Special Report

May 2016

DCIG

Dell Data Protection | Rapid Recovery 6.0 Takes Organizations from Today's Good Backups to Putting them at the Forefront of Achieving Application Availability and Recovery

By Jerome Wendt

DCIG
SPECIAL REPORT

Dell Data Protection | Rapid Recovery 6.0 Takes Organizations from Today's Good Backups to Putting them at the Forefront of Achieving Application Availability and Recovery

Table of Contents

May 2016

DCIG

SPECIAL REPORT

Dell Data Protection | Rapid Recovery 6.0 Takes Organizations from Today's Good Backups to Putting them at the Forefront of Achieving Application Availability and Recovery

Executive Summary

Organizations often want to believe that once they solve their backup challenges recovery should just be an afterthought. Having typically spent many hours, days and even years and untold amounts of money to solve their backup challenges, they may expect and even feel entitled to a recovery experience that matches the investment in time and money they have previously made in backup.

While the ability to recover an application and/or its data is certainly a by-product of first having successfully completed a backup, organizations may find that a good backup in no way guarantees that they will get the type of recovery experience that they want or expect. Rather they may end up with a lot of backups that completed successfully but without a means to effectively recover their applications.

Dell Data Protection | Rapid Recovery 6.0 addresses these challenges that organizations frequently encounter in their environments. Unlike other solutions, Rapid Recovery, previously named AppAssure, has always had a “recover first” design with the availability of its Live Recovery, Verified Recovery and Universal Recovery technologies that:

1. Provide near-zero Recovery Time Objectives (RTOs)
2. Automatically verify that backups are recoverable
3. Enable cross-platform recoveries.

Rapid Recovery 6.0 expands upon these core competencies to give organizations more options to protect their data as well as introduces new ways for them to restore it. Rapid Recovery's introduction of agent-less backup gives organizations the flexibility to protect virtual machines (VMs) in the manner which is most appropriate for each individual machine.

Whether one uses agent-less backup that facilitates the non-disruptive introduction of backup of the VM or uses agent-based backup to provide more robust backup and recovery options, organizations may turn to Rapid Recovery to employ either of these options to centralize backup in their environment within a single solution.

Combining these new features with Rapid Recovery's pre-existing recovery features along with its enhanced abilities to recover data directly from public cloud storage providers, organizations have a powerful new tool at their disposal that does much more than ensure “good” backups occur. It provides them with the breadth of recovery options and tools that they need to achieve the levels of application availability and fast restores that the users in their various lines of business now expect and demand.

Good Backups Do Not Equate to Application Availability and Fast Data Recovery

Application availability and data recovery in their various forms are hot topics of conversation in organizations of all sizes. As organizations put their long standing backup problems to bed, they want to focus on identifying solutions that provide near-instantaneous recoveries of their applications and data with little to no interruption to their business operations.

The introduction of agent-less backups, disk-based backup targets, public storage clouds and the use of snapshots have changed the conversation around data protection. These technologies, present in many modern data protection solutions, contribute to backups succeeding the vast majority of the time. DCIG anecdotally finds that organizations who properly implement and use these technologies will successfully complete their backups up to 98% of the time or more.

This minimally has two positive outcomes.

1. Organizations spend far less time trouble-shooting or even thinking about backups when these jobs complete successfully within their scheduled backup windows. This frees IT staff to focus on tasks that add more value to the organization.
2. Viable backups reside on disk media that facilitate faster, easier recoveries.

Yet backup software rarely stores these "good" backups on disk media in a format that positions organizations to re-purpose them to solve other business challenges. Many products that do backup trace their roots to tape-based backup. As such, they store backup data in a format that is neither optimized for in-place recovery nor do they offer a feature to store the data in such a manner.

For example, if an organization wants the option to use the backup copy residing on disk as a hot standby for instant application and/or data recovery, they cannot do so. They must follow the more traditional course of action of first recovering the data back to another system before they can use the data or re-start the application. Similarly, if they want to perform application testing and development or test patches or fixes, they again must first restore the application and/or its data to another physical or virtual machine before they can do any testing against it.

