DRaaS Providers Come of Age

As more organizations embrace a cloud-first model, everything in their IT infrastructure comes under scrutiny, to include backup and recovery. A critical examination of this component of their infrastructure often prompts them to identify their primary objectives for recovery. In this area, they ultimately want simplified application recoveries that meet their recovery point and time objectives. To deliver this improved recovery experience, organizations may now turn to a new generation of disaster-recovery-as-a-service (DRaaS) offerings.

A Laundry List of DRaaS' Past Shortcomings

DRaaS may not be the first solution that comes to mind to improve their recovery experience. They may not even believe DRaaS solutions can address their recovery challenges. Instead, DRaaS may imply that organizations must first:

1. Figure out how to pay for it
2. Accept there is no certainty of success
3. Do an in-depth evaluation of their IT infrastructure and applications
4. Re-create their environment at a DR site
5. Perform time consuming tests to prove DRaaS works
6. Dedicate IT staff for days or weeks to gather information and perform DR tests

This perception about DRaaS may have held true at some level in the past. However, any organizations that still adhere to this view need to take a fresh view of how DRaaS providers now deliver their solutions.

The Evolution of DRaaS Providers

DRaaS providers have evolved in four principal ways to take the pain out of DRaaS and deliver the simplified recovery experiences that organizations seek.

1. They recognize recovery experiences are not all or nothing events.

In other words, DRaaS providers now make provisions in their solutions to do partial on-premises recoveries. In the past, organizations may have only called upon DRaaS providers when they needed a complete off-site DR of all applications. While some DRaaS providers still operate that way, that no longer applies to all of them.

Now organizations may call upon a DRaaS provider to help with recoveries even when they experience just a partial outage. This application recovery may occur on an on-premises backup appliance provided by the DRaaS provider as part of its offering.

2. They use clouds to host recoveries.

Some DRaaS providers may still make physical hosts available for some application recoveries. However, most make use of purpose-built or general-purpose clouds for application recoveries. DRaaS providers use these cloud resources to host an organization’s applications to perform DR testing or a real DR. Once completed, they can re-purpose the cloud resources for DR and DR testing for other organizations.

3. They gather the needed information for recovery and build out the templates needed for recovery.

Knowing what information to gather and then using that data to recreate a DR site can be a painstaking and lengthy process. While DRaaS providers have not eliminated this task, they shorten the time and effort required to do it. They know the right questions to ask and data to gather to ensure they can recover your environment at their site. Using this data, they build templates that they can use to programmatically recreate your IT environment in their cloud.

4. They can perform most or all the DR on your behalf.

When a disaster strikes, the stress meter for IT staff goes straight through the roof. This stems from, in part, few, if any of them have ever been called upon to do a DR. As a result, they have no practical experience in performing one.

In response to this common shortfall, a growing number of DRaaS providers perform the entire DR, or minimally assist with it. Once they have recovered the applications, they turn control of the applications over to the company. At that point, the company may resume its production operations running in the DRaaS provider's site.

DRaaS Providers Come of Age

Organizations should have a healthy fear of disasters and the challenge that they present for recovery. To pretend that disasters never happen ignores the realities that those in Southern California and Louisiana may face right now. Disasters do occur and organizations must prepare to respond.

DRaaS providers now provide a means for organizations to implement viable DR plans. They provide organizations with the means to recover on-premises or off-site and can do the DR on their behalf. Currently, small and midsize organizations

remain the best fit for today’s DRaaS providers. However, today’s DRaaS solutions foreshadow what should become available in the next 5-10 years for large enterprises as well.

HPE Expands Its Big Tent for Enterprise Data Protection

When it comes to the mix of data protection challenges that exist within enterprises today, these companies would love to identify a single product that they can deploy to solve all their challenges. I hate to be the bearer of bad news, but that single product solution does not yet exist. That said, enterprises will find a steadily improving ecosystem of products that increasingly work well together to address this challenge with HPE being at the forefront of putting up a big tent that brings these products together and delivers them as a single solution.

Having largely solved their backup problems at scale, enterprises have new freedom to analyze and address their broader enterprise data protection challenges. As they look to bring long term data retention, data archiving, and multiple types of recovery (single applications, site fail overs, disaster recoveries, and others) under one big tent for data protection, they find they often need to deploy multiple products.

This creates a situation where each product addresses specific pain points that enterprises have. However, multiple products equate to multiple management interfaces that each have their own administrative policies with minimal or no integration

between them. This creates a thornier problem – enterprises are left to manage and coordinate the hand-off of the protection and recovery of data between these different individual data protection products.

A few years HPE started to build a "big tent" to tackle these enterprise data protection and recovery issues. It laid the foundation with its HPE 3PAR StoreServ storage arrays, StoreOnce deduplication storage systems, and Recovery Manager Central (RMC) software to help companies coordinate and centrally manage:

- Snapshots on 3PAR StoreServ arrays
- Replication between 3PAR StoreServ arrays
- The efficient movement of data between 3PAR and StoreOnce systems for backup, long term retention, and fast recoveries

This week HPE expanded its big tent of data protection to give companies more flexibility to protect and recover their data more broadly across their enterprise. It did so in the following ways:

- ***HPC RMC 6.0 can directly recover data to HPE Nimble storage arrays.*** Recoveries from backups can be a multi-step process that may require data to pass through the backup software and the application server before it lands on the target storage array. Beginning December 2018, companies can use RMC to directly recover data to HPE Nimble storage arrays from an HPE StoreOnce system without going through the traditional recovery process just as they can already do to HPE 3PAR StoreServ storage arrays.
- ***HPE StoreOnce can directly send and retrieve deduplicated data from multiple cloud providers.*** Companies sometimes fail to consider that general purpose cloud service providers such as Amazon Web Services (AWS) or Microsoft Azure make no provisions to

optimize data stored with them such as deduplicating it. Using HP StoreOnce's new direct support for AWS, Azure, and Scality, companies can use StoreOnce to first compress and deduplicate data before they store the data in the cloud.

- ***Integration between Commvault and HPE StoreOnce systems.*** Out of the gate, companies can use Commvault to manage StoreOnce operations such as replicating data between StoreOnce systems as well as moving data directly from StoreOnce systems to the cloud. Moreover, as this relationship between Commvault and HPE matures, companies will also be able to use HPE's StoreOnce Catalyst, HPE's client-based deduplication software agent, in conjunction with Commvault to backup data on server clients where data may not reside on HPE 3PAR or Nimble storage. Using the HPE StoreOnce Catalyst software, Commvault will deduplicate data on the source before sending it to an HPE StoreOnce system.

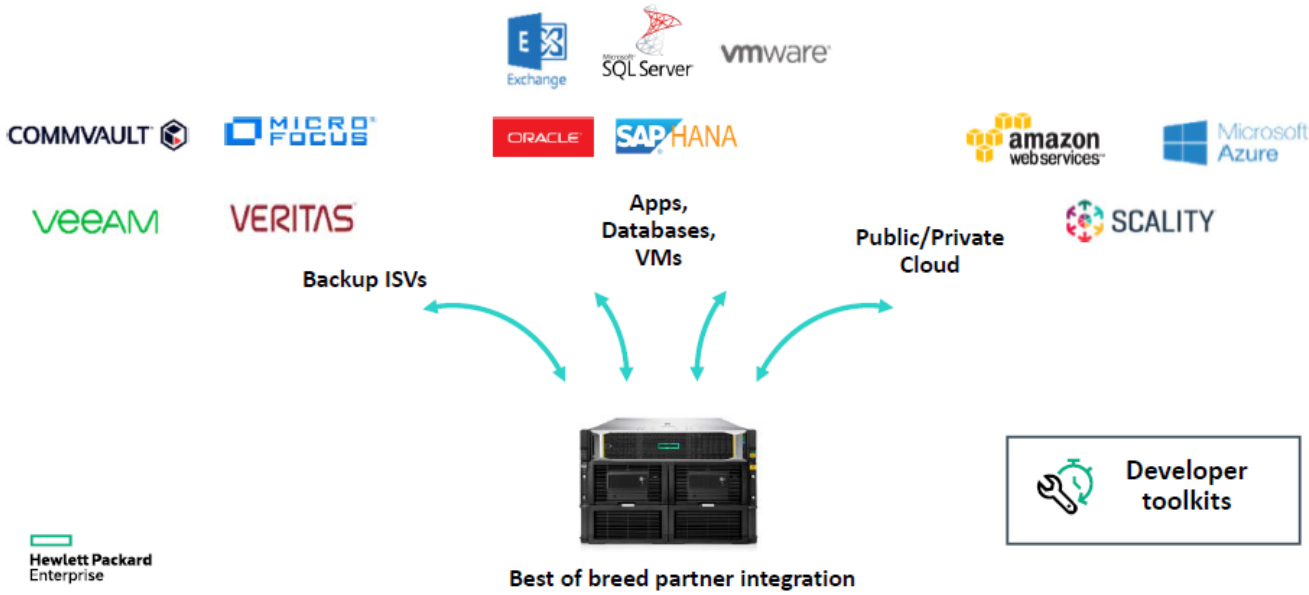

Source:HPE

Of these three announcements that HPE made this week, this new relationship with Commvault that accompanies its pre-existing relationships with Micro Focus (*formerly HPE Data Protector*)

and Veritas demonstrate HPE's commitment to helping enterprises build a big tent for their data protection and recovery initiatives. Storing data on the HPE 3PAR and Nimble and using RMC to manage their backups and recoveries on the StoreOnce systems certainly accelerates and simplifies these functions when companies can do so. But by working with these other partners, it illustrates that HPE recognizes that companies will not store all their data on its systems and that HPE will accommodate companies so they can create a single, larger data protection and recovery solution for their enterprise.

Glitch is the Next Milestone in Recoveries

No business – and I mean no business – regardless of its size ever wants to experience an outage for any reason or duration. However, to completely avoid outages means spending money and, in most cases, a ***lot*** of money. That is why, when someone shared with me earlier this week, that one of their clients has put in place a solution that keeps their period of downtime to what appears as a glitch to their end-users for nominal cost, it struck a chord with me.

