

# Deduplication Still Matters in Enterprise Clouds as Data Domain and ExaGrid Prove

Technology conversations within enterprises increasingly focus on the “*data center stack*” with an emphasis on cloud enablement. While I agree with this shift in thinking, one can too easily overlook the merits of underlying individual technologies when only considering the “*Big Picture*”. Such is happening with deduplication technology. A key enabler of enterprise archiving, data protection, and disaster recovery solutions, vendors such as Dell EMC and ExaGrid deliver deduplication technology in different ways as DCIG’s most recent 4-page Pocket Analyst Report [reveals](#) that makes each product family better suited for specific use cases.

It seemed for too many years enterprise data centers focused too much on the vendor name on the outside of the box as opposed to what was inside the box – the data and the applications. Granted, part of the reason for their focus on the vendor name is they wanted to demonstrate they had adopted and implemented the best available technologies to secure the data and make it highly available. Further, some of the emerging technologies necessary to deliver a cloud-like experience with the needed availability and performance characteristics did not yet exist, were not yet sufficiently mature, or were not available from the largest vendors.

That situation has changed dramatically. Now the focus is almost entirely on software that provides enterprises with cloud-like experiences that enables them to more easily and efficiently manage their applications and data. While this change is positive, enterprises should not lose sight of the technologies that make up their emerging data center stack as they are not all equally equipped to deliver them in the same

way.

A key example is deduplication. While this technology has existed for years and has become very mature and stable during that time, the options in which enterprises can implement it and the benefits they will realize it vary greatly. The deduplication solutions from Dell EMC Data Domain and ExaGrid illustrate these differences very well.

## **DCIG [Pocket Analyst Report](#) Compares Dell EMC Data Domain and ExaGrid Product Families**

Deduplication systems from both Dell EMC [Data Domain](#) and [ExaGrid](#) have widespread appeal as they expedite backups, increase backup and recovery success rates, and simplify existing backup environments. They also both offer appliances in various physical configurations to meet the specific backup needs of small, midsize, and large enterprises while providing virtual appliances that can run in private clouds, public clouds, or virtualized remote and branch offices.



However, their respective systems also differ in key areas that will impact the overall effectiveness these systems will have in the emerging cloud data stacks that enterprises are putting in place. The six areas in which they differ include:

1. Data center efficiency
2. Deduplication methodology
3. Networking protocols
4. Recoverability
5. Replication
6. Scalability

The most recent 4-page DCIG [Pocket Analyst Report](#) analyzes these six attributes on the systems from these two providers of deduplication systems and compares their underlying features that deliver on these six attributes. Further, this report identifies which product family has the advantage in each area and provides a feature comparison matrix to support these claims.

This [report](#) provides the key insight in a concise manner that enterprises need to make the right choice in deduplication solutions for their emerging cloud data center stack. This [report](#) may be [purchased](#) for \$19.95 at [TechTrove](#), a new third-party site that hosts and makes independently developed analyst content available for sale.

Cloud-like data center stacks that provide application and data availability, mobility, and security are rapidly becoming a reality. But as enterprises adopt these new enterprise clouds, they ignore or overlook technologies such as deduplication that make up these stacks at their own peril as the underlying technologies they implement can directly impact the overall efficiency and effectiveness of the cloud that one is building.

---

# BackupAssist 10.0 Brings Welcomed Flexibility for Cloud Backup to Windows Shops

Today's backup mantra seems to be backup to the cloud or bust! But backup to the cloud is more than just redirecting backup streams from a local file share to a file share presented by a cloud storage provider and clicking the "Start" button. Organizations must examine to which cloud storage providers they can send their data as well as how their backup software packages and sends the data to the cloud. BackupAssist 10.0 answers many of these tough questions about cloud data protection that businesses face while providing them some welcomed flexibility in their choice of cloud storage providers.

Recently I was introduced to [BackupAssist](#), a backup software company that hails from Australia, and had the opportunity to speak with its founder and CEO, [Linus Chang](#), about Backup Assist's 10.0 release. The big news in this release was BackupAssist's introduction of [cloud independent backup](#) that gives organizations the freedom to choose any cloud storage provider to securely store their Windows backup data.

The flexibility to choose from multiple cloud storage providers as a target when doing backup in today's IT environment has become almost a prerequisite. Organizations increasingly want the ability to choose between one or more cloud storage providers for cost and redundancy reasons.

Further, availability, performance, reliability, and support can vary widely by cloud storage provider. These features may

even vary by the region of the country in which an organization resides as large cloud storage providers usually have multiple data centers located in different regions of the country and world. This can result in organizations having very different types of backup and recovery experiences depending upon which cloud storage provider they use and the data center to which they send their data.

These factors and others make it imperative that today's backup software give organizations more freedom of their choice in cloud storage providers which is exactly what BackupAssist 10.0 provides. By giving organizations the freedom to choose from Amazon S3 and Microsoft Azure among others, they can select the "best" cloud storage provider for them. However, since the factors as to what constitute the "best" cloud storage provider can and probably will change over time, BackupAssist 10.0 gives organizations the flexibility to adapt to any changes in conditions at the situation warrants.



**Compress and  
Deduplicate**



**Insulated  
Encryption**



**Resilient  
Transfers**

**Source:BackupAssist**

To ensure organizations experience success when they backup to the cloud, it has also introduced three other cloud-specific features as well, which include:

1. ***Compresses and deduplicates data.*** Capacity usage and network bandwidth consumption are the two primary factors that drive up cloud storage costs. By

introducing compression and deduplication into this release, BackupAssist 10.0 helps organizations better keep these variable costs associated with using cloud storage under control.

2. ***Insulated encryption.*** Every so often stories leak out about how government agencies subpoena cloud providers and ask for the data of their clients. Using this feature, organizations can fully encrypt their backup data to make it inaccessible to anyone.
3. ***Resilient transfers.*** Nothing is worse than having a backup two-thirds to three-quarters complete only to have a hiccup in the network connection or on the server itself that interrupts the backup and forces one to restart the backup from the beginning. Minimally, this is annoying and disruptive to business operations. Over time, restarting backup jobs and resending the same backup data to the cloud can run networking and storage costs. BackupAssist 10.0 ensures that if a backup job gets interrupted, it can resume from the point where it stopped while only sending the required amount of data to complete the backup.

In its 10.0 release, BackupAssist makes needed enhancements to ensure it remains a viable, cost-effective backup solution for businesses wishing to protect their applications running on Windows Server. While these businesses should keep some copies of data on local disk for faster backups and recoveries, the value of efficiently and cost-effectively keeping copies of their data offsite with cloud storage providers cannot be ignored. The 10.0 version of BackupAssist gives them the versatility to store data locally, in the cloud, or both with new flexibility to choose a cloud storage provider at any time that most closely aligns with their business and technical requirements.

---

# DCIG 2016-17 Hybrid Cloud Backup Appliance Buyer's Guide Now Available

DCIG is pleased to announce the availability of the *2016-17 Hybrid Cloud Backup Appliance Buyer's Guide* developed from the backup appliance body of research. Other Buyer's Guides based on this body of research include the recent *DCIG 2016-17 Deduplicating Backup Appliance Buyer's Guide* and the forthcoming *2016-17 Integrated Backup Appliance Buyer's Guide*.

As core business processes become digitized, the ability to keep services online and to rapidly recover from any service interruption becomes a critical need. Given the growth and maturation of cloud services, many organizations are exploring the advantages of storing application data with cloud providers and even recovering applications in the cloud.

Hybrid cloud backup appliances (HCBA) are deduplicating backup appliances that include pre-integrated data protection software *and* integration with at least one cloud-based storage provider. An HCBA's ability to replicate backups to the cloud supports disaster recovery needs and provides essentially infinite storage capacity.

The *DCIG 2016-17 Hybrid Cloud Backup Appliance Buyer's Guide* weights, scores and ranks more than 100 features of twenty-three (23) products from six (6) different providers. Using ranking categories of *Recommended*, *Excellent* and *Good*, this Buyer's Guide offers much of the information an organization should need to make a highly informed decision as to which hybrid cloud backup appliance will suit their needs.

Each backup appliance included in the *DCIG 2016-17 Hybrid Cloud Backup Appliance Buyer's Guide* meets the following criteria:

- Be available as a physical appliance
- May also ship as a virtual appliance
- Includes backup and recovery software that enables seamless integration into an existing infrastructure
- Stores backup data on the appliance via on premise DAS, NAS or SAN-attached storage
- Enables connectivity with at least one cloud-based storage provider for remote backups and long-term retention of backups in a secure/encrypted fashion
- Provides the ability to connect the cloud-based backup images on more than one geographically dispersed appliance
- Be formally announced or generally available for purchase on July 1, 2016

It is within this context that DCIG introduces the *DCIG 2016-17 Hybrid Cloud Backup Appliance Buyer's Guide*. DCIG's succinct analysis provides insight into the state of the hybrid cloud backup appliance marketplace. The Buyer's Guide identifies the specific benefits organizations can expect to achieve using a hybrid cloud backup appliance, and key features organizations should be aware of as they evaluate products. It also provides brief observations about the distinctive features of each product. Ranking tables enable organizations to get an "at-a-glance" overview of the products; while DCIG's standardized one-page data sheets facilitate side-by-side comparisons, assisting organizations to quickly create a short list of products that may meet their requirements.

End users [registering](#) to access this report via the [DCIG Analysis Portal](#) also gain access to the DCIG Interactive Buyer's Guide (IBG). The IBG enables organizations take the next step in the product selection process by generating

custom reports, including comprehensive side-by-side feature comparisons of the products in which the organization is most interested.

By using the DCIG Analysis Portal and applying the hybrid cloud backup appliance criteria to the backup appliance body of research, DCIG analysts were able to quickly create a short list of products that meet these requirements which was then, in turn, used to create this Buyer's Guide Edition. DCIG plans to use this same process to create future Buyer's Guide Editions that further examine the backup appliance marketplace.

Additional information about each buyer's guide edition, including a download link, is available on the [DCIG Buyer's Guides](#) page.

