that focused on random data being generated in a series of files that range from 1MB to 1GB in size. This ensured that the deduplication engine wouldn't be matching against an extremely high series of individual blocks and there would be no cross matching from client to client.

There were four types of operations performed in the lab:

- Full or Baseline
  - Baseline operations are a singular operation per client where the Commvault MediaAgent has no other protected datasets, in the case of this lab from these clients or any other, so that all data blocks will be transferred and stored for protection.
  - This will typically be the slowest operation because of the transfer of all the data from the clients to the MediaAgent. As this is not a continual event the baseline operation isn't used in the metrics for calculating the long-term speeds of the solution
- Full on Full
  - Shown for comparison purposes only, the full on full protection operation was the only operation that was done with no data change rates.
  - Full on full protection operations are significantly faster than the baseline operations because there would be no data to be transferred. Of course, the faster the clients the faster the full on full operations would be.
  - In the test lab there are two types of clients ones that are on SSD on some that were on slower HDDs, the overall speed differences between these two types of clients represent a more realistic real-world scenario that not all clients will be fast.
- Incremental
- Each incremental operation has a 1% data change rate associated with it.
- This is almost identical to the baseline operation, as only the changed or "unique" blocks of data would be sent across the network.
- DASH Full
  - DASH Full operations provide the outcome of a traditional synthetic full operation without having to reprocess all the data. The previous baseline and incremental operations store all the relevant blocks that need to be protected but also storage the metadata and index information about the changes during each operation. Processing these indexes and metadata components creates a new full protection point without having to re-read and rehydrate all the backed-up data to protect it again.

Performance Numbers

The chart below provides the outputs of this limited scale test, there are some wide ranges in performance numbers as would be expected in the real world. Some of the clients were much slower hard drivers, while other systems were made up of Flash/NVMe drives. This provides a mixture of speeds from specific clients. Additionally, there were only 32 clients that were in use for these tests, the infrastructure of a single Extra-Large MediaAgent can protect typically hundreds of clients.

	Full (Baseline)	Full on Full	Incremental	DASH Full
Operation Count	32 Clients	32 Clients	32 Clients	32 Clients
Client Averages	350 GB/hr.	800 GB/hr.	350 GB/hr.	5,875 GB/hr.
Roll-Up Numbers	11.2 TB/hr.	23.4 TB/hr.	11.2 TB/hr.	188.1 TB/hr.

The baseline on these clients was an average of 11.2TB/hr., or 350GB/hr. per client. The SSD clients were obviously much faster than the hard drive clients, but the roll-up number is the more important element here. The client data was shifted across the network and protected on the MediaAgent. This was the only operation that required the full dataset to be transmitted.

The Full on Full operation was run without the addition of any data, it was 23.4TB/hr. which translates to over a 100% increase in throughput. Since there was no data that was changed at the source, the clients were the bottleneck because they still had to re-read all the information and check for any new unique data to be sent across the network for protection. In this sense the MediaAgent and the network were mostly idle because there was only small lookups to do instead of actual data to be protected.

The incremental datasets changed 1% of the data, across several runs the average throughput was consistent with the baseline. Of course, only about 1% of the original dataset was changed. In an incremental forever world though, that's all the needs to be sent, the changes.

The DASH Full numbers represent 188TB/hr., the averages of course have dramatically increased because the clients themselves don't have to re-read all the data. As outlined in the schedule, the DASH Full in this case was run at the end of each week to roll-up all the changes and represent a new "full" backup for recovery. Even with 32 clients this represents over a 700% increase in protection speed due to the way that the DASH Full operation works.

As noted earlier, but worth repeating, Commvault does not require a DASH full to perform a full recovery operation, it will automatically pull together all the relevant blocks to always perform a recovery operation as required, there is never a need to perform multiple incremental copies to have the most recent full.

All the testing was performed with a standard build of Commvault Complete Backup & Recovery Software, with the patch level of v11 SP11. Standard performance tuning options which are built into Commvault were utilized, there was no custom code or modules in place.

Better Performance with Intelligent Processing

Numbers on datasheets can skew the information in the most favorable way possible, they can utilize extremely fast clients, or clients that are sitting on an all flash array to achieve a compelling number. These tests were done with a mixture of both fast and slow clients, and not done saturating the MediaAgent with hundreds of clients which can make other numbers better. For instance, if there were additional clients added to maximize those throughput numbers all the numbers would benefit from that. It would look significantly faster.

So why would we highlight that fact?

To prove the point that the ultimate speeds and feeds, or brute force of the infrastructure, no matter how fast it is, can benefit from intelligent processing. The Cisco UCS S3260 has dual connectivity into the Nexus 9000 switches at 40Gbps. These were bonded both for bandwidth and availability providing a staggering 80Gbps. But how fast is the raw throughput of 80Gbps?

Network Link	Speed
1Gbps	450 GB/hr.
2 Gbps	900 GB/hr.
3 Gbps	1.35 TB/hr.
5 Gbps	2.25 TB/hr.
10 Gbps	4.5 TB/hr.
25 Gbps	11.25 TB/hr.
40 Gbps	18 TB/hr.
80 Gbps	36 TB/hr.

80Gbps provides a raw throughput of 36TB/hr., by itself that's a truly impressive number, however in terms of overall throughput the DASH Full operations performed at 188.1TB/hr., if that would be required to be transferred over the network that would require about 417Gbps of network connectivity. Of course, even if the network infrastructure was at that speed, the number of other requirements would have to greatly increase also. The MediaAgent would likely have to be on NVMe technology instead of spinning disks, which would greatly increase the cost, and then there's the client infrastructure. If it's only able to send 23TB/hr. today how much faster would the source data have to be to provide that amount of throughput?

It simply wouldn't be realistic to perform operations that way.

The intelligent processing method provides a way to process data at scale without solely relying on the infrastructure itself. The concept of a full backup dates back to days well before disk-based backups were even an idea, back to tape drives. Soon after the tape drives couldn't handle doing full copies each night, which gave way to incremental and full copies. All the technology allows us to move forward into a smarter and more efficient future.

Raw speeds and deduplication ratios can paint a picture, albeit an incomplete one.

Infrastructure is a key element in data protection operations, speed can obviously affect certain characteristics of a modern data protection solution. The Cisco and Commvault solution provides a modern data protection solution that has the infrastructure where necessary to move the data, and the intelligence to know when there is simply a more complete approach.

Conclusion

Ninety percent of the data that will exist in two years does not exist today. The exponential increase in the quantity and density of data being protected over the past few years has made it pertinent for IT managers to focus on the elimination of duplicate information to squeeze out costly redundant storage and still find ways to meet the stringent service level agreements required for achieving recovery time objectives.

The combination the state of the art low latency, lossless, dual 40 Gbps unified network fabric from Cisco UCS with the intelligent processing, tracking with data management and data protection from Commvault provides the most efficient and at the same time the highest performant solution that allows a mechanism for processing data at scale to 188 TB/hr. Recovery time for end-user data requests is reduced from multiple hours in traditional environments to minutes with the joint solution from Cisco UCS and Commvault.

The Cisco | Commvault combination is the only solution available with enterprise-class data management services that takes full advantage of industry-standard infrastructure together with Cisco UCS. Increase your infrastructure flexibility, remove data silos and costly purpose-built appliances and introduce elastic economics for your data – no matter where it lives.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-742356-00 7/19