The word outage does not sit well with anyone in any size organization. It conjures up images of catastrophes, chaos, costs, lost data, screaming clients, and uncertainty. Further, anyone who could have possibly been involved with causing the outage often takes the time to make sure they have their bases covered or their resumes updated. Regardless of the scenario, very little productive work gets done as everyone scrambles to first diagnose the root cause of the outage, fix it, and then

takes steps to prevent it from ever happening again.

Here’s the rub in this situation: only large enterprises with money to buy top-notch hardware and software backed by elite staff to put solutions in place that come anywhere near guaranteeing this type of availability. Even then, these solutions are usually reserved for a handful of mission critical and maybe business critical applications. The rest of their applications remain subject to outages of varying lengths and causes.

Organizations other than large enterprises daily face this fear. While their options for speed of recovery have certainly improved in recent years thanks to disk-based backup and virtualization, recovering any of their applications from a major outage such as hardware failure, ransomware attack, or just plain old unexpected human error, it may still take hours or longer to complete the recovery. Perhaps worse, everyone knows about it and cursing out the IT staff for this unexpected and prolonged interruption in their work day.

Here’s what caught my attention on the phone call I had this week. While this company in question retains its ideal of providing uninterrupted availability for its end-users as its end game, its immediate milestone is to reduce the impact of outages down to a glitch from the perspective of their end-users.

Granted, a temporary outage of any applications for even a few minutes is neither ideal nor will end-users or management greet any outage with cheers. However, recovering an application in a few minutes (*say in 5-10 minutes*,) will be more well-received than communicating that the recovery will take hours, days, or replying with an ambiguous "*we are making a best faith effort to fix the problem.*"

This is where setting a milestone of having any application recovery appear as a glitch to the organization starts to make sense. Solutions that provide uninterrupted availability and instant recoveries often remain out of reach financially for all but the wealthiest enterprises. However, solutions that provide recoveries that can make outages appear as only a glitch to end-users are now within reach of almost any size business.

No one likes outages of any type. However, if IT can in the near-term turn outages into glitches from a corporate visibility perspective, IT will have achieved a lot. The good news is that data protection solutions that span on-premises and the cloud are readily available now that when properly implemented can well turn many applications outages into a

glitch.

2017 Reflects the Tipping Point in IT Infrastructure Design and Protection

At the end of the year people naturally reflect on the events of the past year and look forward to the new. I am no different. It is as I reflect on the past year and look ahead on how IT infrastructures within organizations have changed and will change, 2017 has been as transformative as any year in the past decade if not the past 50 years. While that may sound presumptuous, 2017 seems to be the year that reflects the tipping point in how organizations will build out and protect their infrastructures going forward.

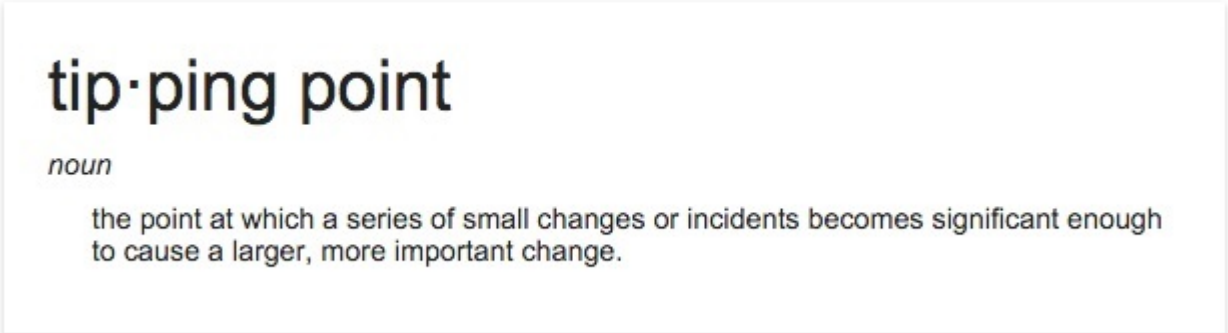

Over the last few years technologies have been coming to market that challenge two long standing assumptions regarding the build out of IT infrastructures and the protection of the data stored in that infrastructure.

1. The IT infrastructure stack consists of a server with its own CPU, memory, networking, and storage stack (or derivations thereof) to support it

2. The best means of protecting data stored in that stack is done at the file level

Over the last two decades, organizations of all sizes have been grappling with how best to accommodate and manage the introduction of applications into their environment that automate everything. They have been particularly stressed on the IT infrastructure side with each application needing its own supporting server stack. While managing one or even a few (less than 5) applications may be adequately achieved using the original physical server stack, more than that starts to break the stack and create new inefficiencies.

These inefficiencies gave rise to virtualization at the server, networking, and storage levels which helped to somewhat alleviate these inefficiencies. However, at the end of day, one still had multiple physical servers, storage arrays, and networking switches that now hosted virtual servers, storage arrays, and fabrics. This virtualization solved some problems but created its own set of complexities that made managing these virtualized infrastructures even worse if one did proactively put in place or have in place frameworks to automate the management of these virtualized infrastructures.

Further aggravating this situation, organizations also needed to protect the data residing on this IT infrastructure. In protecting it, one of the underlying assumptions made by both providers of data protection software and those who acquired it was that data was best protected at the file level. While this premise largely worked well when applications resided on physical servers, it begins to break down in virtualized environments and almost completely falls apart in virtualized environments with hundreds or thousands of virtual machines (VMs).

These inefficiencies associated with very large (and even not so large) virtualized environments have resulted in the

following two trends coming to the forefront and transforming how organizations manage their IT infrastructures going forward.

1. Hyper-converged infrastructures will become the predominant way that organizations will deploy, host, and manage applications going forward
2. Data protection will predominantly occur at the volume level as opposed to the host level

I call out hyper-converged infrastructures as this architecture provides organizations the means to successfully manage and scale their IT infrastructure. It does so with minimal to no compromise on any of the features that organizations want their IT infrastructure to provide: affordability, availability, manageability, reliability, scalability, or any of the other abilities I mentioned in my blog entry from last week.

The same holds true with protecting applications at the volume level. By primarily creating copies of data at the volume level (aka virtual machine level) instead of the file level, organizations get the level of recoverability that they need with the ease and speed at which they need it.

I call out 2017 as a tipping point in the deployment of IT infrastructures in large part because the combination of hyper-converged infrastructures and the protection of data at the volume level enables the IT infrastructure to finally get out of the way of organizations easily and quickly deploying more applications. Too often organizations hit a wall of sorts that precluded them from adopting new applications as quickly, easily, and cost-effectively as they wanted because the existing IT infrastructures only scaled up to a point. Thanks to the availability and broad acceptance of hyper-converged infrastructures and volume level data protection, it appears the internal IT infrastructure wall that prevented the rapid adoption of new technologies has finally fallen.

Veritas Delivering on its 360 Data Management Strategy While Performing a 180

Vendors first started bandying about the phrase "*cloud data management*" a year or so ago. While that phrase caught my attention, specifics as what one should expect when acquiring a "*cloud data management*" solution remained nebulous at best. Fast forward to this week's Veritas Vision 2017 and I finally encountered a vendor that was providing meaningful details as to what cloud data management encompasses while simultaneously performing a 180 behind the scenes.

Ever since I heard the term cloud data management a year or so ago, I loved it. If there was ever a marketing phrase that captured the essence of how every end-user secretly wants to manage all its data while the vendor or vendors promising to deliver it commits to absolutely nothing, this phrase nailed it. A vendor could shape and mold that definition however it wanted and know that end-users would listen to the pitch even if deep down the users knew it was marketing spin at its best.

Of course, Veritas promptly blew up these pre-conceived notions of mine this week at Vision 2017. While at the event, Veritas provided specifics about its cloud data management strategy that rang true if for no other reason that they had a high degree of veracity to them. Sure, Veritas may refer to its current strategy as "360 Data Management." But to my ears it sure sounded like someone had finally articulated, in a meaningful way, what cloud data management means and the way in which they could deliver on it.

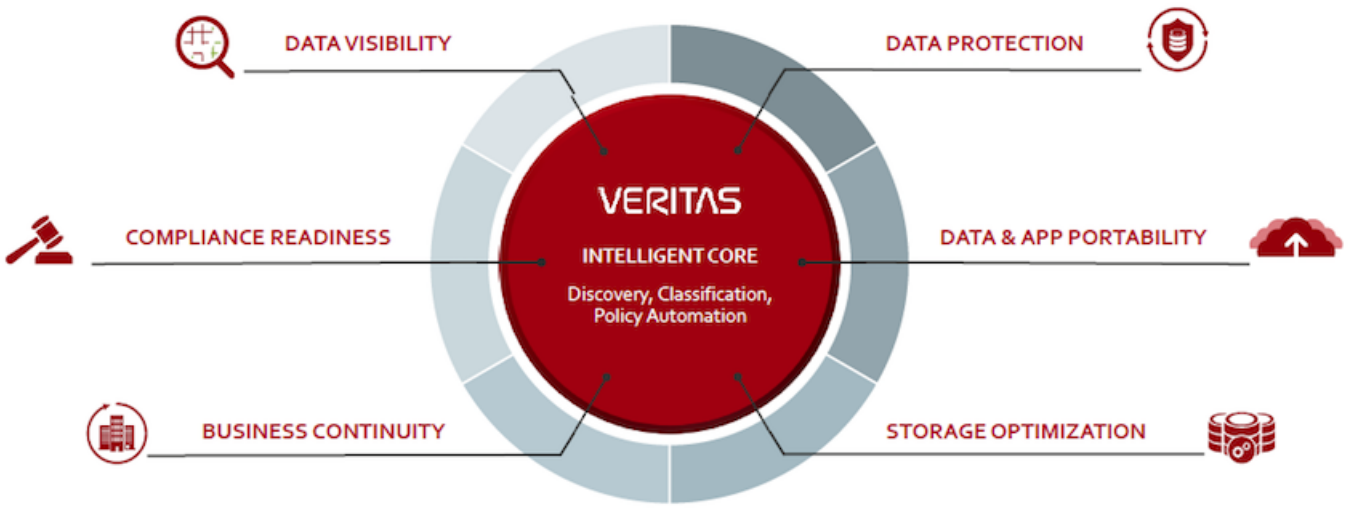

Source: Veritas

The above graphic is the one that Veritas repeatedly rolls out when it discusses its 360 Data Management strategy. While notable in that it is one of the few vendors that can articulate the particulars of its data management strategy, it more importantly has three important components to it that currently makes its strategy more viable than many of its competitors. Consider:

1. ***Its existing product portfolio maps very neatly into its 360 Data Management strategy.*** One might argue (*probably rightfully so*) that Veritas derived its 360 Data Management strategy from its existing product portfolio that it has built-up over the years. However, many of these same critics have also contended that Veritas has been nothing but a company with an amalgamation of point products with no comprehensive vision. Well, guess what, the world changed over the past 12-24 months and it bent decidedly bent in the direction of software. Give Veritas some credit. It astutely recognized this shift, saw that its portfolio aligned damn well with how enterprises want to manage their data going forward, and had the hutzpah to craft a vision that it could deliver based upon the products it had in-house.
2. ***It is not resting on its laurels.*** Last year when Veritas

first announced its 360 Data Management strategy, I admit, I inwardly groaned a bit. In its first release, all it did was essentially mine the data in its own NetBackup catalogs. Hello, McFly! Veritas is only now thinking of this? To its credit, this past week it expanded the list of products to which to which its Information Map connectors can access to over 20. These include Microsoft Exchange, Microsoft SharePoint, and Google Cloud among others. Again, I must applaud Veritas for its efforts on this front. While this news may not be momentous or earth-shattering, it visibly reflects a commitment to delivering on and expanding the viability of its 360 Data Management strategy beyond just NetBackup catalogs.

3. ***The cloud plays very well in this strategy.*** Veritas knows that plays in the enterprise space and it also knows that enterprises want to go to the cloud. While nowhere in its vision image above does it overtly say "cloud", guess what? It doesn't have to. It screams, "Cloud!" This is why many of its announcements at Veritas Vision around its CloudMobility, Information Map, NetBackup Catalyst, and other products talk about efficiently moving data to and from the cloud and then monitoring and managing it whether it resides on-premises, in the cloud, or both.

One other change it has made internally (*and this is where the 180 initially comes in*,) is how it communicates this vision. When Veritas was part of Symantec, it stopped sharing its roadmap with current and prospective customers. In this area, Veritas has made a 180, customers who ask and sign a non-disclosure agreement (NDA) with Veritas can gain access to this road map.

Veritas may communicate that the only 180 turn it has made in the last 18 months or so since it was spun out of Symantec is its new freedom to communicate its road map to current and/or

prospective customers. While that may be true, the real 180 it has made entails it successfully putting together a cohesive vision that articulates the value of products in its portfolio in a context that enterprises are desperate to hear. Equally impressive, Veritas' software-first focus better positions it than its competitors to enable enterprises to realize this ideal.

Full Potential of Disk-based Backup Finally Becoming a Reality with Cohesity DataPlatform 4.0

Organizations have come to the realization that using disk as a backup storage target does more than simply solve backup problems. It creates entirely new possibilities for recovery. But as they recognize these new opportunities, they also see the need for backup solutions that offer them new options for application availability and recoverability backed by ease of management. The latest DataPlaform 4.0 release from Cohesity moves organizations closer to this ideal.

Using tape as a primary backup target is largely dead but the best practices, technologies, and the possibilities to capitalize on using disk as a backup target and as a source for recoveries are still emerging. For instance, secondary storage solutions that only offer "scale-up" architectures create management problems. Additionally, organizations want to do more with their long neglected second or third copies of

data so they want to use these secondary storage solutions to host applications or VMs for the purposes of recovery.

Cohesity's latest DataPlatform 4.0 release illustrates the potential of what the current generation of secondary storage targets can do for organizations to improve their abilities to recover while simultaneously making it easier for them to manage and scale their infrastructure.

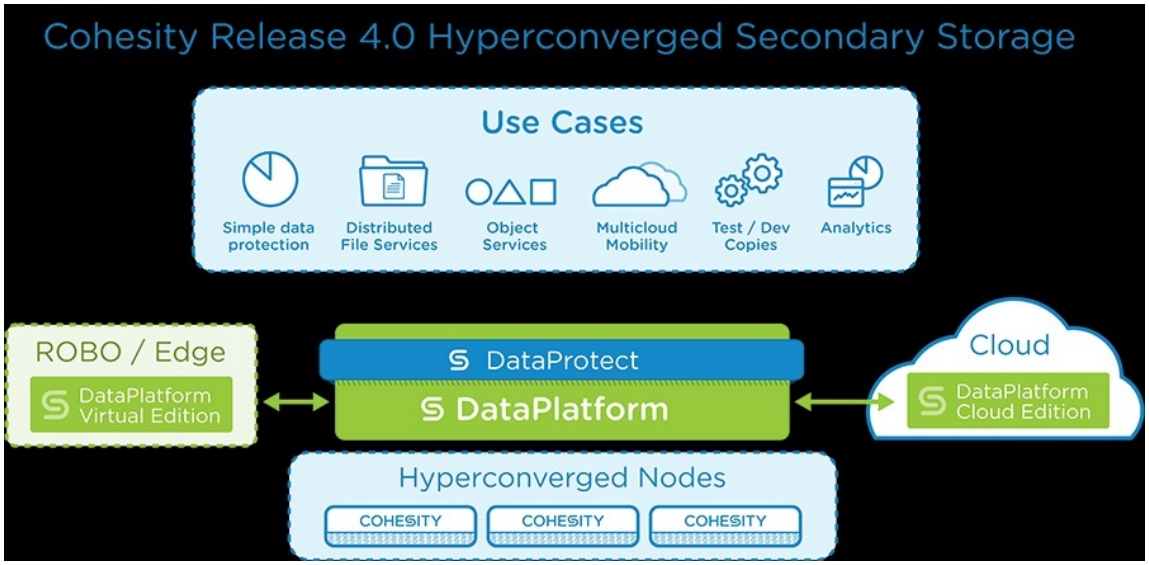

Source: Cohesity

Consider:

- ***Integration with the Pure Storage FlashArray//M series.*** Making snapshots of applications and/or virtual machines (VMs) on your Pure Storage production arrays is a great approach to data protection and instant recovery until one starts to run out of capacity on these arrays. Aggravating this situation, flash costs money. Through its recently announced integration with Pure Storage, organizations can seamlessly move snapshots via SAN or NAS protocols from Pure Storage FlashArray//M arrays to the Cohesity DataPlatform. This frees up availability capacity on Pure Storage arrays while making it possible for organizations to retain snapshots for longer periods of time.

- ***More usable capacity using the same amount of raw capacity.*** Everyone ideally wants something for nothing and Cohesity's latest 4.0 DataPlatform release delivers on this ideal. Previously, it mirrored data between disk drives for data redundancy. Using its new erasure coding technology, organizations can achieve 40% or more storage efficiency when compared to its previous generation product. Further, organizations can achieve this increase in storage capacity by installing this latest software realize on its existing platform.
- ***New options for remote and branch office locations.*** Remote and branch offices are not going away anytime soon yet organizations do not have any more time to manage and protect them. To provide them with higher levels of protection while reducing the time required to manage them, Cohesity introduced its smaller C2100 appliance as well as rolled out a Virtual Edition of its software. The Virtual Edition can be used on traditional backup servers to support current backup and recovery operations or even operate in the cloud when it can serve as a backup target.
- ***Your choice of cloud providers.*** The Cohesity Virtual Edition can operate with multiple cloud providers to include Microsoft Azure and Amazon. In this way, organizations can extend their Cohesity deployment into the cloud to provide instant backup and recovery to ensure uninterrupted operations.

Organizations are now quite acquainted with using disk as a backup target but many still find themselves on the outside looking in when it comes to realizing disk's full potential as a backup target… such as offering fast, simple recoveries that they can deliver at an enterprise scale. The Cohesity DataPlatform 4.0 changes that perspective. Cohesity's use of hyperconverged technology as part of a secondary storage offering solves the key pain points that organizations have for quickly recovering either locally or in the cloud while

simultaneously making their backups easier to manage.

SaaS Provider Pulls Back the Curtain on its Backup Experience with Cohesity; Interview with System Architect, Fidel Michieli, Part 3

Usually when I talk to backup and system administrators, they willingly talk about how great a product installation was. But it then becomes almost impossible to find anyone who wants to comment about what life is like ***after*** their backup appliance is installed. This blog entry represents a bit of anomaly in that someone willingly pulled back the curtain on what their experience was like after they had the appliance installed. In this third installment in my interview series with system architect, Fidel Michieli, describes how the implementation of Cohesity went in his environment and how Cohesity responded to issues that arose.

Jerome: *Once you had Cohesity deployed in your environment, can you provide some insights into how it operated and how upgrades went?*

Fidel: We have been through the upgrade process and the process of adding nodes twice. Those were the scary milestones that we did not test during the proof of concept (POC). Well, we did cover the upgrade process, but we did not cover adding

nodes.

Jerome: *How did those upgrade go? Seamlessly?*

Fidel: The fact that our backup windows are small and we can run during the night essentially leaves all of our backup infrastructure idle during the day. If we take down one node at a time, we barely notice as we do not have anything running. But as software company, we expect there to be a few bumps along the way which we encountered.

Jerome: *Can you describe a bit about the "bumps" that you encountered?*

Fidel: We filled up the Cohesity cluster much faster than we expected which set its metadata sprawling. We went to 90-92 percent very quickly so we had to add in nodes in order to get the capacity back which was being taken up by its metadata.

Jerome: *Do you control how much metadata the Cohesity cluster creates?*

Fidel: The metadata size is associated with the amount of duplicated data it holds. As that grew, we started seeing some services restart and we got alerts of services restarting.