---

## **Data Visualization, Recovery, and Simplicity of Management Emerging as Differentiating Features on Integrated Backup Appliances**

Enterprises now demand higher levels of automation, integration, simplicity, and scalability from every component deployed into their IT infrastructure and the integrated backup appliances found in the DCIG's forthcoming Buyer's Guide Editions that cover integrated backup appliances are a clear output of those expectations. Intended for organizations that want to protect applications and data and then keep it

behind corporate fire walls, these backup appliances come fully equipped from both hardware and software perspectives to do so.

Once largely assembled and configured by either IT staff or value added resellers (VARs), integrated backup appliances have gone mainstream and are available for use in almost any size organization. By bundling together both hardware and software, large enterprises get the turnkey backup appliance solution that was just a few years ago primary reserved for smaller organizations. In so doing, large enterprises can eliminate the need to spend days, weeks, or even months they previously had to spend configuring and deploying these solutions into their infrastructure.

The evidence of the demand for backup appliances at all levels of the enterprise is made plain by the providers who bring them to market. Once the domain of providers such as [STORServer](#) and Unitrends, "software only" companies such as [Commvault](#) and [Veritas](#) have responded to the demand for turnkey backup appliance solutions with both now offering their own backup appliances under their respective brand names.



***Commvault Backup Appliance***



## ***Veritas NetBackup Appliance***

In so doing, any size organization may get any of the most feature rich enterprise backup software solutions on the market, whether it is IBM Tivoli Storage Manager ([STORServer](#)), Commvault ([Commvault](#) and [STORServer](#)), [Unitrends](#) or Veritas [NetBackup](#), delivered to them as a backup appliance. Yet while traditional all-software providers have entered the backup appliance market, behind the scenes new business demands are driving further changes on backup appliances that organizations should consider as they contemplate future backup appliance acquisitions.

- ***First, organizations expect successful recoveries.*** A few years ago, the concept of all backup jobs completing successfully was enough to keep everyone happy and giving high-fives to one another. No more. Organizations recognize that they have reliable backups residing on a backup appliance and these appliances may largely sit idle during off-backup hours. This gives the enterprise some freedom to do more with these backup appliances during these periods of time such as testing recoveries, recovering applications on the appliance itself, or even presenting these backup copies of data to other applications to use as sources for internal testing and development. DCIG found that a large number of backup appliances support one or more vCenter Instant Recovery features and the emerging crop of backup appliances can also host virtual machines and recover applications on them.
- ***Second, organizations want greater visibility into their data to justify business decisions.*** The amount of data residing in enterprise backup repositories is staggering. Yet the **lack** of value that organizations derive from that stored data combined with the potential risk it presents to them by retaining it is equally staggering. Features that provide greater visibility

into the metadata of these backups which then analyze it and help turn it into measurable value for the business are already starting to find their way onto these appliances. Expect these features to become more prevalent in the years to come.

- ***Third, enterprises want backup appliances to expand their value proposition.*** Backup appliances are already easy to deploy but maintaining and upgrading them over time or deploying them for other use cases gets more complicated over time. To address these concerns, emerging providers such as [Cohesity](#), which is making its first appearance in DCIG Buyer's Guides as an integrated backup appliance, directly addresses these concerns. Available as a scale-out backup appliance that can function as a hybrid cloud backup appliance, a deduplicating backup appliance and/or as an integrated backup appliance, it provides an example of how an enterprises can more easily scale and maintain it over time while giving them the flexibility to use it internally in multiple different ways.

The forthcoming DCIG 2016-17 Integrated Backup Appliance Buyer's Guide Editions highlight the most robust and feature rich integrated backup appliances available on the market today. As such, organizations should consider these backup appliances covered in these Buyer's Guides as having many of the features they need to protect both their physical and virtual environments. Further, a number of these appliances give them early access to the set of features that will position them to meet their next set of recovery challenge,s satisfy rising expectations for visibility into their corporate data, and simplify their ongoing management so they may derive additional value from it.

---

# DCIG 2016-17 Deduplicating Backup Appliance Buyer's Guide Editions Now Available

DCIG is pleased to announce the availability of the following DCIG 2016-17 Deduplicating Backup Appliance Buyer's Guide Editions developed from the backup appliance body of research. Other Buyer's Guide Editions based on this body of research will be published in the coming weeks and months, including the 2016-17 Integrated Backup Appliance Buyer's Guide and 2016-17 Hybrid Cloud Backup Appliance Buyer's Guide Editions.

Buyer's Guide Editions being released on September 20, 2016:

- *DCIG 2016-17 Sub-\$100K Deduplicating Backup Appliance Buyer's Guide*
- *DCIG 2016-17 Sub-\$75K Deduplicating Backup Appliance Buyer's Guide*
- *DCIG 2016-17 Sub-\$50K Deduplicating Backup Appliance Buyer's Guide*
- *DCIG 2016-17 US Enterprise Deduplicating Backup Appliance Buyer's Guide*

*DCIG 2016-17 Deduplicating Backup Appliance Buyer's Guide* Edition had to meet the following criteria:

- Be intended for the deduplication of backup data, primarily target-based deduplication
- Includes an NAS (network attached storage) interface
- Supports CIFS (Common Internet File System) or NFS (Network File System) protocols
- Supports a minimum of two (2) hard disk drives and/or a minimum raw capacity of eight terabytes
- Be formally announced or generally available for purchase on July 1, 2016

The various Deduplicating Backup Appliance Buyer's Guide Editions are based on at least one additional criterion, whether list price (Sub-\$100K, Sub-\$75K and Sub-\$50K) or being from a US-based provider.

By using the DCIG Analysis Portal and applying these criteria to its body of research into backup appliances, DCIG analysts were able to quickly create a short list of products that meet these requirements which was then, in turn, used to create Buyer's Guide Editions to publish and release. DCIG plans to use this same process to create future Buyer's Guide Editions that examine hybrid cloud and integrated backup appliances among others.

End users registering to access any of these reports via the [DCIG Analysis Portal](#) also gain access to the DCIG Interactive Buyer's Guide (IBG). The IBG enables organizations take the next step in the product selection process by generating custom reports, including comprehensive side-by-side feature comparisons of the products in which the organization is most interested.

Additional information about each buyer's guide edition, including a download link, is available on the [DCIG Buyer's Guides](#) page.

---

## **Dell NetVault and vRanger are Alive and Kicking; Interview with Dell's Michael Grant,**

# Part 3

Every now and then I hear rumors in the market place that the only backup software product that Dell puts any investment into is Dell Data Protection | Rapid Recovery while it lets NetVault and vRanger wither on the vine. Nothing could be further from the truth. In this third and final part of my interview series with Michael Grant, director of data protection product marketing for Dell's systems and information management group, he refutes those rumors and illustrates how both the NetVault and vRanger products are alive and kicking within Dell's software portfolio.

**Jerome:** *Can you talk about the newest release of NetVault?*

**Michael:** Dell Data Protection | [NetVault Backup](#), as we now call it, continues to be an important part of our portfolio, especially if you are an enterprise shop that protects more than Linux, Windows and VMware. If you have a heterogeneous, cross-platform environment, NetVault does the job incredibly effectively and at a very good price. Netvault development keeps up with all the revs of the various operating systems. This is not a small list of to-dos. Every time anybody revs anything, we rev additional agents and provide updates to support them.



*Source: Dell*

In this current rev we also improved the speed and performance of NetVault. We now have a protocol accelerator, so we can keep less data on the wire. Within the media server

itself, we also had to improve the speed and we wanted to address more clients. Customers protect 1,000's of clients using NetVault and they want to add even more than that. To accommodate them, we automate the installation so that it's effective, easily scalable and not a burden to the administrator.

To speed up protection of the file system, we put multi-stream capability into the product, so one can break up bigger backup images into smaller images and then simultaneously stream those to the target of your choice. Obviously, we love to talk to organizations about putting the DR deduplication appliances in as that target, but because we believe in giving customers flexibility and choice, you can multi-stream to just about any target.

Re-startable VMware backup is another big pain point for a lot of our customers.. They really bent our development team's ear and said, *"Listen, going back and restarting the backup of an entire VMDK file is a pain if it doesn't complete. You guys need to put an automatic restart in the product."*

Think about watching a show on DVR. If you did not make it all the way through the show in the first sitting, you don't want to have to go back to the beginning and re-watch the entire thing the next time you watch it. You want to pick up where you left off.

Well, we actually put similar capability in NetVault. We can restart the VM backup from wherever the backup ended. Then you can just pick back up knowing that you have the last decently mountable restore point at a point in time when it trailed off. Just restart the VM and get the whole job done. That cuts hours out of your day if you did not get a full backup of a VM. .

Sadly, backing up VMDK files, particularly in a dynamic environment, can be a real challenge. It is not unusual to

have one fail midway through the job or not have a full job when you go to look in the queue. Restarting that VM backup just made a lot of sense for the IT teams.

Those new features really highlight what is new in the NetVault 11 release that we just announced. Later in the first half of this year, you will see the accompanying updates to the agents for NetVault 11 so that we remain in sync with the latest releases from everybody from Oracle through Citrix and VMware, as well as any other agents that need to be updated to align with this NetVault 11 release.

**Jerome:** *Are the functionality of vRanger and AppAssure now being folded under the Rapid Recovery brand?*

**Michael:** That's a little too far. We are blending the technologies, to be sure. But we are still very much investing in [vRanger](#) and it remains a very active part of our portfolio. To quote the famous Mark Twain line, "*the tales of vRanger's death are greatly exaggerated.*"



*Source: Dell*

We are still investing in it and it's still very popular with customers. In fact, we made an aggressive price change in the fall to combine vRanger Pro with the standard vRanger offering. We just rolled in three years of service and made it all vRanger Pro. Then we dropped the price point down several hundred dollars, so that's it less than any of the other entry

level price points for virtualized backup in the industry. We will continue to invest in that product for dynamic virtual environments.

So, yes, you will absolutely still see it as a standalone product. However, even with that being the case, there is no reason that we should not reach in there and get some amazing code and start to meld that with Rapid Recovery. As DCIG has pointed out in its research and, as our customers tell us frequently, they would like to have as few backup tools in their arsenal as possible, so we will continue to blend those products to simplify data protection for our customers. The bottom line for us is, wherever the customer wants to go, we can meet them there with a solution that fits.