These are the dilemmas that organizations encounter more frequently when they go to recover data. While almost any data management and protection software product

Putting a Fork in Backup

A decade ago, and maybe even as recently as five years ago, organizations could and rightfully did complain about their backups not completing and/or failing for reasons such as:

- Inadequately equipped backup software that did not offer the right features to protect their environment
- Using disk as a backup target to deliver faster backups that completed successfully was either too expensive or impractical to implement as part of the backup process
- New technologies that addressed their backup concerns were too difficult and/or time-consuming to configure, implement and manage

These reasons rarely apply any more. New technologies and a maturation in backup processes put organizations in a position to succeed with backup regardless of the applications that need protection or the type of backups (physical or virtual machine) that they need to perform.

The introduction of disk into the backup process arguably contributes the most to organizations finally having the tools they need to put a fork in backup. The benefits of using disk as a backup target are well-known though its continuing drop in price has largely removed any objections to its use. By way of example, many utility storage arrays price out at below $1/GB for raw disk storage capacity with some providers making disk capacity available on their arrays for as low as .25/GB raw.

Data deduplication further amplifies the argument for using disk as a backup target. When coupled with disk, organizations may reasonably expect to achieve data reduction ratios of 10:1 or greater in their environments. This combination of disk and data deduplication helps to drive down the effective cost per GB even below the cost of tape on a per GB basis.

As disk has emerged as a backup target, data protection software has similarly evolved to better leverage disk as part of the backup process. Minimally this software recognizes disk as either a disk volume or file share (as opposed to a virtual tape device.) This serves to simplify the introduction and use of disk in the backup environment.

This flexibility to use disk more easily has given rise to new features in data protection software. Agent-less backups and snapshot technology combine to accelerate backups even as they reduce server overhead. Further, more backup software offers all-inclusive pricing to facilitate organizations fully leveraging these features. Equally important, many solutions have enhanced their management consoles to equip organizations to deploy these features both locally and remotely while still centrally managing them.

All organizations, regardless of their size, have more choices available to them and are better equipped than ever before to finally put a fork in their backup challenges. Any organization that does not achieve 98% or greater backup success rates needs to re-examine its processes and tools to determine if it is properly leveraging them as the days for excuses for continually failing or broken backup processes should officially be over.

.

can successfully backup data, the options these products offer for recovery do not align with growing organizational expectations for near-immediate access to data so they may quickly perform application recoveries locally, remotely or in the cloud to meet their growing need for uninterrupted business operations.

The Devil is in the Details

Organizations looking for products that offer these next generation capabilities have no shortage of solutions from which to choose. Almost all of these products have, in recent years, made significant enhancements to their feature set to better leverage disk as part of the backup process as well as deliver the breadth of functionality that organizations want and expect. In fact, at first glance, leading products such as these may look very similar when comparing the critical features shown in Table 1.

However the devil is in the details in terms of how these products implement each of these features. For example, all of these products support agent-less backups in environments virtualized using the VMware ESX hypervisor. By leveraging the VMware vSphere APIs for Data Protection (VADP), organizations may protect virtual machines (VMs) without needing to get inside the VM itself. Using these APIs, these products can protect inactive VMs or detect and automatically backup new VMs as they are created which simplifies and eases the protection of virtual machines (VMs).

But enterprise environments, even fully virtualized ones, can rarely be thoroughly protected using just one approach. While agent-less backups offer a base line methodology for protecting VMs, many organizations continue to need agent-based backup to perform more specialized data protection and recovery tasks for VMs hosting business critical applications. They also need to protect any physical machines yet remaining in their environment.

Instant recoveries illustrate this nuance. Instant recoveries give organizations the flexibility to quickly recover and re-start an application residing on a VM to include the mounting and re-booting of the VM. Further, instant recoveries may be done using the agent-less backup feature found in most backup products. However for organizations to perform live recoveries (the recovery of an application ***without*** the need to start and stop them,) they can, for now, only accomplish this task using an agent-***based*** approach.