Jerome: *You corrected the out of capacity condition by adding more nodes?*

Fidel: Temporarily, yes. Cohesity recognized we were not in a stable state and they did not want us to have a problem so they shipped us eight more nodes for us to create a new cluster. [*Editor's Note: Cohesity subsequently issued a new software release to store dedupe metadata more efficiently, which has since been implemented at this SaaS provider's site.*]

Jerome: *That means a lot that Cohesity stepped up to the plate to support its product.*

Fidel: It did. But while it was great that they shipped us the new cluster, I did not have any additional Ethernet ports to connect these new nodes as we did not have the additional port count in our infrastructure. To resolve this, Cohesity agreed to ship us the networking gear we needed. It talked to my network architect, found out what networking gear we liked, agreed to buy it and then shipped the gear to us overnight.

Further, my Cohesity system engineer, calls me every time I open a support ticket and shows up here. He replies and makes sure that my ticket moves through the support queue. He came down to install the original Cohesity cluster and the upgrades to the cluster, which we have been through twice already. The support experience has been fantastic and Cohesity has taken all of my requests into consideration as it has released software upgrades to its product, which is great.

Jerome: *Can you share one of your requests that Cohesity has implemented into its software?*

Fidel: We needed to have connectivity to Iron Mountain's cloud. Cohesity got that certified with Iron Mountain so it works in a turnkey fashion. We also needed support for SQL Server which Cohesity into its road map at the time and which it recently delivered. We also needed Cohesity to certify support for Exchange 2016 so they expedited support for that so it is also now certified.

In part 1 of this interview series, Fidel shares the challenges that his company faced with its existing backup configuration as well as the struggles it encountered in identifying a backup solution that scaled to meet a dynamically changing and growing environment.

In part 2 of this interview series Fidel shares how he gained a comfort level with Cohesity prior to rolling it out enterprise-wide in his organization.

In part 4 of this interview series Fidel shares how Cohesity

functions as both an integrated backup software appliance and a deduplicating target backup appliance in his company's environment.

SaaS Provider Decides to Roll Out Cohesity for Backup and DR; Interview with System Architect, Fidel Michieli, Part 2

Evaluating product features, comparing prices, and doing proofing of concepts are important steps in the process of adopting almost any new product. But once one completes those steps, the time arrives to start to roll the product out and implement it. In this second installment of my interview series with System Architect, Fidel Michieli, he shares how his company gained a comfort level with Cohesity for backup and disaster recovery (DR) and how broadly it decided to deploy the product in the primary and secondary data centers.

Jerome: *How did you come to gain a comfort level for introducing Cohesity into your production environment?*

Fidel: We first did a proof of concept (POC). We liked what we saw about Cohesity but we had a set of target criteria based on the tests we had previously run using our existing backup software and the virtual machine backup software. As such, we had a matrix of what numbers were good and what numbers were bad. Cohesity's numbers just blew them out of the water.

Jerome: *How much faster was Cohesity than the other solutions you had tested?*

Fidel: Probably 250 percent or more. Cohesity does a metadata snapshot where it essentially uses VMware's technology, but the way that it ingests the data and the amount of compute that it has available to do the backups creates the difference, if that makes sense. We really liked the performance for both backups and restores.

We had two requirements. On the Exchange side we needed to do granular message restores. Cohesity was able to help us achieve that objective by using an external tool that it licensed and which works. Our second objective was to get out of the tape business. We wanted to go to cloud. Unfortunately for us we are constrained to a single vendor. So we needed to work with that vendor.

Jerome: *You mean single cloud vendor?*

Fidel: Well it's a tape vendor, Iron Mountain. We are constrained to them by contract. If we were going to shift to the cloud, it had to be to Iron Mountain's cloud. But Cohesity, during the POC level, got the data to Iron Mountain.

Jerome: *How many VMs?*

Fidel: We probably have around 1,400 in our main data center and about 120 hosts. We have a two-site disaster recovery (DR) strategy with a primary and a backup. Obviously it was important to have replication for DR. That was part of the plan before the 3-2-1 rule of backup. We wanted to cover that.

Jerome: *So you have Cohesity at both your production and DR sites replicating between them?*

Fidel: Correct.

Jerome: *How many Cohesity nodes at each site?*

Fidel: We have 8 and 8 at both sites. After the POC we started to recognize a lot of the efficiencies from management perspective. We knew that object storage was the way we wanted to go, the obvious reason being the metadata.

What the metadata means to us is that we can have a lot of efficiencies sit on top of your data. When you are analyzing or creating objects on your metadata, you can more efficiently manage your data. You can create objects that do compression, deduplication, objects that do analysis, and objects that hold policies. It's more of a software defined data, if you will. Obviously with that metadata and the object storage behind it, our maintenance windows and backups windows started getting lower and lower.

In part 1 of this interview series, Fidel shares the challenges that his company faced with its existing backup configuration as well as the struggles it encountered in identifying a backup solution that scaled to meet a dynamically changing and growing environment.

In part 3 of this interview series, Fidel shares some of the challenges he encountered while rolling out Cohesity and the steps that Cohesity took to address them.

In the fourth and final installment of this interview series, Fidel describes how he leverages Cohesity's backup appliance for both VM protection and as a deduplicating backup target for his NetBackup backup software.

If We Cannot Scale Our Backup

Solution, We Die; Interview with SaaS Provider System Architect Fidel Michieli, Part I

Every year at VMworld I have conversations that broaden my understanding and appreciation for new products on the market. This year was no exception as I had the opportunity to talk at length with Fidel Michieli, a System Architect at a SaaS provider, who shared his experiences with me about his challenges with backup and recovery and how he came to choose Cohesity. In this first installment in my interview series with Fidel, he shared the challenges that his company was facing with his existing backup configuration as well as the struggles that he had in identifying a backup solution that scaled to meet his dynamically changing and growing environment.

Jerome: *Fidel, thanks for taking time out of schedule here at VMworld to meet and talk with me about how you came to choose Cohesity for your environment. To begin, please tell me about your role at your company.*

Fidel: I work as a system architect at a software-as-a-service (SaaS) provider that pursues a very innovative, agile course of development which is very good at adopting new technology and trends.

My job is on the corporate infrastructure side. I do not work with the software delivery to our customers. Our software is more of a cookie cutter environment. It is very scalable but it is restricted to our application stack. I work on the corporate side where we have all the connectivity, email, financial, and other applications that the enterprise needs, including some of our customers' applications.

I am responsible for choosing and deploying the technology and the strategy to get us to where we need to go and try to develop this division and the strategy to see where things are going. We used Veritas NetBackup and Dell deduplication appliances for backups. Using these solutions, we were constrained as they did not scale to match the demands of our business that grows at the rate we were going.

One of the biggest things that we have to worry about is scale. Often by the time we architect and set up a new solution, we always end up short. If we do not scale, we die.

We were at a crossroads with our previous strategy where we did not scale. It was very expensive to grow and manage. The criticality of the restore is huge and we had horrible restore times. We had a tape strategy. The tape guy came once a week. You could ask for a tape and it would come the next time he stopped by so we would potentially wait six days for the tape to get there. Then you had to move the data, get it off of tape, and convert it to a disk format. Our recovery SLAs were horrible.

I was tasked with finding a new solution. For back-end storage, we looked at Data Domain as we were an EMC shop. For backup software, we looked at Gartner and their magic quadrant

and we chose the first three. With EMC (*now Dell Technologies*) we saw what the ecosystem looked like 12 years ago. A bunch of acquisitions integrated into one solution. It does not get one out of the silo scaling. There were some efficiencies but, honestly, we were not impressed with the price.

Jerome: *Did you find EMC expensive for what it offered?*

Fidel: Yes. It was ridiculously expensive. We also looked at Commvault and just with the first quote we realized this is way too complicated. We are a smaller organization, so we do not have people dedicated to jobs. Commvault quoted us 30 days for implementation engineers. We would have a guy from Commvault in our office for 30 days implementing and migrating jobs. That speaks about the complexity about what we are doing and it speaks to how, when their implementation engineer leaves, who is going to take on these responsibilities and how long is that going to take.

We decided that we should find a more sensible approach.

Jerome: *How virtualized is your environment?*

Fidel: 98 percent. All VMware. This led us to look at a virtual machine backup solution. We had heard very good things about this product but the only problem we had was the back-end storage. How do we tackle the back-end storage? My background is on the storage side so I started looking at solutions like Swift, which is an open source object-based storage as well as ScaleIO. Yet when we evaluated this virtual machine backup solution using this storage, we were not impressed with it.

Jerome: *Why was that? Those solutions are specifically tailored for backup of virtual machines.*

Fidel: To be very honest, NetBackup performed better which I did not expect. I was very invested in the virtual machine

backup solution. We did a full analysis on times and similar testing using different back ends. We found that the virtual machine backup software was up to 37 percent slower and more expensive because of its licensing model so it was not going to work for us.

Jerome: *What did you decide to do at that point?*

Fidel: We talked with our SHI International representative. We explained that we experienced a very high rate of change and that we needed to invest in a solution that in 2-3 years could be supporting an environment that may look radically different than today. Further, we did not want to delay deploying it because we were concerned how competitive we would be. If we delayed, the impact could be huge.

He recommended Cohesity. We recognized that it was obviously scale-out. One of the things that I particularly really liked about its scale-out architecture is that since you originate all of your data copies from the storage, you can have multiple streams from all your nodes. In this way, you are not only scale-out on capacity, but also performance and the amount of data streams that you can have.

In part 2 of this interview series Fidel shares how he gained a comfort level with Cohesity prior to rolling it out enterprise-wide in his organization.

In part 3 of this interview series, Fidel shares some of the challenges he encountered while rolling out Cohesity and the steps that Cohesity took to address them.

In the fourth and final installment of this interview series, Fidel describes how he leverages Cohesity's backup appliance for both VM protection and as a deduplicating backup target for his NetBackup backup software.

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.

---

3 Questions Small and Midsized Enterprises Should Ask to Choose the Right Cloud Hosting Provider

Organizations of all sizes now look to host some or all of their applications with cloud hosting providers and for good reason. Organizations may eliminate the upfront capital costs associated with technology purchases; the overhead associated with managing this technology over time; and, the hassles associated with scaling the infrastructure once it is in place. Yet organizations should not assume all cloud hosting

providers are created equal. If anything, small and midsized enterprises (SMEs) may be particularly susceptible and even find themselves unnecessarily exposed to unexpected outages or extended periods of downtime if they do not carefully choose their cloud hosting provider.