**Jerome:** *How are you positioning each of these three products in terms of market segment?*

**Michael:** I do want to emphasize that we focus very much on the midmarket. We define midmarket as 500 to 5,000 employees. When we took a look at who really buys these products, we found that 90 plus percent of our solutions are being deployed by midmarket firms. The technologies that we have just talked about are well aligned to that market, and that makes them pretty unique. The midmarket is largely under served when it comes to IT solutions in general, but especially when it comes to backup and recovery. We are focusing on filling a need that has gone unfilled for too long.

*In [Part 1](#) of this interview series, Michael shares some details on the latest features available in Dell's data protection line and why organizations are laser-focused on recovery like never before.*

*In [Part 2](#) of this interview series, Michael elaborates upon how the latest features available in Dell's data protection line enable organizations to meet the shrinking SLAs associated with these new recovery objectives.*

---

# HP StoreOnce Deduplicating Backup Appliances Put Organizations on Path to Ending Big Data Backup Headaches

During the recent HP Deep Dive Analyst Event in its Fremont, CA, offices, HP shared some notable insights into the percentage of backup jobs that complete successfully (*and unsuccessfully*) within end-user organizations. Among its observations using the anonymized data gathered from hundreds of backup assessments at end-user organizations of all sizes, HP found that over 60% of them had backup job success rates of 98% or lower, with 12% of organizations showing backup success rates of lower than 90%. Yet what is more noteworthy is through HP's use of Big Data analytics, it has identified large backups (*those that take more than 12 hours to complete*) as being the primary contributor to the backup headaches that organizations still experience.

About once every nine (9) months (*give-or-take*) HP invites storage analysts to either its Andover, MA, or Fremont, CA, offices to have a series of in-depth discussion about its portfolio of products in its Storage division. During these 2-day events, the product managers from the various groups (3PAR StoreServ, [StoreOnce Backup](#), StoreAll Archive, StoreVirtual, etc.) are given time to present to the analysts in attendance. It is during these times that candid and frank discussions ensue where each HP product is examined in-depth with the HP product managers providing context as to why they made the

product design decisions that they have.

One of the more enlightening pieces of information to come out of these sessions was the amount of data that HP has collected from organizations into which its StoreOnce appliances are being considered for deployment. To date, HP has assessed environments with more than **half an exabyte** of backup data with the vast majority of backup data analyzed comprised of file system backups, either performed directly or thru NDMP.

This amount of data gives HP a rather unique perspective on backup successes and failures. For instance, HP shared that of the approximately 4.5 million backup jobs for which it has collected data, 94.7% of them have completed successfully.

HP also revealed that organizations in particular struggle with long-running backups. Over 50% of the assessed environments had backup windows of 24 hours or more. Of these, 30% of the organizations that it had assessed had at least one backup that ran in excess of 192 or more hours – or 8 days or more. Further, the data indicates a correlation between file system backups and long backup windows.

Granted, these statistics from HP are by no means “official” and subject to some interpretation. However they possibly provide some of the first, large scale empirical evidence that for the vast majority of organizations that data growth goes hand-in-hand with elongated backup windows and is a major contributor if not the primary source of why backups still fail today.

Organizations moving to StoreOnce appliances, which provide high levels of performance in conjunction with source-side deduplication, are addressing this common organizational pain point as they both shorten backup windows and increase the probability that backups complete successfully. Further, using HP’s StoreOnce [Recovery Manager Central](#) solution, organizations may perform virtual machine and file system

backups based on block level changes as backup data flows from HP 3PAR StoreServ to StoreOnce. This combination of solutions provides the keys that organizations need to solve backup in their environments as many organizations using the HP StoreOnce deduplicating backup appliances have already discovered.

---

## **Advanced Encryption and VTL Features Give Organizations New Impetus to Use the Dell DR Series as their “One Stop Shop” Backup Target**

To simplify their backup environments, organizations desire backup solutions that essentially function as “one-stop shops” to satisfy their multiple backup requirements. To succeed in this role, they should provide needed software, offer NAS and virtual tape library (VTL) interfaces, scale to high capacities and deliver advanced encryption capabilities to secure backup data. By Dell introducing advanced encryption and VTL options into its latest DR 3.2 OS software release for its [DR Series](#), it delivers this “one-stop shop” experience that organizations want to implement in their backup infrastructure.

### **The More Backup Changes, the More It Stays the Same**

Deduplicating backup appliances have replaced tape as a backup target in many organizations. By accelerating backups and restores, increasing backup success rates and making disk-

based backup economical, these appliances have fundamentally transformed backup.

Yet their introduction does not always change the underlying backup processes. Backup jobs may still occur daily; are configured as differential, incremental or full; and, are managed centrally. The only real change is using disk in lieu of tape as a target.

Even once in place, many organizations still move backup data to tape for long term data retention and/or offsite disaster recovery. Further, organizations in the finance, government and healthcare sectors typically encrypt data such as SEC [Rule 17a-4](#) specifies or the 2003 [HIPAA Security Rules](#) and more recent [2009 HITECH Act](#) strongly encourage.

### **Continued Relevance of Encryption and VTLs in Enterprises**

This continued widespread use of tape as a final resting place for backup data leads organizations to keep current backup processes in place. While they want to use deduplicating backup appliances, they simply want to swap out existing tape libraries for these solutions. This has given rise to the need for deduplicating backup appliances to emulate physical tape libraries as virtual tape libraries (VTLs).

A VTL requires minimal to no changes to existing backup-to-tape processes nor does it require many changes to how the backup data is managed after backup. The backup software now backs up data to the VTL's virtual tape drives where the data is stored on virtual tape cartridges. Storing data this way facilitates its movement from virtual to real or physical tape cartridges and enables the backup software to track its location regardless of where it resides.

VTLs also accelerate backups. They give organizations more flexibility to keep data on existing SANs which negates the need to send data over corporate LANs where it has to contend with other network traffic. SAN protocols also better support

the movement of larger block sizes of data which are used during backup.

Finally, VTLs free backup from the constraints of physical tape libraries. Creating new tape drives and tape cartridges on a VTL may be done with the click of a button. In this way organizations may quickly create multiple new backup targets to facilitate scheduling multiple, concurrent backup jobs.

Encrypting backup data is also of greater concern to organizations as data breaches occur both inside and outside of corporate firewalls. This behooves organizations to encrypt backup data in the most secure manner regardless if the data resides on disk or tape.

### **Advanced Encryption and VTL Functionality Central to Dell DR Series 3.2 OS Release**

Advanced encryption capabilities and VTL functionality are two new features central to Dell's 3.2 operating system (OS) release for its [DR Series](#) of deduplicating backup appliances. The 3.2 OS release provides organizations a key advantage over competitive solutions as Dell makes all of its software features available without requiring additional licensing fees. This applies to both new DR Series appliances as well as existing Dell DR Series appliances which may be upgraded to this release to gain full access to these features at no extra cost.

The 3.2 OS release's advanced encryption capabilities use the FIPS 140-2 compliant 256-bit Advanced Encryption Standard (AES) standard to encrypt data. By encrypting data that conforms to this standard ensures that it is acceptable to federal agencies in both Canada and the United States. This also means that organizations who are in these countries and need to comply with their regulations are typically, by extension, in compliance when they use the DR Series to encrypt their backup data.

The 3.2 OS release implements this advanced encryption capability by encrypting data after its inline deduplication of the backup data is complete. In this way, each DR Series appliance running the 3.2 OS release deduplicates backup data as it is ingested to achieve the highest possible deduplication ratio as encrypting data prior to deduplication negatively impacts deduplication's effectiveness. Encrypting the data after it is deduplicated also reduces the amount of overhead associated with encryption since there is less data to encrypt while keeping the overhead associated with the encryption on the DR Series appliance. In cases where existing DR4100s are upgraded to the 3.2 OS release, encryption may be done post-process on those data volumes that have previously been stored unencrypted in the DR4100's storage repository.

The VTL functionality that is part of the 3.2 OS release includes options to present a VTL interface on either corporate LANs or SANs. If connected to a corporate LAN, the NDMP protocol is used to send data to the DR Series while, if it is connected to a corporate SAN, the iSCSI protocol is used.

Every DR Series appliance running the 3.2 OS release may be configured to present up to four (4) containers that each operate as separate VTLs. Each of these individual VTL containers may emulate one (1) StorageTek STK L700 tape library or an OEM version of the STK L700; up to ten (10) IBM ULT3580-TD4 tape drives; and, up to 10,000 tape cartridges that may each range in size from 10GB to 800GB.

As each individual VTL container on the DR Series appears as an STK L700 library to backup software, the backup software manages the VTL in the same way it does a physical tape library: it copies the data residing on virtual tape cartridges to physical tape cartridges and back again, if necessary. With this functionality available on leading enterprise backup software products such as Dell NetVault, CommVault Simpana, EMC Networker, IBM TSM, Microsoft Data

Protection Manager (iSCSI only), Symantec Backup Exec and Symantec NetBackup, each of these can recognize and manage the Dell DR Series VTL as a physical STK L700 tape library, carry forward existing tape copy processes, implement new ones if required, and manage where copies of tape cartridges—physical or virtual—reside.

## **Dell's 3.2 OS Release Gives Organizations New Impetus to Make Dell DR Series Their "One Stop Shop" Backup Target**

All size organizations want to consolidate and simplify their backup environments and using a common deduplicating backup appliance platform is one excellent way to do so. Dell's 3.2 OS release for its DR Series gives organizations new impetus to start down that path. The introduction of advanced encryption and VTL features along with the introduction of 6TB HDDs on expansion shelves for the DR6000 and the availability of Rapid NFS/Rapid CIFS protocol accelerators for the DR4100 provide the additional motivation that organizations need to non-disruptively introduce and use the DR Series in this broader role to improve their backup environments even as they keep existing backup processes in place.

---

## **Four New Features that Will Keep Backup Software "Sticky"**

Backup software has traditionally been one of the stickiest products in organizations of all sizes in part because it has been so painful to deploy and maintain. After all, once it was installed and sort of working, no organization wanted to subject itself to that torture again. But in recent years as backup has become easier to install and maintain, swapping it

out for another or consolidating multiple backup software solutions down to single one has become much more palatable. This puts new impetus on backup software providers to introduce new features into their products to keep them relevant and “sticky” in their customer environments longer term.