Similar nuances exist in protecting and recovering applications such as Microsoft Exchange and SQL Server. Each of these applications continually create metadata as part of the management of their respective databases. This metadata must be captured to keep the data in sync and create application-consistent backups.

Capturing this data can currently only be done when placing an agent on the physical or virtual machine on which these applications reside. The differences between the products show up in how and when this agent is placed on the physical or virtual machine.

Some products place an agent on the physical or virtual machine where it perpetually resides. Others only deploy an agent when the backup of the application actually occurs. Each of these approaches to the timing and duration of the agent deployment has its own pros and cons that organizations need to understand and weigh before selecting one.

Even how these products leverage cloud services providers becomes more critical to understand. Many products already connect to cloud services providers and use their cloud storage as a backup target to archive or store data.

Table 1: Critical Features

	Dell Data Protection | Rapid Recovery 6.0	CommVault Data Platform	StorageCraft ShadowProtect	Unitrends Enterprise Backup	Veritas Backup Exec 15
Agent-less Backups	✓	✓	✓	✓	✓
Application Consistent Backups (Microsoft)	✓	✓	✓	✓	✓
Backup Verification	✓	✓	✓	✓	✓
Connectivity to Public Cloud Storage Providers	✓	✓	✓	✓	✓
Deduplication	✓	✓	✓	✓	✓
Encryption	✓	✓	✓	✓	✓
Instant Recoveries	✓	✓	✓	✓	✓
Multi-Hypervisor Support (ESX & Hyper-V)	✓	✓	✓	✓	✓

However many organizations envision using these cloud services providers to do much more than simply archiving and storing data in the future.

They anticipate more quickly and effectively recovering data from these providers as well as laying the groundwork for performing recoveries in the cloud. To achieve these longer term objectives, organizations need to identify products that offer these needed cloud connectivity and management options. While many products offer these features, their maturity in making them usable in a practical manner varies widely.

Dell Data Protection | Rapid Recovery 6.0 Bridges the Backup and Recovery Gap

Bridging this gap between existing backup requirements and new expectations for recovery dictates that organizations identify and implement products positioned to do so. This demands that organizations look beyond the basic check box feature functionality that these various products support to fully understand how they deliver them and what differentiates their implementation from other products.

Dell Data Protection | Rapid Recovery 6.0 in particular sets itself apart from competitors by providing organizations with a single solution to:

- Centralize data protection and recovery across their environment
- Optimize available capacity and performance during the backup process
- Recover from public cloud storage providers

Centralizing Data Protection and Recovery across the Environment

Organizations want to ease the protection and recovery of applications across their environment. This starts by identifying and implementing products that address the many nuances associated with effectively protecting and recovering applications in today's modern data centers. To accomplish these tasks, a product must minimally offer multiple forms of data protection.

In today's increasingly virtualized world where VMware is the predominant hypervisor, it is almost an imperative that the solution supports its VADP APIs. Using VADP, organizations may much more easily and non-intrusively deploy a backup solution into their virtualized environment, protect their VMs and then scale backup of these VMs with minimal to no application or end-user interruption going forward.

Yet the introduction of agent-less backup should primarily be viewed as a complement to agent-based protection. Agent-based data protection remains relevant as it grants organizations the flexibility to use a single solution to manage data protection and recovery across their entire physical and virtual environment.

In most situations, physical machines still require the use of agents to protect their applications and/or operating systems they host. Using

Five (5) Key Benefits and Drivers for Agent-less Backup

The inclusion of the VMware vStorage APIs for Data Protection (VADP) in the VMware ESX hypervisor in 2010 opened the doors for organizations to do agent-less backup of their VMs. These APIs have largely become the default way that organizations look to protect their virtualized environment as they enable organizations to easily deploy, manage and scale their protection of VMs in five (5) ways:

1. ***Comprehensive data protection.*** Using these APIs, organizations may implement policies that automatically protect every VM on every ESX server in their environment. These policies may even extend to protecting dormant or idle VMs that would normally never be protected as an agent-based approach to backup only works if the VM itself is active.
2. ***Cost-effective.*** Organizations of almost all sizes have remote locations with many of them having VMs in one or more of those sites. Agent-less backup centralizes the management of the protection and recovery of the VMs in these remote sites. This, in turn, eliminates the need to send individuals out to configure backup, recover data or VMs or train people at those sites to perform and manage these backup and recovery tasks.
3. ***Less intrusive.*** By eliminating the need to install agents on VMs, administrators no longer need to schedule time or get permission from an application owner to login to a VMs to install agents on one. Further, installing an agent on a VM after the VM is already in production may require a restart of the VM. Agent-less backups eliminate that requirement.
4. ***Protects all guest operating systems.*** The VMware ESX hypervisor can host VMs that may contain one of any number of different guest operating systems (OSes) to include Apple OS X, Linux (various releases,) Novell NetWare and Microsoft Windows. Using agent-less backup, organizations may protect and recover any of these OSes without the need to deploy specific products that offer backup agents for them.
5. ***Satisfy compliance, regulatory and/or internal policy requirements.*** Many organizations have VMs to which they may need to restrict and/or limit access due to sensitive data residing on them. In these instances, it becomes difficult if not impossible to centrally protect and manage the data residing on these VMs. Agent-less backup opens the door for organizations to manage these VMs since there is no need for backup administrators to install agents on them.

APIs to perform agent-less backup on physical machines rarely surfaces as an option. Even in virtualized environments where APIs are available for agent-less backup, VMs may require agents to perform application consistent backups and/or granular restores.

Application consistent backups and granular restores specifically come into play when protecting and recovering database and email applications. Whether one recovers a specific table in Microsoft SQL Server or a specific Microsoft Exchange mailbox or email message within that mailbox, the backup product needs to have visibility into these applications at their core to give organizations the flexibility to restore these individual components.

While many backup and recovery products offer these baseline agent-less and agent-based backup and recovery features, Rapid Recovery takes recovery a step beyond what almost any other product in the market offers today.

Using its Live Recovery feature, organizations may initiate immediate application recoveries for their Windows-based applications that have a Rapid Recovery agent installed on them. Using this option, the application recovery begins even through the bulk of the application data still resides in Rapid Recovery's backup repository. As the application runs in production, Rapid Recovery in the background restores the data from Rapid Recovery's backup repository to the server hosting the application, whether that server is a physical or virtual machine.

Organizations may also use Rapid Recovery to centrally protect the multiple versions of operating systems that likely exist in their environment. While Rapid Recovery supports all Microsoft operating systems to include the latest Windows 10 and Windows Server 2016 releases, it also supports the leading Linux platforms. These include the Community ENTerprise Operating System (CentOS), Oracle Linux, Red Hat Enterprise Linux (RHEL) and Ubuntu, among others.

Using Rapid Recovery to protect Linux and/or Windows platforms also gives organizations the flexibility to protect, recover and export Unified Extensible Firmware Interface (UEFI) partitions that the latest versions of these operating systems use. Poised to replace BIOS, UEFI manages interactions between host operating systems such as Linux and Windows and the hardware on which these operating systems reside.

As this largely hidden changeover in the underlying firmware used by server hardware occurs, backup software

Continued Benefits and Need for Agent-based Backup

The advent of agent-less backup makes it easy to believe that the end of agent-based backup is nigh. Nothing is further from the truth. While agent-less backup addresses many challenges around the protection and recovery of VMs, compelling reasons persist for organizations to continue to offer agent-based backup as an alternative to agent-less backup. Consider:

1. ***Agents can better monitor server CPU and tune backup performance accordingly.*** Agent-less backups still incur overhead on the underlying physical machine as they consume server CPU and memory while the backup occurs. The consumption of these resources may, in turn, impact other VMs running on the host. Using an agent-based approach, organizations can better monitor the CPU and/or memory consumption on the host by the backup process. If the backup process does impact applications on other VMs on that host, an agent-based backup is better positioned to throttle the backup process to lessen its effect on other VMs on that host.