Hosting applications with cloud providers is rapidly gaining momentum but it has particular appeal for SMEs who recognize their need to use technology to efficiently and effectively operate their business but are not technologists, per se. This makes them particularly apt to use cloud hosting providers since these providers can satisfy the technology requirements of many SMEs.

Yet the trap SMEs can easily fall is into is that just because a cloud hosting provider has a physical data center and can host their applications does not necessarily mean these cloud hosting provider is well-positioned to deliver the levels of services and expertise that SMEs may naively assume they possess. While there is no foolproof way to ensure a cloud hosting provider will offer all of the services and expertise to the level that an SME might need or natively expect, there are three questions that SMEs can and should ask to choose the right cloud hosting provider to host their applications.

1. How does the cloud hosting provider implement change control?

Change control is instrumental to successfully running a data center of any size. In essence, it requires creating a check list of tasks that must be completed as well as individuals internally and externally that should be notified and even sign off before a change actually occurs. While this process may sound like routine practice, IT organizations can be lax in implementing and adhering to such procedures. Even if they do exist, these processes are, too often, informal or poorly documented.

These problems do not automatically disappear when one selects a cloud hosting provider. Their1 IT staff do not suddenly and magically obtain the necessary skills and discipline to effectively manage the cloud hosting provider's data center just because they work for a "cloud hosting provider."

Verifying the provider has change control processes in place that are documented and strictly adhered should be viewed as almost a prerequisite prior to hosting one's applications with a cloud hosting provider.

2. How does the cloud hosting provider host and manage application workloads from different customers?

One of the great benefits of using a cloud hosting provider is that rather than having to build your own data center and buy your own computer hardware and software, you can be part of a group that collectively accesses, shares and better utilizes the infrastructure that the cloud hosting provider owns at a lower cost than what you can do by yourself.

However the trick is that the cloud hosting provider must be able to effectively manage application workloads from different organizations. These workloads must be serviced in a way that meet the expectations of all of the organizations accessing and using that same piece of hardware. Should one organization's workloads start to consume and negatively impact the applications of the other organizations also hosted on that computer hardware, how does the cloud hosting provider initially detect and then mitigate the situation to everyone's satisfaction?

If the cloud hosting provider cannot satisfactorily answer that question, be wary about hosting your applications with that provider.

3. How long does it take you to recover my application(s) from an outage?

Organizations did not want to think about backing up and recovering their applications when they hosted them in-house. Now that they host them with a cloud provider, many would prefer to stop thinking about backup and recovery altogether. If so, that would be a mistake.

I had a conversation just this last week with a cloud hosting provider who told me a competitor use Carbonite, an online cloud-base backup service, to backup their clients' applications. While there is certainly nothing wrong with Carbonite, per se, the issue becomes the time it takes to recover their applications should an outage occur. In this case, this provider was aware of an outage that affected some of his competitor's clients. It took them a week or more to first pull back their data from the Carbonite cloud and then to restore their applications.

In short, verify that the cloud hosting provider has a means to keep some recent backups onsite. Ideally they will keep recent backups/copies of data on site for at least the last 24 hours and ideally 7-30 days. While there is nothing fundamentally wrong with backing up to the cloud, quickly recovering from it is a much different proposition. As such, organizations should know up front how quickly they need to recover their applications and verify that the cloud hosting provider has a solution in place to deliver on those expectations.

Expanded Use Cases for Hybrid Cloud Backup Appliances

Viewing hybrid cloud backup appliances strictly in the context

of "*backup and recovery*" is a mindset that organizations must strive to overcome. While these appliances certainly fulfill this traditional role, new use cases are constantly emerging for these appliances. Hybrid cloud backup appliance have now matured to the point where organizations may use them in multiple roles besides just backup.

Hybrid cloud backup appliances minimally solve a challenge that confronts many organizations. They provide onsite backups and give them the flexibility to store backup copies of their data in the cloud for data recovery purposes. Instead of needing to purchase and install backup appliances at two or more locations for data recovery, organizations may use a hybrid cloud backup appliance in conjunction with a public cloud storage provider as a means to:

- Move and store data offsite
- Keep a retention copy or copies of the backup data with the provider long term
- Set the stage for organizations to recover their applications at the provider's site

Once these standard data protection requirements are met, organizations may now look to leverage some of the other features that a number of hybrid cloud backup appliances offer to address their broader business continuity needs.

For example, some hybrid cloud backup appliances give organizations the flexibility to create one or more VMs on the appliance that can host the protected applications and/or their data. Using this features, these organizations can, with comparative levels of ease and simplicity and without disrupting their production environment, test and verify that they can restore protected applications and data.

Some appliances even offer the flexibility to run these applications on a VM in a standby state. In this configuration, if the production application goes offline, the

application running on the standby VM on the hybrid cloud backup appliance can keep the application operational until the production server or VM comes back online.

Restoring applications on a standby VM also gives organizations new flexibility to test application and operating system fixes, patches and upgrades before they apply them on the production server. An organization may bring up an application on a VM on the hybrid cloud backup appliance in a configuration that mimics their production environment.

Fixes, patches or upgrades may then be applied to either the OS and/or application to verify that they work. This technique also gives administrators some practice on how to apply the patch and grants them visibility and understanding into what occurs on the system when the fix or patch is applied such as seeing what alerts are generated (if any) and how much time it takes to complete.

Organizations using public cloud storage providers that offer cloud recovery options may even be able to go so far as to simulate a disaster recovery (DR) at the provider's site. Granted, no organization should expect any of the appliances evaluated in DCIG's recently released Hybrid Cloud Backup Appliance Buyer's Guide to provide an out-of-the-box, turnkey DR solution. However using these appliances and the partnerships they have built with various public cloud storage providers, organizations may realistically look toward creating a viable DR solution much more easily than they have in the past.

Most hybrid cloud backup appliances provide the out-of-the-box backup experience that organizations expect when they acquire them. However for organizations to strictly view and use these appliances strictly in that context is to fail to fully realize the additional value that these appliances now bring to the table. By giving organizations the flexibility to stand-up VMs, test fixes, patches and upgrades and even

simulate disaster recoveries, these appliances give organizations the opportunity and foundation to begin to implement some level of business continuity in their day-to-day operations.

The Dell DL4300 Puts the Type of Thrills into Backup and Recovery that Organizations Really Want

Organizations have long wanted to experience the thrills of non-disruptive backups and instant application recoveries. Yet the solutions delivered to date have largely been the exact opposite offering only unwanted backup pain with very few of the types of recovery thrills that organizations truly desire. The new Dell DL4300 Backup and Recovery Appliance successfully takes the pain out of daily backup and puts the right types of thrills into the backup and recovery experience.

Everyone enjoys a thrill now and then. However individuals should want to get their thrills at an amusement park, not when they backup or recover applications or manage the appliance that hosts their software. In cases like these, boring is the goal when it comes to performing backups and/or managing the appliance that hosts the software with the excitement and thrills appropriately reserved for fast, successful application recoveries. This is where the latest Dell DL4300 Backup and Recovery Appliance introduces the right mix of boring and excitement into today's organizations.

Show Off

Being a show off is rarely if ever perceived as a "*good thing.*" However IT staff can now in good conscience show off a bit by demonstrating the DL4300's value to the business as it quickly backs up and recovers applications without putting business operations at risk. The Dell DL4300 Backup and Recovery Appliance's AppAssure software provides the following five (5) key features to give them this ability:

- ***Near-continuous backups.*** The Dell DL4300 may perform application backups as frequently as every five (5) minutes for both physical and virtual machines. During the short period of time it takes to complete a backup, it only consumes a minimal amount of system resources – no more than 2 percent. Since the backups occur so quickly, organizations have the flexibility to schedule as many as 288 backups in a 24 hour period which helps to minimize the possibility of data loss so organizations can achieve near-real time recovery point objectives (RPOs).
- ***Near-instantaneous recoveries.*** The Dell DL4300 complements its near-continuous backup functionality by also offering near-instantaneous application recoveries. Its Live Recovery feature works across both physical and virtual machines and is intended for use in situations where application data is corrupted or becomes unavailable. In those circumstances, Live Recovery can within minutes present data residing on non-system volumes to a physical or virtual machine. The application may then access that data and resume operations until the data is restored and/or available locally.
- ***Virtual Standby.*** The Dell DL4300's Virtual Standby feature complements its Live Recovery feature by providing an even higher level of availability and recovery for those physical or virtual machines that need this level of recovery. To take advantage of this feature, organizations identify production applications

that need instant recovery. Once identified, these applications are associated with the up to four (4) virtual machines (VMs) that may be hosted by a Dell DL4300 appliance and which are kept in a "standby" state. While in this state, the Standby VM on the DL4300 is kept updated with changes on the production physical or virtual VM. Then should the production server ever go offline, the standby VM on the Dell DL4300 will promptly come online and take over application operations.

- ***Helps to insure application consistent recoveries.*** Simply being able to bring up a Standby VM on a moment's notice for some production applications may be insufficient. Some applications such as Microsoft Exchange create check points to ensure it is brought up in an application consistent state. In cases such as these, the DL4300 integrates with applications such as Exchange by regularly performing mount checks for specific Exchange server recovery points. These mount checks help to guarantee the recoverability of Microsoft Exchange.
- ***Open Cloud support.*** As more organizations keep their backup data on disk in their data center, many still need to retain copies of data offsite without either moving it to tape or needing to set up a secondary site to which to replicate the data. This makes integration with public cloud storage providers to archive retention backup copies an imperative. The Dell DL4300 meets this requirement by providing one of the broadest levels of public cloud storage integration available as it natively integrates with Amazon S3, Microsoft Azure, OpenStack and Rackspace Cloud Block storage.