Since the start of 2015 DCIG has been doing background research in anticipation of releasing its first ever [Buyer's Guide](#) on the topic of Hybrid Cloud Backup Appliances. In doing so, it is seeing many new features in the backup software that ships with these appliances that were never part of backup software in the past because these providers were still trying to get their software to successfully backup organizational applications and data.

The good news is that backup has largely been solved. In conferences and analyst summits that I have attended in recent weeks where end users have presented coupled with my own conversations with end users, most openly say their backups challenges are 99% solved. This is resulting in their organizations devoting fewer or even no IT staff to managing backup and re-allocating those individuals to solve more pressing, strategic initiatives in their respective organizations.

The bad news, at least for backup providers, is that solving backup is turning into a bit of mixed blessing. While they are grateful that their backup software now works in their customers' environments and the number of support calls they receive is declining, they also face the “*squeaky wheel gets the oil*” dilemma. Since their backup software no longer “squeaks” in customer deployments, customers start to view it as a commodity solution that works and can be easily replaced by a competing product at a lower cost (*which I already see evidence of occurring.*)

This puts the onus on backup software providers to introduce

new features that will keep their backup software “sticky” by having specific features that add measurable and quantifiable value to the organization over competing products. While the industry remains in the early stages of this transformation, there are four new features that backup software minimally needs to offer going forward for it to remain relevant and keep it “sticky” going forward:

- **Connectivity to multiple cloud providers.** Nearly every organization is looking to store some of its data with public cloud storage providers. While organizations are still in the early stages of moving data off-site, small and mid-sized businesses and enterprises are moving more quickly than larger organizations to adopt cloud connectivity as they are less likely to have a secondary site to store their backup data. Based on early results, organizations are most apt to want connectivity to one or more of the following cloud providers: Amazon Web Services (AWS), Microsoft Azure, Rackspace and Google. Also, organizations should want backup software that supports multiple cloud providers so they have the flexibility to move from one to another should cost reduction or feature functionality justify such a change.
- **Data copy management.** Organizations no longer just want one or more copies of their data simply residing in a repository in anticipation of a recovery that they hope they never have to perform. They want to use the copies of data for other purposes such as testing, development or doing Big Data analysis. To accomplish this, backup software has to store copies of data in a format that may be easily accessed and used by other applications or the copy recovered without needing to use the backup software to recover the data.
- **File sync and share.** File sync and share is a relatively new feature that is already from providers such as Acronis ([Acronis Access Advanced](#)). Putting this feature

in backup software capitalizes on the software footprint that backup software often has on many PCs and servers in the enterprise and utilizes backup software's native ability to copy and replicate data. Further, many organizations would like to move away from file sync and share options such as Dropbox because of the inherent security risks they present. More organizations see backup software as a means to securely deliver file sync and share feature functionality to their users.

- ***Recovery in the cloud.*** Getting data into the cloud is great. However recovering one's applications or even an entire data center is the new end game because if one's existing data center goes away, having backups offsite with no place or means to restore them is pretty much worthless. Being able to recover data, applications or even data centers with a cloud provider and orchestrating the management of those recoveries through the backup software will help to make that backup software almost indispensable to organizations.

---

## **A Single Backup Solution for Today's Multiple Backup and Recovery Challenges; Interview with Dell Software's General Manager,**

# Data Protection, Brett Roscoe

## Part VIII

One of the largest challenges facing enterprises today in respect to backup and recovery is successfully meeting all of the different backup and recovery requirements associated with each application. Physical backups, virtual backups, instant recoveries, application-specific backup requirements and much more make successfully executing upon a comprehensive backup and recovery strategy more difficult than ever before. In this eighth installment of my interview series with Brett Roscoe, General Manager, Data Protection for Dell Software, he shares how Dell has brought together its various data protection products into one backup and disaster recovery suite to make it easier to customers to address these challenges with a single solution.

***Jerome:** Can you discuss this emerging trend in the data protection industry for providers to bundle together different but complementary backup and software together in a single product suite. In fact Dell software recently announced the launch of the Dell Backup and Disaster Recovery Suite. Can you talk about this new suite and how it might benefit customers?*

***Brett:** Absolutely. I'm really excited about the suite. It accomplishes quite a few things for our customers. Most importantly, it allows customers to use and leverage all of the Dell data protection IP with one simple licensing model. This is a great story just from a customer perspective, and that's before we even finish all of the exciting integration projects we currently have in development*

With the Dell Backup & Disaster Recovery Suite, customers have the freedom to leverage the best tool set for whatever their application is or whatever portions of their environment they want or need to protect. You may have a team that's very

focused on virtualization and vRanger is a great fit into that environment. You may have a critical application that you feel like you can have no more than five minutes of down time, in which case, AppAssure can come in and help you build a solution there. You may have traditional, file-based, cross platform protection needs, in which case, NetVault is an outstanding choice. With one license, you get the freedom to mix and match these technologies based on your specific needs. That to me is a great story.

As I look across the industry, I don't know of any vendor that has the broad portfolio capabilities that we do, much less the ability to give customers access to that entire portfolio through a single license.

Not only are we giving you all of the capabilities you need, but we've simplified the purchase by offering a single capacity based license that gives you that broad portfolio capability. You do not have to choose. You do not have to be locked into a certain product. In fact with our portfolio , you can even change your implementation over time without changing your license.

Maybe you start out with a primarily physical environment that has one set of requirements. Then you move to a more virtual or cloud-based environment over time. As your RPO/RT0 requirements shift, you can reconfigure the Dell product set that you're using in order to provide the best fit and value for you changing needs.

There is a lot of flexibility there, but I want to be clear that this is just the start. We have a robust integration road map. We are still doing all the cool things on the development side to make it easy as possible for customers to use all of the IP, but the Dell Backup & Disaster Recovery Suite allows customers to take advantage of all of our capabilities today.

I think the suite is a great value as it can allow a customer

to grow up with us. A customer can start as a small business with one portion of a portfolio and grow larger while having access to a more comprehensive portfolio without any disruption or major forklift upgrades.

***Jerome:** Sounds like some pretty exciting times for Dell. What's the general morale at Dell in terms of where you are at and where you are going with all this?*

***Brett:** I'll tell you, I feel very fortunate to be at Dell, and to be involved with our data protection business specifically. I feel like it is just one of the fun areas right now. Data protection has traditionally been seen as a form of insurance, but the landscape in data protection is changing, and customers are using our products in new ways and finding ways to reduce risk, and it is great to be a part of that change.*

You are always going to *"have to have it"*, but new features and capabilities can often change the way customers use or leverage our products and free up resources for our customers to invest in other areas.

Also, customers feel like data protection is becoming a more critical part of the environment. They know they got to have it. But they also see these new capabilities and these changing tool sets, and how they can now show the value of data protection to their customers and management teams, and free up resources to work on other projects. It is fun to see some of the testimonials we get from our customers, how they are using our products, and the cool things they are doing with it.

Personally, I am very bullish about data protection portion at Dell. I love my job and cannot imagine doing anything else right now. We are just having a lot of fun.

*In [Part I](#) of this interview series, Brett and I discussed the biggest backup and recovery challenges that organizations face today.*

In [Part II](#) of this interview series, Brett and I discussed the imperative to move ahead with next gen backup and recovery tools.

In [Part III](#) of this interview series, Brett and I discussed four (4) best practices that companies should be implementing now to align the new capabilities in next gen backup and recovery tools with internal business processes.

In [Part IV](#) of this interview series, Brett and I discussed the main technologies in which customers are currently expressing the most interest.

In [Part V](#) of this interview series, Brett and I examine whether or not one backup software product can “do it all” from a backup and recovery perspective.

In [Part VI](#) of this interview series, Brett and I discuss Dell’s growing role as a software provider.

In [Part VII](#) of this interview series, Brett provides an in-depth look at Dell’s new Backup and Disaster Recovery Suite.

In [Part IX](#) of this interview series, Brett shares his thoughts as to what he sees as the future of data protection over the next decade.

---

**An In-Depth Look at the Dell  
Data Protection Portfolio;  
Interview with Dell  
Software’s General Manager,  
Data Protection, Brett**

# Roscoe, Part VII

Backup and recovery used to generate as much interest among IT as watching paint dry. But with almost all organizations expecting near-24x7 uptime from all of their applications all of the time and potentially anywhere, that perspective has changed. Agentless backups, disaster recovery and instant recovery features found on backup software have the attention of IT like never before. In this seventh installment of my interview series with Brett Roscoe, General Manager, Data Protection for Dell Software, we take an in-depth look at Dell's data protection portfolio and how it maps to these pressing backup and recovery concerns of IT managers today.

**Jerome:** You have talked about Dell's growing reputation as a software provider. Please talk about how its data protection products as part of Dell's overall software portfolio and what they formally bring to the market.

**Brett:** Absolutely. First thing people need to understand is that we are very focused on integration. We are very focused on delivering an experience whereby no matter what product brings you into the family of Dell data protection customers, you will benefit from the IP that we have across the entire portfolio.

That's a key point. We do not want to keep these products as standalone technologies. We are working very hard to provide the capabilities in each of these products or the advantages and value propositions in each of these products across the portfolio. Having said that, let me quickly talk about where the portfolio came from, and what are all the different pieces of IP that we have developed or acquired, and how they fit together.

The first piece of IP was an [acquisition](#) called Ocarina. Ocarina was a leading deduplication and compression technology

company. At the time we acquired them, their big focus was actually in the primary storage market around vertical markets like imaging and video. Their IP is really very high horsepower kind of stuff that works well against any number of data sets.

The fact that we focused this business on backup and recovery really speaks to the need and the real value that deduplication and compression bring to the backup and recovery market.

Ocarina is certainly an area that we are investing in and you are seeing that technology come to market in the form of the DR line of backup and deduplication appliances. You also see it in our NetVault product and will see it in other places in our portfolio as time goes by.

The next one is AppAssure. AppAssure was an [acquisition](#) designed to meet next generation backup and recovery capabilities. It is our application consistent technology that allows customers to have five minute RPOs and RTOs in minutes, leveraging features like virtual standby, live recovery, and change block tracking. It's the kind of technology that really provides that very high performance recovery capability for customers..