The Thrill of Having Peace of Mind

The latest Dell DL4300 series goes a long way towards introducing the type of excitement that organizations really want to experience when they use an integrated backup

appliance. It also goes an equally long way toward providing the type of peace of mind that organizations want when implementing a backup appliance or managing it long term.

For instance, the Dell DL4300 gives organization the flexibility to start small and scale as needed in both its Standard and High Capacity models with their capacity on demand license features. The Dell DL4300 Standard comes equipped with 5TB of licensed capacity and a total of 13TB of usable capacity. Similarly, the Dell DL4300 High Capacity ships with 40TB of licensed capacity and 78TB of usable capacity.

Configured in this fashion, DL4300 series minimizes or even eliminates the need for organizations to install additional storage capacity at a later date should its existing, available licensed capacity ever run out of room. If the 5TB threshold is reached on the DL4300 Standard or the 40TB limit is reached on the DL4300 High Capacity, organizations only need to acquire an upgrade license to access and use the pre-installed and existing additional capacity. This takes away the unwanted worry about later upgrades as organizations may easily and non-disruptively add 5TB of additional capacity to the DL4300 Standard or 20TB of additional capacity to the DL4300 High Capacity.

Similarly the DL4300’s Rapid Appliance Software Recovery (RASR) removes the shock of being unable to recover the appliance should it fail. RASR improves the reliability and recoverability of the appliance by taking regularly scheduled backups of the appliance. Then should the appliance itself ever experience data corruption or fail, organizations may first do a default restore to the original backup appliance configuration from an internal SD card and then restore from a recent backup to bring the appliance back up-to-date.

The Dell DL4300 Provides the Types of Thrills that Organizations Want

Organizations want the sizzle that today's latest technologies have to offer without the unexpected worries that can too often accompany them. The Dell DL4300 provides this experience. It makes its ongoing management largely a non-issue so organizations may experience the thrills of near-continuous backup and near-instantaneous recovery of data and applications across their physical, virtual and/or cloud infrastructures.

It also delivers the new type of functionality that organizations want to meet their needs now and into the future. Through its native integration with multiple public cloud storage providers and giving organizations the flexibility to use its virtual standby feature for enhanced testing to insure consistent and timely recovery of their data, organizations get the type of thrills that they want and should rightfully expect from a solution such as the Dell DL4300 Backup and Recovery appliance that offers industry-leading self-recovery features and enhanced appliance management.

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-24×7 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 protections products together in a single product suite.

The Imperative to Move Ahead with Next Gen Backup and Recovery Tools; Interview with Dell Software's General Manager, Data Protection, Brett Roscoe Part II

New technology always sounds great on the surface. However the ramifications of implementing and then managing it can be daunting, intimidating or both. Yet in the case of next

generation backup and recovery tools, the improvements it provides over traditional backup can be so dramatic that NOT adopting and implementing them out is worse than trying to make existing backup software work in today's virtualized, real-time environments. In this second installment of my interview series with Dell Software's General Manager, Data Protection, Brett Roscoe, we discuss why it is imperative organizations move ahead with next generation backup and recovery tools.

Jerome: Many end-users have almost become accustomed to operating in crisis mode, whether it's managing tape backups, managing this deluge of data or trying to meet application RTOs or RPOs. So how do they even get started on utilizing next generation backup and recovery tools as, in the back of their mind, they are probably concerned about making some mistakes? Maybe not deliberately, but if they do not understand their full capabilities, they may not necessarily know how to best implement or manage them correctly. Can you talk a little bit about that?

Brett: Absolutely. It's an interesting situation because the biggest mistake I see really isn't related to implementation of new capabilities. Instead, the biggest mistake I see is that too often customers are not thinking about utilizing next-generation backup in the first place. That is usually the bigger problem. We are all creatures of habit, and we all tend to move down a line of, "*Hey, I am looking to decrease my backup windows and speed up my current processes.*" As opposed to wiping the slate clean and saying, "*I am going to re-architect this thing and eliminate backup windows and nightly backups altogether.*" Or, "*I am going to rethink the method and speed of recovery by utilizing some of these new technologies,*" and, "*I need to re-think how to design my solution around my new virtual infrastructure or cloud-based infrastructure.*"

Overlooking the vast capabilities that exist today is probably

the most common mistake. Customers can actually utilize a Dell backup and recovery solution to run their primary application during down time, which is something that you really cannot do with a lot of our competitors. With Dell, customers can do things like:

- Test and verify their backup points.
- Rapid recovery of an entire image or a single mailbox or a single email.
- Actually run their application while they are restoring it.
- Run their entire application from their backup server or cloud provider

Those are really advanced capabilities that not all customers are thinking about just yet. When you start to think about these things, you can see how customers can quickly simplify their environment by moving from a fragmented solution that uses multiple tools and applications to meet business service level agreements (SLAs) for uptime and recovery, to using one or two tools that can actually protect and meet the service level agreements (SLAs). In addition to just being simpler to manage, there's an economic benefit to consolidating backup and DR into one application rather than maintaining separate tools, one that often far outweighs any of the initial concerns a customer might have about re-architecting.

Having said that, Dell has many customers who *are* using next generation backup tools, and for those customers who have taken that key first step of embracing new capabilities, the next step is to determine how to best align those capabilities with your specific business needs. For example, new technologies can provide five minute RPOs, by creating an application snapshot every five minutes. They can even keep all those five minute increments forever. While customers can do that, they may not necessarily want to do that.

In fact, Dell provides a lot of tools for them to tailor how

many copies they keep over time. So the first day you may keep them every five minutes. The first week you may keep them every hour. Beyond the first week or month, they may fall back to daily snapshots, and then continue to throttle it back from there.

But I see customers who say, "*I am going to create snapshots every five minutes on every application running in my environment and keep them forever.*" Even the best compression and deduplication technology backed by the best CPUs and memory probably come under pressure when you start taking snapshots of 1,000 applications every five minutes and keep every one of them forever.

A better approach is take an inventory of the environment and say, "Here are my critical applications, here are my secondary, less critical applications, and then determine how to tailor the level of protection I am providing to those applications." In this way, customers can get the level of data protection and recovery they need for each application, while reducing the amount of IT infrastructure they need and the load on their network to perform all of these operations.

In Part I of my interview series with Brett Roscoe we discuss the biggest backup and recovery challenges that enterprises face.

In Part III of this interview series, Brett and I will discuss 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 will discuss 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 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 protections products together in a single product suite.

When it Comes to Backup, the Smart Money is on Rapid Reliable Recovery; Interview with StorageCraft's Chief Evangelist Matt Urmston, Part 3

Matt Urmston, StorageCraft's Chief Evangelist and Director of Product Management, has worked in a variety of roles in backup, archiving, data recovery and high availability. In this third blog entry of this interview series, Matt emphasizes that StorageCraft's value is in the recovery process—getting systems back online quickly and efficiently, and having that work every time.

Charley: What do you hear as customers' reasons for choosing your company?

Matt: I can tell you that across the board the one thing that we always hear back from StorageCraft customers is, "*It just*

works." The product does what it is supposed to do.

Your question is interesting because a lot of people are out evaluating backup software. Often times their testing will be heavy on the backup aspect. "How long does it take to run a backup job? Do I see a lot of additional processor cycles being used? Is there a lot of IO? Is it dragging my system to its knees?" This becomes the main focus and extent of the evaluation—it is all about the backup and the performance impact backup is going to have on the systems. They say, "Okay, I have my backup and that all works great.

Many companies will make buying decisions based on the backup aspects of a particular solution, where really they need to be evaluating recovery. In most cases when they compare the recovery process between StorageCraft and one of its competitors, they quickly realize that *StorageCraft's value is in the recovery process—getting systems back online quickly and efficiently, and having that work every time.*

StorageCraft also has had partners who occasionally get wooed away to a competitor because of pricing. We've seen some of our competitors get pretty aggressive, especially when they're going after managed service providers (MSPs). StorageCraft had a few who have come to us and said, "*Listen, I'm negotiating prices with this competitor, is there any way you guys can try to match their pricing?*"

Although StorageCraft works very closely with our partners during sales cycles and do what we can to help them win business, there have been times that we've just had to say 'no'. StorageCraft is comfortable where it is at with its pricing and technology and what it is offering. In most cases those partners have come back after they have had a customer's system go down and they had to perform a recovery.

Charley: What products do you compete against?

Matt: It kind of goes in waves. It is really interesting. We

will have some competitor crop up, we will have an account rep say, "*Hey listen, I am really getting beat up by AppAssure*." For a while AppAssure was a big competitor. There was a time when they were making a big push, but then they kind of went away. We do not run into them much anymore.

Acronis traditionally is a competitor. StorageCraft would go head to head in competitive situations with them. They have also just kind of drifted away. I do not know if they have made some decisions to try to go after some larger environments or not, but StorageCraft does not really run into them a whole lot anymore.

Symantec's System Recovery has been a direct competitor in the past as well. But, as with so many other competitors we have been comfortable with our ability to win deals against them just based on our reputation for reliability when it comes to recovery.

Those are StorageCraft's traditional competitors. Some that are cropping up now are in the Backup and Disaster Recovery (BDR) space where they are offering backup appliances. Then there are a lot of smaller players or newcomers that will pop up on certain deals. It's rare to have a week go by that I don't get that email with a company name in the subject line asking "Have you ever heard of these guys?"

But today if you were to ask our sales team who they are running into most often, it is probably going to be Unitrends, Veeam, or Axcient. There really seems to be a trend toward the hardware solutions.

Charley: Are you looking to move into hardware?

Matt: This is a recurring topic of discussion internally. I believe the first time StorageCraft talked about offering a BDR device it was about two years ago. StorageCraft decided it is not interested in getting into the hardware business because of low margins, refreshes, maintenance, etc. Plus we

have some great partners who are already providing hardware solutions based on StorageCraft Technology.

As we include our partners in this discussion, especially our MSP partners, for the most part they are saying, "*Do not force me into using hardware that I am not familiar with or a device that doesn't meet my standards. I already have relationships with Dell or whoever it might be. Let me buy the hardware.*" Most MSPs are already providing hardware to their customer. They know what type of hardware they like to use and support.

In Part I of DCIG's executive interview with StorageCraft's Chief Evangelist, Matt Urmston, he explained the features that ShadowProtect offers to minimize or even eliminate the possibility of users encountering BSODs when conducting a recovery.

In Part II of this interview series, we expanded on how StorageCraft uses ImageManager to provide a full DR solution that can offer rapid recovery in less than five minutes, and also how ShadowProtect performs equally well whether it's placed in a physical or virtual environment.

In Part IV of this interview series, we delve into how StorageCraft fits itself into the cloud storage landscape and what cloud replication looks like for backup and recovery.

The Cloud, Deduplication and Replication are Must-Have

Features on Backup Appliances; Interview with STORServer President Bill Smoldt, Part III

There is backup and then there is backup. To meet the backup and recovery needs of today's organizations, they need to verify that the selected backup appliance includes the features needed to protect their environment today and positions them to meet their needs into the foreseeable future. In this third installment of DCIG's interview with STORServer President Bill Smoldt, he describes the new must-have features that backup appliances must offer.

Jerome: *What do you consider the new must-have features of backup appliances?*

Bill: The two big features that jump right out would be deduplication and replication. These go hand in hand. STORServer can replicate between appliances or replicate to a public cloud or private cloud without doing deduplication.

However, it makes a lot more sense with deduplication where appliances have like deduplicating algorithms at both sites and they only have to send data that has not been sent before in some duplicate fashion. Using deduplication, we can detect a duplicate chunk of data and just send a pointer into a hashing table instead of sending the data.

Those are features that customers certainly demand at this point. While they are not always necessities, they are certainly on their checklists so we built them into the appliances. We will see a lot more deduplication in different places in the future because it does permit our appliance to store much more data and we have more of that feature moved

down further down into our devices.

Instant recovery is another one of these features that is becoming more important to many customers. This takes two forms. One, they want to bring up a machine while it is still in our backup appliances, such as a virtual machine. Alternatively, they want to mount a disk or a volume right out of our appliance so they can browse the files on the device or use the device. That is another feature that customers are demanding at this point.

Jerome: *What steps is STORServer taking to accommodate public cloud storage connectivity?*

Bill: It is absolutely critical that STORServer offers that service. We offer it on our own cloud that we run. But we also have managed service providers (MSPs) that have bought appliances from us and who offer cloud services. That permits us to exchange data with those MSPs to resell their services and connect our customers with them, and then they also resell our appliances.

That part gets really interesting from an MSP perspective because they can also offer a complete disaster recovery (DR) service within their own cloud and help a customer run their data center inside their cloud. They really like that aspect of DR. So if you lose a facility, you can be back up and running with the MSP more quickly than if you had to try and rebuild your own environment in your data center.

In regards to cloud adoption, most of the customers we deal with directly are reticent to send data to a public cloud at this point. Part of that has to do with some of the things that have gone on with some of the cloud providers. But others are adopting it just fine.

We see a real adoption rate at the low end, where perhaps a customer does not have enough data that requires an appliance at their site and they are able to depend on a backup straight

to the cloud. That is what our MSPs typically offer. We have certainly seen it in the high end where end users can afford the bandwidth that they need to replicate all of their data.

Yet, what is enabling much of the cloud connectivity is the continuing drop in communication costs coupled with its increasing speeds. For example, SMB sites now have communication speeds to the Internet that are equivalent to what local area networks (LANs) had 10 years ago. That is making a big difference.

From strictly a restore perspective, the restore speeds that a customer typically needs to restore data to their own environment is still too slow, so they often need an appliance at their own site. In these cases, we do what we call disaster recovery to the cloud.

There are several layers here. There is backup as a service, disaster recovery as a service, and a whole infrastructure as a service where part of the DR infrastructure is restored into a cloud facility.

STORServer can offer all those different levels at different prices. Most of the adoption rate has happened at a different level where customers want to create their own hybrid cloud or even a private cloud of their own.

Quite often, we will have one of our existing appliance customers during a technology refresh put the new appliance on their own site and then move their old appliance off to a remote site as a target for replication. Or, they have a site that perhaps has a few computers and they can do their own recovery. Most of our customers are more comfortable with that because they still own the data and it is never out of their control.

Many still exhibit a bit of hesitation when putting data out in the public cloud, but that is likely going to change fairly quickly. Public cloud storage is going to get less expensive

and that is happening rapidly. As that public cloud storage gets less expensive, organizations will put more and more data there, particularly for long-term archive. That just makes perfect sense plus it provides a DR layer.

Given fast enough communication speeds, we will see more backup as a service in the cloud. We see that in our own customer base as they go from 100 megabits to 1 gigabit to, in some cases, 10 gigabits infrastructure speeds between sites. Then it makes more sense to provide centralized backup and recovery too.

Part I of this interview series covers why large organizations can get up and running faster using STORServer's backup appliances as they have the knowledge and confidence that they can backup data on any file or operating system.

In Part II of this interview series, Bill Smoldt 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.

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.

StorageCraft Gets a Headstart on the Competition by Enabling Recovery in Minutes

Instead of Hours; Interview with StorageCraft's Chief Evangelist Matt Urmston, Part 2

Companies all want more reliable backup and recovery, with short recovery times when things go awry. *In part II of this interview series with StorageCraft's Chief Evangelist Matt Urmston, we expand on how StorageCraft uses StorageCraft ImageManager and StorageCraft Headstart Restore technology to provide a full DR solution that can offer recovery in as little as five minutes, and also how ShadowProtect performs equally well in physical and virtual environments.*

Charley: You can access ImageManager from anywhere?

Matt: Correct. And that anywhere access to ImageManager was a rudimentary DR solution. StorageCraft included in ImageManager some replication technology based around FTP. Basically, StorageCraft told its customers at that time, that they could take that data and move it wherever they wished. In a lot of cases StorageCraft's partners were taking customer data and moving that either to a colocation facility, or to storage in their own offices.

StorageCraft found it was continually being asked: "*What's the best practice, what's the best way to do this, what should I do once I have the data at the remote site, is there a way that I can now run those images at my remote site for that customer?*" So StorageCraft introduced some new technologies as it grew into that DR phase of development.

One of those DR technologies is VirtualBoot, which provides really rapid recovery. Take any recovery point, simply right click on it, and then select an option to VirtualBoot. It will

spin that image up as a virtual machine, often in less than five minutes.

StorageCraft also introduced a technology called Headstart Restore that is designed for much larger data sets. As StorageCraft grew as a company and grew into some larger environments, it was finding that a traditional restore just wasn't fast enough to meet business requirements. As the data volumes grew beyond the TB size, it could take hours to do a recovery because of disk speed limitations.

StorageCraft wanted to provide a way for its customers to stage recoveries, so it introduced Headstart Restore. In a nutshell, Headstart Restore takes a base backup image and converts it to a virtual disk—VMDK or VHD—and then drip feeds incremental backups to that virtual disk. In the event that a production server crashes, now there is a virtual disk already in place sitting over on an ESX host or on a Hyper-V host. When recovery is needed, Headstart Restore can quickly finalize the process and get that system up and running as a virtual instance without forcing the customer to restore data back to some physical device.

It's really all about recovery time objectives. StorageCraft finds that many managed service providers (MSPs) are competing on their service level agreements (SLAs). StorageCraft technology allows MSPs to get very aggressive with those SLAs, and talk about recovery time objectives in terms of minutes as opposed to hours or days.

Charley: Are you focused on SMB or enterprise?

Matt: StorageCraft is almost exclusively focused on SMB. It also sells almost exclusively through the channel, its network of value-added resellers and managed service providers. StorageCraft also works with some distributors who have some direct to market customers that they work with. But yes, StorageCraft is focused on SMB. Later on, StorageCraft will

look at products and services that might appeal to larger companies.

Charley: Physical and virtual backup. StorageCraft addresses both?

Matt: Absolutely. The StorageCraft backup agent, ShadowProtect, is run in both physical and virtual environments. The agent that runs on an endpoint does not know the difference between a physical machine and a virtual machine. Once ShadowProtect creates an image, the file format is exactly the same whether the source server is a physical box or virtual server. That is really what allows ShadowProtect to take those backups and then restore to wherever the customer prefers.

At the hypervisor level, if a customer has a VMware environment and wants to migrate to a Hyper-V environment, they can migrate by taking backup images of guest machines that are in an ESX environment and then performing a restore of those images to a Hyper-V host. With ShadowProtect, migrating is that simple.

In Part I of DCIG's executive interview with StorageCraft's Chief Evangelist, Matt Urmston, he explained the features that ShadowProtect offers to minimize or even eliminate the possibility of users encountering BSODs when conducting a recovery.

In Part III of this interview series, we discuss what the competitive landscape looks like for StorageCraft and why its customers choose StorageCraft over other solutions.

Eliminating the Dreaded Blue Screen of Death from the Recovery Process; Interview with StorageCraft's Chief Evangelist Matt Urmston, Part 1

The one screen that no system admin ever wants to see is the dreaded blue screen of death (BSOD), especially when doing a recovery. Yet when recovering an application on a different hardware platform, BSODs become a distinct possibility. In this first installment of DCIG's executive interview with StorageCraft's Chief Evangelist, Matt Urmston, he explains the features that ShadowProtect offers to minimize or even eliminate the possibility of users encountering BSODs when conducting a recovery.

Charley: Tell me about your background and position and ShadowProtect.

Matt: I'm Matt Urmston, and I've been with StorageCraft for six years. I've spent 20 years in IT and the bulk of that has been in HA, Archiving, and Backup and DR. I currently play a dual role at StorageCraft, I'm both the Chief Evangelist and Director of Product Management.

Charley: When did StorageCraft start?

Matt: The Company was founded in December 2003. For the first four or five years we funded the company through our OEM arrangements with other folks, basically licensing them our technology so they could create their own backup and disaster recovery solutions.

We are proud that StorageCraft continues to be completely self-funded, continues to be profitable.

Charley: How has StorageCraft moved away from being a point product?