The next product is vRanger, which is our leading product for agentless backup of VMware ESX and Microsoft Hyper V. It is designed to meet the needs of the virtualization IT administrator who is very centered around the VMware or Hyper-V environments and management tools.

This individual can really leverage the vRanger product because we an integrated and focused approach on that ecosystem. We provide agentless backup of VMware and HyperV, we provide plug-ins, and we can work within the VMware and Hyper-V toolset. Our look and feel is very much like the Hyper V and VMware products, so customers who are used to those

hypervisor management tools get up and running very quickly with vRanger.

Then there's NetVault. NetVault is our product that, in terms of OS and application support, has the broadest portfolio support of any of our products. It comes from a more traditional backup and recovery product background, but it's one we are heavily investing in order to ensure it evolves to continue meeting the needs of the modern customer... Over time, you'll see NetVault as a great example of Dell leveraging capabilities from other parts of our portfolio to enhance existing offerings for customers.

NetVault has been around for a long time and was part of the Quest [acquisition](#). It continues to be a very popular product among customers who are looking to augment or maybe centralize their data protection environment from multiple Independent Software Vendors (ISVs), where maybe one piece of backup software is supporting one OS and another is supporting another OS. You can consolidate them on NetVault and meet all of your application and OS protection requirements with one tool.

Each of these products was acquired at a different time, but there is a lot of history in terms of how these products came about. Almost all of them came up through startups, from people thinking about how to be disruptive and create unique capabilities. If you look at our portfolio, I believe we have the youngest, most IP-rich portfolio in the industry. Now we're focusing on integrating and provide as much value as we can to customers. But you can't integrate great technologies unless you have great technologies to begin with, so I'm very excited that we have these tools in our toolset in order to make that initiative successful.

*In [Part I](#) of this interview series, Brett and I discussed the biggest backup and recovery challenges that organizations face today.*

In [Part II](#) of this interview series, Brett and I discussed the imperative to move ahead with next gen backup and recovery tools.

In [Part III](#) of this interview series, Brett and I discussed four (4) best practices that companies should be implementing now to align the new capabilities in next gen backup and recovery tools with internal business processes.

In [Part IV](#) of this interview series, Brett and I discussed the main technologies in which customers are currently expressing the most interest.

In [Part V](#) of this interview series, Brett and I examine whether or not one backup software product can “do it all” from a backup and recovery perspective.

In [Part VI](#) of this interview series, Brett and I discuss Dell’s growing role as a software provider.

In [Part VIII](#) of this interview series, Brett and I discuss the trend of vendors bundling different but complementary data protection products together in a single product suite.

---

## **Answering The Question of Whether One Backup Product Can Do It All; Interview with Dell Software’s General Manager, Data Protection, Brett Roscoe Part V**

Data protection has evolved well beyond the point where one can backup and recover data doing once a day backups. Continuous data protection, array-based snapshots,

asynchronous replication, high availability, disaster recovery, backup and recovery in the cloud and long term backup retention are now all part of managing backup.

However, the real question becomes, *“Can one product even manage all of these different facets of backup and recovery? Or should a backup solution even try to accomplish this feat?”* In this fifth installment of my interview series with Brett Roscoe, General Manager, Data Protection for Dell Software, we discuss this very important question of whether one backup product can do it all in today’s data center.

**Jerome:** There are a lot of demands being placed on backup and recovery software these days so the question I have you is this, can one backup and software product still do it all to meet these different customer demands? If so, why? If not, why not?

**Brett:** That’s a great question and it is hard to provide a yes or no answer to that question. Due to the rapid pace of change in IT, we see lots of variables that are changing the landscape including software defined data center, container based application rollout and the ongoing trend of virtualization and cloud adoption. As a result, customers’ requirements for data protection are changing, and that in turn is changing what they look for and need from data protection vendors like [Dell](#) and others.

Given that the needs of customers are rapidly evolving, we as a company spend a lot of time working to make sure we provide the new technologies and unique capabilities that can help them meet those needs. That’s one of the core things Dell drives for with every decision we make. As a general manager, I need to make sure that my development teams are constantly working to ensure that our technologies keep up with the changing marketplace.

To tie it back to the initial question, there are certainly

ways to consolidate and simplify data protection and disaster recovery. So, for example, I talked about our DR line of target-based disk backup and deduplicating appliances. Those products today not only work seamlessly with the other products in our portfolio, but they can also work with all other backup products a customer might already have in their environment.

If customers want to look for way to consolidate technology, the DR series is a great place to start. The DR products are designed to run in a heterogeneous environment with all applications, any OS, and all backup software. But there are certainly advantages to start consolidating in some of those areas. We have a broad portfolio, , which really has one of the broadest capabilities in the industry and we've really worked to tune those products to work better together. Many of our products like AppAssure and vRanger provide very rapid recovery times and provide native replication tools that can extend traditional backup and recovery to more of a business continuity solution.

We are also really driving to integrate across that product line. You are starting to see more and more capabilities of each of these different products within each of the other product lines. We have a lot of integration going on between those products and, over time, you will be able to do more and more to address different use case scenarios within these products.

When we talk to customers, we certainly see an interest in consolidation. Customers are moving away from individual replication tools, high availability tools, and tools that they use for offsite data management, and at Dell, we've moved to a place where we can now provide all of that in one tool.

We can do things like data protection using traditional backup and recovery. We can replicate each of those snapshots to an offsite location. We can stand up each of those snapshots in

an offsite location or onsite. You can see how that might start moving you to centralize more of your capabilities into the Dell data protection tool set, and to that end, we recently introduced our backup and disaster recovery suite that provides a capacity based license by which you can use all of the products in our portfolio and consolidate their respective capabilities there.

*In [Part I](#) of this interview series, Brett and I discussed the biggest backup and recovery challenges that organizations face today.*

*In [Part II](#) of this interview series, Brett and I discussed the imperative to move ahead with next gen backup and recovery tools.*

*In [Part III](#) of this interview series, Brett and I discussed four (4) best practices that companies should be implementing now to align the new capabilities in next gen backup and recovery tools with internal business processes.*

*In [Part IV](#) of this interview series, Brett and I discussed the main technologies in which customers are currently expressing the most interest.*

*In [Part VI](#) of this interview series, Brett and I discuss Dell's growing role as a software provider.*

*In [Part VII](#) of this interview series, Brett provides an in-depth explanation of Dell's data protection portfolio.*

*In [Part VIII](#) of this interview series, Brett and I discuss the trend of vendors bundling different but complementary data protection products together in a single product suite.*

---

**DCIG 2014-15 Deduplicating**

# Backup Appliance Buyer's Guide Now Available

DCIG is pleased to [announce](#) the availability of its DCIG 2014-15 Deduplicating Backup Appliance Buyer's Guide that weights, scores and ranks over 100 features on 47 different deduplicating backup appliances from 10 different providers. This Buyer's Guide provides the critical information that all size organizations need when selecting deduplicating backup appliances to protect environments ranging from remote offices to enterprise data centers.



Deduplication is a proven data reduction technology that removes redundant data by only storing one copy of unique data. Reducing storage consumption by up to 20X or more, everyone from small businesses to enterprise data centers benefits through lower storage costs, shortened backup windows and improved backup success rates.

The plug-and-play nature of deduplicating backup appliances has contributed to their success as these appliances quickly and easily fit into almost any size corporate network. They also help organizations better keep up with their ever increasing amounts of production data as they finally have a solution that gives them a means to control the large volumes of backup data that all of this production data generates. This has led to the rapid adoption of these appliances and

them becoming a mainstay in many data centers with organizations now [spending](#) upwards of \$2 billion annually on purpose built backup appliances such as these.

The continuing adoption of these appliances in both small and large organizations is contributing to ongoing innovation in deduplicating backup appliances. Vendors have revamped their product lines by introducing new appliances that are faster, more scalable, more versatile and less expensive. Consider:

- EMC has reintroduced into its Data Domain line the capability to detach a controller from the backup storage so that existing storage shelves can be used with new controllers.
- Quantum has simplified its lineup and now only sells a single system into enterprise shops and another line for midrange backup.
- Dell, HP and Quantum have made their deduplicating backup appliances available as virtual appliances. These typically operate in tandem with their hardware counterparts and provide organizations the option to put a virtual deduplicating backup appliance into highly virtualized small and remote offices without needing to deploy a physical hardware appliance.

It is in this context that DCIG presents its DCIG 2014-15 Deduplicating Backup Appliance Buyer's Guide. As prior Buyer's Guides have done, it puts at the fingertips of organizations a comprehensive list of deduplicating backup appliances and the features they offer in the form of detailed, standardized data sheets that can assist them in this important buying decision.

The DCIG 2014-15 Deduplicating Backup Appliance Buyer's Guide accomplishes the following objectives:

- Provides an objective, third party evaluation of deduplicating backup appliances that evaluates and scores their features from an end user's perspective.

- Scores and ranks the features of each deduplicating backup appliance based upon the criteria that matter most to end users and then presents these results in an easy to understand tables that displays the products' scores and rankings so they can quickly ascertain which deduplicating backup appliance is the most appropriate for their needs.
- Provides a standardized data sheet for each of the 49 deduplicating backup appliances from 10 different providers so users may do quick comparisons of the features that are supported and not supported on each product.
- Gives any organization a solid foundation for getting competitive bids from different deduplicating backup appliance providers that are based on "apples-to-apples" comparisons.

The DCIG 2014-15 Deduplicating Backup Appliance Buyer's Guide Top solutions include (in alphabetical order): EMC [Data Domain 7200](#), 4500 and 990; HP StoreOnce 6500, Quantum [DXi8500](#), [6900](#), and [6802](#); and NEC HYDRAsTOR [HS8-4104R-7920](#), HS8-4006R-720, and HS8-4002S-192.

HP [StoreOnce 6500](#) earned the "Best-in-Class" ranking for the first time. Having revamped its product line over the past year, HP has set a high standard with its flagship StoreOnce 6500 backup appliance to which others are now compared.