Matt: StorageCraft was very development driven in its early days. The CTO and current director of development were "*The Guys*." The two of them, along with one or two other developers, pretty much developed ShadowProtect around a technology that they had OEM'd at a previous company.ShadowProtect is an image-based backup product that does a really good job of capturing systems in a stable state. It integrates really well on Windows systems into their VSS framework to make sure that when we're capturing those snapshots, that they are in a very clean, stable state, which lends itself well to a reliable restore.

One of the things that the founders wanted to focus on, and that they had noticed in the industry, was that performing the restore was often a very painful and time consuming process, especially when restoring from tape. Companies had to gather all their tapes and catalog all of the data before even starting the restore process. In a lot of cases that tape was not very reliable, so companies lost data at the time of restore.

Early Focus on Rapid Reliable Recovery and Hardware Independent Restore

The focus in the development of ShadowProtect was really on the recovery. We wanted to make sure that when it came time to recover, it was going to be able to do so reliably and successfully and in a timely fashion.

Our focus was on getting rid of that backup window and providing the ability to run backups continuously throughout the day by capturing changes as they occur. This allowed those backups to be taken as frequently as every 15 minutes, then

once snapshots were taken, we wanted to give users the ability to quickly do restores.

Part of the recovery goal that StorageCraft set out to accomplish was to give end users the flexibility to recover to any hardware, not just the exact same hardware that the system was running on.

We went to market as a very stable "*hardware independent restore.*" Customers could take the backup images they had created and restore them to disparate hardware, including going from a physical system to a virtual system or even from one hypervisor to another.

With ShadowProtect, it really does not matter where recovery takes place. ShadowProtect takes images from one system and drops them on another. Our technology makes sure all of the necessary drivers are in place at boot up time so that the system does not blue screen and force the customer to do additional work to get it up and running. ShadowProtect takes care of all of that for them.

The Move from Recovery Product to DR Solution

StorageCraft went to market and quickly realized that if it was going to grow as a company, it would have to move away from just being selling a backup and recovery product and move to providing a disaster recovery (DR) solution—the StorageCraft® Recover-Ability™ solution

To do so, we introduced a product called ImageManager. ImageManager's role, first and foremost, is to manage the incremental chain and to verify images, protecting the integrity of the backup images created by ShadowProtect. Additionally, ImageManager provides the replication technologies that take those images, replicate them offsite, and create a remote DR location. Now in the event of a site disaster, an organization can use ShadowProtect to spin machines up at a remote site and provide for that true

disaster recovery.

In part II of this interview series, we expand on how StorageCraft uses ImageManager to provide a full DR solution that can offer rapid recovery in less than five minutes, and also how ShadowProtect performs equally well whether it's placed in a physical or virtual environment.

VMware Takes Big Step Up in Virtual Server Backup with VDP Advanced 5.5

VMware recently announced the enhancement of its VMware vSphere Data Protection (VDP) Advanced product at the European edition of VMworld. The features and developments included in the 5.5 release decisively provide a robust backup and recovery package for SMBs, both on the high and low end, while becoming a viable alternative for enterprises looking to protect remote datacenters and office locations.
Surveys of the VMware customer base have shown that ***Business Continuity*** and ***Disaster Recovery*** are the highest priorities among nearly half of the respondents they polled. Now with the enhancements in its VDP Advanced 5.5 solution, VMware for the first time has a real answer to these concerns.

Consider this: in 2012 the virtual backup market represented a $1.1B space with a 21% CAGR (*Combined Annual Growth Rate*.) Notably, SMB backup makes up 40% of that space. Therefore it makes sense that VMware would address any shortfalls in its backup and recovery functionality.

The main 5.5 enhancements of note include deduplication, replication, backup verification, application aware backups,

cloud-based backups, EMC integration, and management. VDP Advanced is delivered in the form of a virtual appliance with it now able to effectively protect up to 200 VMs (versus only about 50 VMs in VDP.) More notably, VDP Advanced can now be deployed in conjunction with multiple virtual appliances so that it can realistically protect up to a total of 800 VMs

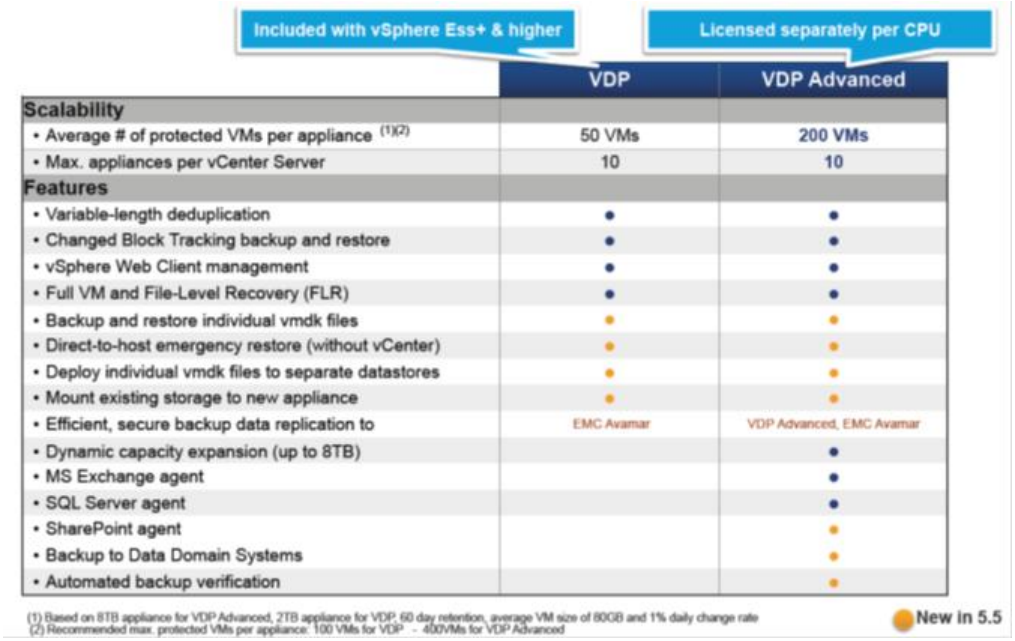

	Included with vSphere Ess+ & higher	Licensed separately per CPU
	VDP	VDP Advanced
Scalability		
• Average # of protected VMs per appliance (1)(2)	50 VMs	200 VMs
• Max. appliances per vCenter Server	10	10
Features		
• Variable-length deduplication	•	•
• Changed Block Tracking backup and restore	•	•
• vSphere Web Client management	•	•
• Full VM and File-Level Recovery (FLR)	•	•
• Backup and restore individual vmdk files	•	•
• Direct-to-host emergency restore (without vCenter)	•	•
• Deploy individual vmdk files to separate datastores	•	•
• Mount existing storage to new appliance	•	•
• Efficient, secure backup data replication to	EMC Avamar	VDP Advanced, EMC Avamar
• Dynamic capacity expansion (up to 8TB)		•
• MS Exchange agent		•
• SQL Server agent		•
• SharePoint agent		•
• Backup to Data Domain Systems		•
• Automated backup verification		•

(1) Based on 8TB appliance for VDP Advanced, 2TB appliance for VDP, 60 day retention, average VM size of 80GB and 1% daily change rate
(2) Recommended max. protected VMs per appliance: 100 VMs for VDP - 400VMs for VDP Advanced

New in 5.5

Source: VMware

Tunneling into the details on each one of these enhancements it is apparent that this version goes a long way to addressing end-user concerns from past releases.

- ***Backup data replication***. Asynchronous replication of backups to a disaster recovery (DR) location in an optimized fashion is new to VDP Advanced. This approach saves network bandwidth and shortens the time to have a VM backup replicated and available at a DR location. Organizations may also independently configure retention periods of the backup at source and destination backup images For example, they may want to keep a local backup for 30 days and retain a DR copy for 60 days.

Replication is performed at the VM level, not at the datastore level, which allows end users to only do DR for the VMs that are the most important to there business. Replicated backup images may also be encrypted if that is a necessity.

Replication streams are encrypted and replication topology options include 1:1, 1:M, and M:1 with VDP Advanced RPOs (Recovery Point Objectives) of 24 hours..

- ***Automated backup verification.*** VDP Advanced provides the ability to schedule, in an automated fashion, a backup verification in a temporary or test area to verify VMs can actually be recovered. This functionality includes post verification reports as well the ability to power-on, boot OS, and restart the application all in an area that will not impact your production landscape.
- ***Application aware backups.*** Although application aware backups are not new to VDP Advanced, this capability has been enhanced greatly. These still require VSS/Application integration, which does involve placing an agent on the VM. However with Exchange VDP Advanced now offers the recovery of individual user mailboxes as well provides support for SQL Server as well as a new agent for SharePoint.
- ***Cloud-based backups.*** Organizations that do not have a DR location or simply want to keep copies of there backup images at a completely separate place can take advantage of its cloud based backup capabilities. Out of the gate VDP Advanced offers two options. They may backup to an Avamar service provider such as Sunguard or to a VDP Advanced replication target at a service provider location. The advantage of this second option is there are no specific hardware requirements for either party involved.
- ***EMC Data Domain integration.*** Many VMware shops have Data Domain deduplicating backup appliances deployed. Using VDP Advanced, they can now take advantage of it DDBoost for Avamar plug-in to further optimize backup to their existing or new Data Domain appliances. This provides the ability to leverage the Data Domain deduplication as

well as the basic and advanced replication capabilities built into Data Domain systems.

It’s clear that VMware has stepped up to the table in this release as organizations continue to get not only their virtualized compute resources from VMware while facilitating the ability to backup, recover and replicate those VM’s all without any assistance from a 3rd party virtual backup software provider. The inclusion of these features in VDP Advanced means VMware can effectively compete with other virtualized backup products in this space though VMware specifically has Veeam in its sights.

These advances are so pronounced versus its prior VDP release that it behooves organizations with fewer than 1000 VMs to take a long, hard look at VDP Advanced as a viable replacement to their existing backup software as this latest release may well meet all of their needs. Many will find it does which could contribute to greatly centralizing and simplifying their internal business continuity and disaster recovery processes as organizations may turn to VMware as a one-stop shop for virtualization as well as the backup, recovery and replication of their virtualized environment.