Others vying for the top included Quantum, with its revamped product line, and NEC where, through a combination of hybrid and storage nodes, its HYDRAsTOR HS8-4000 lines can scale to a massive 7.9PB in backup storage capacity. However it was through its combination of deduplication, hardware, management and support capabilities that the HP StoreOnce 6500 came out on top.

In doing its research for this Buyer's Guide, DCIG uncovered some interesting statistics about deduplicating backup appliances in general:

- All systems compress data after it is deduplicated.
- 100% offer backup acceleration software. Support for Symantec OST was the most prevalent though others offer support for Accent, AIR and Dell's Rapid Data Access (*RDA*).
- All deduplicating backup appliances deduplicate incoming data while concurrently replicating to another system.
- Almost all bundle deduplication technology with their backup appliance at no extra charge.

As with prior DCIG Buyer's Guides, it accomplishes the following objectives for end users:

- Lists each deduplicating backup appliance by vendor
- Lists out features of each deduplicating backup appliance showing key features supported or not supported
- Scores the features most relevant to end users
- Provides "*at a glance*" reference for companies evaluating specific deduplicating backup appliances or their features
- Provides a deduplicating backup appliance ranking showing how products compare against similar products on the market
- Offers recommendations as to which deduplicating backup appliance rankings and products best align with their specific backup objectives
- Provides 47 deduplicating backup appliance data sheets from 10 different vendors so organizations may compare solutions from one or many technology providers.
- Facilitates and accelerates the process of organizations obtaining bids on competitive products

The DCIG 2014-15 Deduplicating Backup Appliance Buyer's Guide is immediately available. Subscribing users of the DCIG Analysis Portal may access and download the Guide by following this [link](#). Individuals who have not yet subscribed to the DCIG Analysis Portal may test drive the DCIG Analysis Portal for 30

days as well as download this Guide by following this [link](#).

---

# Virtual Backup Appliances Often Do NOT Make Sense Except in Small Environments; Interview with STORServer President Bill Smoldt, Part V

Choosing the right backup appliance – physical or virtual – does not have to be complicated so long as an organization knows the right questions to ask and gathers the appropriate information. However, as organizations are gathering this information, most conclude that a virtual backup appliance is NOT the right answer in most circumstances. In this fifth and final installment of DCIG's interview with STORServer President Bill Smoldt, he explains how to choose the most appropriate backup appliance for your environment and why a virtual backup appliance is probably not the choice you will be making.

***Jerome:** What advice do you provide for choosing the software and sizing the appliance?*

***Bill:** I would turn those around. It really has more to do with choosing the right size. STORServer has tried to make it simple to help a customer through the concept of:*

- I have this much storage
- I have this long of a retention requirement
- I want to deliver this type of recovery experience to my

end users

Those requirements dictate the size of the appliance you need.

[STORServer](#) has a tool for doing just what you are suggesting on our website. But, quite a few of our customers come in already knowing that they want either a TSM engine or a CommVault engine. Some of that is based on their experience, a previous job, or the reputation of the products. But, they often have already made a choice.

From our perspective, other than some of the fringe things that we have to do, it is fine for them to use either product. We have selected these products because they are the best and we know we can solve problems with them. There are a few things that are available on one that are not on the other, but we can use both to supplement features across appliance lines.

What is interesting is that though our customer already typically knows which platform they want, we help them through the decision making process. We help them decide on the size of the appliances and which features they will use, such as instant restore or automated disaster recovery. We also make them aware of features like backing up mobile devices.

**Jerome:** *Is STORServer seeing demands for a virtual appliance?*

**Bill:** There was a time early on in virtualization that I thought there was going to be a significant wave in virtual appliances. We actually started shipping a virtual appliance in 2009. That was our V-Series of appliances.

At that time, it made sense for us to ship our virtual appliance on a pre-built ESX server so we were shipping this on our appliance hardware. We are a VMware system builder, so we were allowed to build and then ship them as a single unit.

We did that because when you think about backing up an

enterprise and putting this in a virtual environment, if this is the primary and only backup for that enterprise, we are trying to back up on the same environment that we're backing up. There is always a little bit of a danger with that approach so you have to design for those issues, and we did that.

But in this case, because of the performance demands, we were going to pretty well take up resources of the entire physical ESX server. The ESX server that we were shipping was the same model—the same hardware as our appliance—so, at the time, we consumed just about all the available I/O, CPU cycle and network bandwidth on that system.

It still gave us the advantage of integrating in with the customer's environment. But when we got into some of the updates and matching their current virtual environment and putting that into a cluster, it did not make quite as much sense to ship it as an ESX server.

Subsequent to the V-Series, we still ship a virtual appliance. We still have customers that will buy a virtual appliance for very specific needs within their environment. However, back to my point from our V-Series, what we learned is that if a customer takes our virtual appliance to back up their entire environment, there are several reasons you virtualize.

- First, you have shared computer, network and storage resource that you are not consuming. If you want to share that among multiple machines, we are going to consume all of the computer and a lot of your data storage and network bandwidth.
- Second, to solve that problem, we have to put that data right back on the same big expensive SAN storage device that we are backing up. You do not want to do that because we are backing that up to account for the possibility of that going down.

In short, most of the advantages of virtualization don't exist when it comes to putting the backup server in a virtual appliance.

This led us back to using the appliance as most companies are going to add cheaper storage somewhere else anyway. We can do this even more efficiently by adding a physical appliance with our own less expensive storage. Again, we are going to consume all the CPU, network and storage on the ESX server anyway. Using an appliance, you now have less expensive storage that is removed from the expensive SAN storage that all of your virtual environments are sharing.

We map those LUNs directly onto our hardware appliance so all of the data movement goes directly over the SAN into our appliance, not disturbing your ESX server. Your whole virtual environment is offloaded which just makes more sense.

There was another issue that really became problematic in the virtual environment. If the physical machine itself fails, I could put my backup appliance on another machine. But, any time we have a tape library or any other device directly attached using NPIV, we could not move that machine off of one ESX server as well.

Having said of all of that, it still makes perfect sense in some environments such as remote offices where there is a virtual environment, to put a small physical or a virtual appliance on that environment, and then replicate that data back to the main data center. Otherwise, using a virtual appliance as a solution for backing up the entire environment does not make sense.

---

# A Backup Appliance Designed to Stand the Test of Time; Interview with STORServer President Bill Smoldt, Part 2

*"The more things change, the more things stay they stay the same."* That nearly 200 year old French proverb still has relevance even in today's modern technology era when one looks at today's backup appliances and how they have both changed and stayed the same since coming on the scene a little over 10 years ago. In this second installment of DCIG's interview series with STORServer's President, Bill Smoldt, he provides some insight into how backup appliances have evolved over the last decade as well as the features they must offer to stand the test of time.

**Jerome:** *How have you seen backup appliances evolve over the last decade? For instance, compare the year 2000 to right now.*

**Bill:** That is really interesting that you should choose the year 2000 as that is about the same year in which [STORServer](#) started shipping its first production appliances. We had some beta products out that year though the basic concept of that appliance and the way STORServer delivers them today has really not changed a lot.

Even at that time we picked the latest hardware because we wanted the appliance to last as long as possible before having to do any kind of a data migration. We used the optimal version of backup software available because, as everyone knows, there are always issues with new versions of software. As such, we tend not to use the latest software releases, but releases that have been in the field for a while.

Even now we still send a consultant on site to install most of

our appliances. Not all of them, but most of them. That is what we started out doing even 10 years ago. From a delivery standpoint, a build standpoint, and a selection standpoint, all those factors are largely the same.

That is not at all true with the features and the complexities and different ways of doing backups. Those are really different than they were 10 years ago.

**Jerome:** *How long are your appliances typically deployed and used in the field? What should customers expect?*

**Bill:** When we first started delivering appliances, the typical refresh rate for technology, for hardware, was about three years. As the economy changed, that grew to four years and then jumped to five years. That was a typical refresh rate.

Of course what is important in our appliances is the data. We have ways of going through a refresh of hardware without having to rewrite all the data. But then along with the economic conditions too, we have many customers who had to extend that refresh rate and could not change out their hardware.

One of the more notable examples was a customer who had an appliance for more than 10 years and had not renewed the technology. The technology that this customer was using on one of our early appliances was a tape drive that could be in a computer museum at this point.

It was really amazing that he was still able to accomplish his primary mission of getting his backups done. In this particular case, it got to the point where we could not even buy the tape drives anymore. Fortunately his appliance got to the point where he had to do a technology refresh and get the latest equipment. That customer then went from using a tape based system to replicating over the internet and now uses the most modern features. However I suspect this same customer may

keep that appliance another 10 years.

We typically design our appliances for growth. Quite often some of those appliances have had to grow far beyond their original design. But because we are very careful in that design, some customers are able to go far beyond the original design and at least keep running.

One of the problems with that approach is, of course, that as the new features come out that have really made backup more exciting in the last few years, it is less likely that we will be able to run that on the older hardware. These new features take a lot more compute power, a lot more memory, and a lot more IO to run.

*In [Part I](#) of this interview series, we explore how large organizations can get up and running faster using STORServer's backup appliances with the knowledge and confidence that they can backup data on any file or operating system.*

*In [Part III](#) of this interview series, Bill shares why the cloud, deduplication and replication are the new "must-have" features that backup appliances must offer.*

*In [Part IV](#) of this interview series, we discuss the new paradigms of backup and recovery and how they are making these activities routine events.*

---

# **The Right Deduplication Method for the Right Data:**

# Interview with Sepaton's Director of Product Management, Peter Quirk, Part III

Anyone who is close to backup recognizes that some types of data deduplicate better than others. However trying to translate that understanding of the environment into meaningful backup policies is almost impossible since it is both complicated and time consuming to successfully implement. Using the new Sepaton Virtuoso platform, it is able to choose the best form of deduplication for each backup stream on the fly. In this third part of my interview series with Sepaton's Director of Product Management, Peter Quirk, we discuss how its Virtuoso platform detects the nature of incoming backup data and then automatically invokes the best deduplication method to deduplicate the data.

**Jerome:** *Can you elaborate on how the Virtuoso platform evaluates incoming application data and then selects the best method in which to deduplicate it?*

**Peter:** The default behavior of the system is to attempt to do inline deduplication on incoming data, unless the data type is specified in our policy engine as to avoid doing inline deduplication on that data.

As the inline engine samples the first inrush of data from the data source, it looks at the success rate of finding hashes for the incoming chunks of data against the hash dictionary. If it is getting a reasonable hit rate, it concludes that the deduplication method it is using is appropriate for doing continued inline deduplication. If it gets zero matches, it might say, "Hmm, this data is either so unique that I have

*never seen it before, or it is not meant to be inline. This is data that would be better deduplicated using post-process.*" In this case it will defer the deduplication to the post-process phase.

Another thing that happens is it may be ingesting data and everything is looking good. But then the inline detects that there is a region of data within the backup that is really unique and does not hash well against existing data. It will mark that range for post-processing, simply bypass it and leave that data for post-process deduplication.

Now there is a whole class of data that you know will be better deduplicated in post-process mode. So we can set those up at the policy engine to say, *"Do not attempt to inline deduplicate that data. Always post-process it."*

This has two advantages. One, if you bypass inline deduplication, you can provide very consistent ingest rates. At this point, you are just writing the data to the back end in a very predictable fashion. Inline deduplication on anyone's system has a variable ingest rate.

As long as the data change rate is low, you are writing relatively little data to the back end. But as the change rate in the data increases or the system gets data types that it has never seen before, it has to write a lot more IO to the back end, and the ingest rate drops. If your goal is to achieve a constant, predictable ingest rate, you might select post-process as the preferred deduplication approach.

Another reason to use post-process deduplication is if you are doing multi-streamed, multiplexed Oracle RMAN backups. They do not work particularly well using inline deduplication. You are better off deferring that data to the post-process engine where we can do byte level deduplication, which will be more efficient than any method of inline deduplication and really reduced the size of the data on disk.

The last case is turning deduplication off altogether. This applies to situations where you have got pathological data such as encrypted data that might have come from an Oracle secure backup. You need to protect it or get it off of its primary storage onto another media for technological diversity or even geographic diversity by means of replication. But you know you cannot deduplicate it as deduplicating it will not save any space. We can handle that workload by designating that data type so it does not deduplicate at all.

**Jerome:** *You mentioned the term “reasonable” in the context of “rate hit.” Can you elaborate on how your system arrives at a “reasonable rate hit”?*

**Peter:** An example use case might be backing up a set of home directories which come through the inline deduplication engine. They deduplicate very nicely every night. But then one day a system administrator adds a bunch of users to the home director hierarchy whose names sort early in the alphabet.

The next time I do a backup, I have effectively prepended a whole set of new data to the backup that was not there in the prior backups. That prepend pushes all of the data down from a matching point of view. It is also highly likely that this data does not lie on hash boundaries which will result in a lower deduplication ratio than I might expect on that next backup.

This can trigger our post-process deduplication engine to look for matches because it uses a sliding window. This sliding window can find the match between today’s backup and yesterday’s backup in the files that existed in both of them. It is very effective at doing that.

This is a classic case where the [VirtuoS0](#) platform automatically invokes post-process in regions of a backup that did not result in particularly good deduplication ratios, or gave us no deduplication ratio. It did not detect any

similarities so it had a deduplication ratio of one. However when we invoke post-processing, it will check to see if it can find some similarities between this and other data.

*In [part I](#) of this interview series, we discuss how databases and virtual machines (VMs) are just beginning to take full advantage of the benefits that disk offers as a backup target.*

*In [part II](#) of this interview series, we discuss what features Sepaton brought forward from its existing S2100 product line and what new features its Virtuoso platform introduced.*

*In [part IV](#) of this interview series, we discuss the challenges of backing up Oracle environments and what new options the Virtuoso platform offers to simplify and ease those challenges.*

---

# **Bringing Together the Best of Today's and Yesterday's Backup Technologies; Interview with Sepaton's Director of Product Management, Peter Quirk, Part II**

A trend that DCIG is seeing among more new products being introduced into the enterprise space is the proclivity to use the best of what has been previously developed in the past and combining that with new technologies that meet the emerging requirements of today's organizations. The new Virtuoso offering from Sepaton reflects this broader industry trend. In

this second part of my interview series with Sepaton's Director of Product Management, Peter Quirk, we discuss what features Sepaton brought forward from its existing S2100 product line and what new features its Virtuoso platform introduced.

**Jerome:** *What elements did Sepaton bring forward from its S2100 product and what new elements did it added to the Virtuoso platform?*

**Peter:** The OptiScale™ architecture at the heart of [Virtuoso](#) is a combination of technologies. First, it is a distributed file system running across many nodes, whose shared storage model deeply integrates deduplication technologies for both inline and post-process deduplication.

Secondly, its process management model and fine-grained instrumentation allow it to control these distributed services as if they form a single system. By using a modern sharded database, distributed across the nodes, the system is able to ride through failures of parts of the database implementing the global dictionary, whether due to hardware or software problems. The management infrastructure provides for a highly manageable and well instrumented system.

We brought forward the physical architecture from our VTL product, the [S2100](#). Both the S2100 and Virtuoso use shared storage and multiple ingest nodes in a very similar fashion.

The difference is that the S2100 did not really support a file system abstraction. It used a very high performance, extent-based storage system since it was storing blocks of virtual tape, whereas Virtuoso presents a file system via its NAS protocols, which is built on top of an object store which will be exposed to applications in a later release.

In the Virtuoso platform, [Sepaton](#) had to implement a fully distributed file system which it did from the ground up. The

software stack is completely new with respect to the file system and inline deduplication. It is really based on an early project we did around a big data platform, which we did not bring to market. There are a lot of HDFS (Hadoop File System) concepts in the file system used by Virtuoso, since backup deals in the main with very large sequential file transfers, a single-writer per stream and seldom more than one reader, which is similar to Hadoop workloads.

Septon implemented its own inline deduplication engine closely coupled to the file system, while the post-process deduplication engine was ported across from its VTL.

Below the Virtuoso file system layer are several components, the most important of which are the data movers. The data mover design is source- and target-aware, and extensible to new sources and targets. Early extensions that we plan to deliver include support for a dedupe source on the client, and support for cloud targets.

At the lowest layers are the services which coordinate the nodes and storage in the cluster, provide for journaling and recovery, and support performance and health instrumentation.

With a new software stack we were able to introduce a completely new approach to managing the system through a web-based interface implementing the latest responsive techniques to support modern browsers on any desktop or mobile device. The web interface is built open REST APIs which will be exposed for partners and customers to use for integration with third-party tools and home-grown automation scripts.

One other element Septon did bring across in part from the VTL was its OST stack for supporting NetBackup and BackupExec. Much of that code has been layered on the file system with few changes, while the replication features like opt\_dup and A.I.R. will interface to the unified replication engine in Virtuoso.

**Jerome:** So it sounds like you brought forward the best of what Sepaton already had on its S2100 and then added some new features as well?

**Peter:** That is largely true but there is more color to it. Sepaton did have this other project going on, as it was intent on building a new product to complement its existing solution to take us into some new markets. Along the way we realized that it could actually be the foundation for a new file-based backup appliance. We had an option to combine the new and old technologies in one product, but prudently decided not to destabilize the VTL platform by grafting a lot of new features onto it to implement a file system

The software foundations of the S2100 were well-suited to a VTL design, but not for the distributed file system that this new backup appliance needed. So Sepaton leveraged its investment in hardware, and knowledge of very high-end large scale backup applications, and built another software platform using mostly the same hardware. Over time, we'll add the VTL protocol to Virtuoso to provide S2100 customers a way to migrate to the Virtuoso if they need a mix of VTL, NAS and OST protocols in one system.

*In [part I](#) of this interview series, we discuss how databases and virtual machines (VMs) are just beginning to take full advantage of the benefits that disk offers as a backup target.*

*In [part III](#) of this interview series, we discuss how Sepaton's Virtuoso platform examines the nature of the application data being backed up and then automatically implements the best methodology to deduplicate it.*

---

# Databases and VMs Only Now Fully Leveraging Disk-based Backup Targets: Interview with Sepaton Director of Product Management Peter Quirk Part I

Ever since using disk as a preferred backup target gained momentum in the late 2000's, there have been those who opine that disk's life in this role would be short lived. But those providers who deliver disk-based backup solutions and are betting their future on them see no slowdown in their adoption. In this first interview with Sepaton's Director Product Management, Peter Quick, we discuss how databases and virtual machines (VMs) are just beginning to take full advantage of the benefits that disk offers as a backup target.

**Jerome:** *Peter, thanks for joining me today. For my readers who may be unfamiliar with you, can you provide a brief background about yourself?*

**Peter:** I'm a veteran of the IT industry. The last five years I have been director of product management at Sepaton. Prior to Sepaton, I did a couple of tours of duty at EMC as well as a small software company in between my stints at EMC. I was also involved in product management and product marketing for a variety of products at Data General. During that time I gained significant exposure to the storage industry to include storage software, graphical interfaces, operating systems and God knows what else.

**Jerome:** *Thanks for that background. Sepaton has been in disk-based backup space for a long time. In that vein, I'd like to*

*get your take on disk's long term viability as a primary backup target (as a replacement for tape) and how that role played into the significant investment that Sepaton made into its newly announced [VirtuoSO](#) platform.*

**Peter:** If you think about backup, there is essentially a disk model and a tape model out there today. Disk and tape technologies serve quite different and complementary roles. Disk is the fastest and most flexible device to land data on. When you add deduplication technology to it, it brings the cost of storage to a competitive level with tape when you take into account various other management costs for tape. For instance, there is more labor involved in managing tape.

But disk does not really have the same retention capabilities as tape. You can store tape for a very long time. We see disk-to-disk backup playing a role in short to medium term retention environments, where a medium retention term might be one to two years.

If you need to retain data for a really long time, seven years or longer, then you are probably going have to include tape in the discussion, just because the average lifetime of a disk these days is around five years. Technology churns are probably even more frequent than that. Tape is going to give you a long term stable storage solution for those very long retention requirements.

But beyond what you are actually storing the data on, you get to the backup model that the backup apps are using, and there are both tape oriented and disk oriented backup applications with many of them supporting both behaviors.

What is interesting is in the very large backup world such as you find with Oracle [RMAN](#) and SQL Server, these databases have their own integrated backup capabilities which allows them to write to disk or remote disk, without the involvement of another application which behaves as a media manager.

Prior to the use of disk-based backup for Oracle, etc., people would use [NetBackup](#) or [Tivoli Storage Manager](#) as a media manager to position tapes and manage the tape catalogs for RMAN or SQL Server. At the moment [Sepaton](#) sees a strong desire to separate database backup from general backup. People can avoid licensing costs for the backup software to act as the media manager application if they use the inbuilt backup capabilities of these big database apps.

The second motivation for using a disk as opposed to a tape model is that in virtual environments there are a lot of snapshots involved. It is easy to snapshot to disk but very difficult to do snapshots to tape. It's just not a natural conversion. When you look at the virtualization of data centers and the increasing role of virtual machines (VMs), clearly a disk model is the way to go.

In summary, disk-based backup provides a more natural operational model for many applications and end-users, the performance is great and can be scaled by adding controllers and spindles. The economics are more than competitive with tape when you add in space-saving technologies like deduplication and data can be replicated to multiple remote sites at speeds limited only by your communications infrastructure.

*In [part II](#) of this interview series, we discuss Sepaton's Virtuoso platform and how it was specifically architected to meet these emerging demands of disk-based backup targets.*

---

# CommVault and STORServer Poised to Deliver the Best Backup Appliance Experience Possible

CommVault – the #1 solution in Virtual Server Backup Software. STORServer – the #1 solution in Backup Appliances. Putting these two together would, on paper, create a very powerful data management and protection solution for organizations. Now it is no longer on paper but “for real.” Today CommVault and STORServer jointly [announce](#) the availability of a new series of STORserver backup appliances powered by CommVault Simpana that are poised to deliver a better backup appliance experience for organizations.

One of the fastest growing areas of data protection today is backup appliances. Yet in the midst of this growth one of the most prominent names in data protection – [CommVault](#) – had no such backup appliance solution to offer its current or prospective customers.

CommVault has for many years hitched its wagon to [Dell](#) who provided CommVault-powered backup appliances. But Dell’s acquisitions of [AppAssure](#) and [Quest Software](#) who each had their own backup software solutions left serious questions as to how the CommVault-Dell relationship would play out.

That question has now been largely answered. While the relationship between Dell and CommVault is still evolving as CommVault [stated](#) on its recent Q1 2014 earnings call, Dell has ceased to sell backup appliances powered by CommVault Simpana.

This transition has put CommVault in a bit of a predicament. Even though it still gets about 20 percent of its revenue from Dell, not shipping its software as a backup appliance even as

the backup appliance space heats up puts in a bad spot with current and prospective customers as well as its resellers. Aggravating the situation, many of these organizations now want and expect providers such as CommVault to provide them with a viable backup appliance option.

Today CommVault answers their call. By aligning with [STORServer](#), CommVault does more than once again make Simpana available as a backup appliance; it makes it available as one of the most robust backup appliance configurations on the planet.

STORServer differs from Dell in an important way: it is laser-focused on providing enterprise quality backup appliances and knows each one in its product line inside out. These strengths contributed to its backup appliance being the **only one** to achieve an **Enterprise** ranking in the most recent DCIG Backup Appliance Buyer's Guide

Together CommVault and STORServer combine to deliver an even better backup appliance experience than what either of these companies could previously do on their own. In STORServer, CommVault gets a company that offers backup appliances built for the enterprise as they run on hardened IBM hardware that is tested and ready for deployment in these environments. In STORServer's case, it gets enterprise, best-in-class software that is built for the specific data management and protection requirements of today's enterprises.

Yet what is maybe most important in these new CommVault-powered backup appliances from STORServer is what organizations receive. They get choice as STORServer makes these backup appliances available in everything from the Under \$10K BA600 to its enterprise EBA2200 model. They get enterprise hardware with the availability, performance, reliability, and stability that they need. They get software that consistently tops analyst's charts and exceeds user expectations. In short, organizations stand to get more than

just a better backup appliance experience; they may get one of the best ones possible.

---

# DCIG 2013 Midrange Deduplicating Backup Appliance Buyer's Guides Now Available

DCIG is pleased to announce the availability of its DCIG 2013 Midrange Deduplicating Backup Appliances Buyer's Guides. In these two [Buyer's Guides](#), DCIG weights, scores and ranks 20 and 29 midrange deduplicating backup appliances respectively from nine (9) different providers.

What is "new" about these two [Buyer's Guides](#) is that DCIG for the first time breaks out these solutions based upon their starting price to better serve the increasingly varied storage needs and pricing constraints that organizations now have. Like all previous DCIG Buyer's Guides they provide the critical information that organizations need when selecting a midrange deduplicating backup appliance to help protect the large amounts of data in their environments.

The storage volume necessary for digital media is already at incredibly high levels and the rate of growth is not slowing down. [Predictions](#) by Gartner show information managed by enterprise data centers will increase 50 times from 2011 to 2015.1 A similar statistic [issued](#) by IDC estimates that the amount of data reached more than 1.8 zettabytes in 2011 and is more than doubling every two years.

As such, organizations need smarter ways to address the problem of runaway storage volumes. Deduplication technology fills this important gap by providing storage reduction ratios of 10 – 20 times of what can be achieved using standard storage technologies. Given the ever-growing need for more storage, it should come as no surprise that purpose built backup appliances (PBBAs) such as deduplicating backup appliances poised for large growth.

Recently [issued](#) 1QCY13 statistics from IDC show the PBBA market is growing faster than all other areas of the external disk storage and data protection markets. PBBAs are increasing at 17% quarterly year-over-year and shipments have increased by 45% on a year-over-year basis.

Deduplication—maybe more than any other technology—has transformed the backup storage space. Other technologies like thin provisioning, utility computing and virtualization have promised and largely delivered high levels of costs savings and operating efficiency in many environments. However, deduplicating PBBAs often provide these types of results almost immediately in nearly every organization into which they are introduced.

Deduplication may be implemented at the source, the target or both in the backup process, but target deduplicating PBBAs are generally the least disruptive to introduce and may be installed into an existing backup environment. They provide the key features that organizations need today more than ever as they enable organizations to:

- Deduplicate data after it is already backed up
- Non-disruptively introduce data deduplication into their backup environment
- Optionally deduplicate data on the server – physical or virtual – using “*upstream processing*” that divides the deduplication between the server and the appliance

The ability of deduplicating midrange appliances to deliver on most if not all of these features is particularly important for those organizations that opt to implement private cloud storage arrays to host many and/or all of their production applications. When deployed into these environments, any planned or unplanned downtime or disruption in service at any time for any reason can have potentially catastrophic consequences for the entire business.

It is in this context that DCIG presents its 2013 Midrange Deduplicating Backup Appliance Under \$50K and 2013 Midrange Deduplicating Backup Appliance Under \$100K [Buyer's Guides](#). As prior Buyer's Guides have done, these two Buyer's Guides puts at the fingertips of organizations a comprehensive list of midrange deduplicating backup appliances and the features they offer in the form of detailed, standardized data sheets that can assist them in this important buying decision.

The 2013 Midrange Deduplicating Backup Appliance Under \$50K and 2013 Midrange Deduplicating Backup Appliance Under \$100K [Buyer's Guides](#) accomplish the following objectives:

- Provide an objective, third-party evaluation of currently available midrange deduplicating backup appliances
- Evaluate, score and rank midrange deduplicating backup appliances from an end-user's perspective
- Include recommendations on how to best utilize this Buyer's Guide
- Provide standardized data sheets for each of the 20 midrange deduplicating backup appliances from nine (9) different providers in the Under \$50K Guide and standardized data sheets for each of the 29 midrange deduplicating backup appliances from nine (9) different providers in the Under \$100K Guide so organizations may do a quick comparison of features while having sufficient detail at their fingertips to make an informed decision

- Provide insight into each deduplicating backup appliance's robustness of its hardware, what deduplication options it offers, and what levels of support and integration it offers for various backup software solutions



The DCIG 2013 Midrange Deduplicating Backup Appliance Under \$50K Buyer's Guide Top 10 solutions include (in alphabetical order):

- ExaGrid Systems EX1000
- ExaGrid Systems EX10000E
- ExaGrid Systems EX2000
- ExaGrid Systems EX3000
- ExaGrid Systems EX4000
- ExaGrid Systems EX5000
- ExaGrid Systems EX7000
- NEC HYDRAsTOR HS3-410
- Quantum DXi4601
- Symantec NetBackup 5230



The DCIG 2013 Midrange Deduplicating Backup Appliance Under \$100K Buyer's Guide Top 10 solutions include

(in alphabetical order):

- ExaGrid Systems EX10000E
- ExaGrid Systems EX13000E
- ExaGrid Systems EX2000
- ExaGrid Systems EX3000
- ExaGrid Systems EX4000
- ExaGrid Systems EX5000
- ExaGrid Systems EX7000
- NEC HYDRAsTOR HS8-3002S
- Quantum DXi6700 Series
- Symantec NetBackup 5230

The ExaGrid Systems EX10000E achieved the “**Best-in-Class**” ranking in the Under \$50K Buyer’s Guide of midrange deduplicating backup appliances. Its companion appliance, the EX13000E, achieved the “**Best-in-Class**” ranking in the Under \$100K Buyer’s Guide of midrange deduplicating backup appliances. The ExaGrid Systems EX10000E and EX13000E successfully deliver the broadest range of features for those organizations looking for a single deduplicating backup appliance in either of these two price categories.

In doing its research for these Buyer’s Guides, DCIG uncovered some noteworthy statistics about midrange deduplicating backup appliances in general:

- 100% of appliances provide the option to use variable length block as a means to deduplicate data
- 96% natively include deduplication software with the appliance
- 96% verify data integrity in some way after the data is initially deduplicated
- 85% support 10GbE ports
- 68% deduplicate data inline
- 56% display deduplication ratios by backup job name
- 40% deduplicate data post-process

- 38% support 8Gb FC ports
- 19% give users the option to deduplicate data either inline or post-process

The DCIG 2013 Midrange Deduplicating Backup Appliance Under \$50K and 2013 Midrange Deduplicating Backup Appliance Under \$100K Buyer's Guides are immediately available. They may be downloaded for no charge with registration by following this [link](#).