2. ***Agents better facilitate understanding of database or email applications as well as creating application-consistent backups.*** Almost every database and email application generates logs that must be captured and then applied to ensure their successful recovery. Capturing all of this metadata almost always necessitates the deployment of some type of agent on the VM to ensure this recovery. Even a number of backup solutions that promote themselves as agent-less still employ the use of an agent at the beginning of a backup of these types of applications to capture the data necessary to ensure a successful recovery. The main difference is that at the end of the backup they remove their agent while agent-based solutions permanently place an agent on the VM.

3. ***Agent deployment and management is less of an issue.*** Agent-less backup certainly does eliminate the time, effort and management overhead associated with deploying agents initially and then managing them long term. However agent-based backup solutions have taken numerous steps over the years to automate the initial deployment of agent and then maintain them long term. Through their integration with management consoles such as Microsoft Management Console (MMC) and/or VMware vSphere vCenter, organizations may deploy agents on VMs managed by this software as part of each VM's initial setup and/or ongoing management.

4. ***Agents gather more detailed, technical information about the volumes and applications within the VM.*** Agent-less backup lacks visibility into the data contained within each VM. As such, when it completes backups, it often lacks the necessary metadata to quickly restore specific components within the VM, such as specific files or folders. Rather it must restore the entire VM and mount it before the administrator can navigate to and restore the needed data. Agent-based approaches capture and retain this type of metadata so administrators may more quickly do recoveries that require this type of detail to be at their fingertips.

can capitalize on new feature functionality that UEFI offers. Rapid Recovery already leverages UEFI to accelerate a bare metal restore (BMR) for either a Linux or Windows recovery. In these situations, Rapid Recovery calls upon UEFI to automate the partitioning of disks that must initially be done as part of a bare metal restore.

The Rapid Recovery Universal Recovery Console (URC) facilitates and centralizes recoveries in at least two ways.

- At a basic level, administrators may mount recovery points for Linux systems in the URC GUI to simplify file recovery.
- At an advanced level, the Rapid Recovery URC leverages UEFI to do bare metal restores of applications to hardware dissimilar from the hardware on which the applications originally resided. By communicating with the UEFI of the hardware on the recovery server, the URC can identify and inject the appropriate storage controller drivers into the operating system after the data has been restored to allow the server to boot using the recovery server's hardware.

Using these various Rapid Recovery tools, organizations finally have a solution at their disposal that they may use to consolidate backup cross-organization. More importantly, Rapid Recovery puts at their fingertips a powerful set of tools to recover data locally or remotely to the same or even dissimilar hardware using a centralized management console to perform any of these tasks.

Incorporating Cloud Services into Backup and Recovery

Storing data with cloud services providers has finally come of age. Nearly every organization, regardless of its size, wants to leverage cloud services at some level to reduce their local storage costs and easily scale their available backup capacity.

To satisfy these demands, many backup solutions integrate with multiple cloud services providers to meet corporate expectations for managing cloud services as part of the backup process to include storing and archiving backup data longer term in the cloud. In this respect, Rapid Recovery already natively supports as many or more cloud services providers as other leading backup software providers (Table 2).

Once connected to any of these cloud services providers, Rapid Recovery 6.0 provides organizations with two options to create archives in them: one-time and continuous. It creates one-time archives for specified machines while its continuous archives occur automatically at scheduled dates and times so machine backups may be more frequently created in the cloud.

Yet when organizations look beyond archiving backup data to the cloud to actually recovering it from the cloud further differences between solutions emerge. Meeting heightened corporate expectations to recover their data from the cloud highlights two ways that Rapid Recovery differentiates itself from many of its competitors.

For instance, organizations may recover files directly from archives stored in the cloud. By connecting a cloud archive to a Rapid Recovery Core machine, that cloud archive is then made available for recovery along with local archives. Files residing in the cloud archive may then be selected and restored.

Organizations may even do a Bare Metal Recovery (BMR) from data residing on a cloud archive. To do so, an organization needs to initiate the recovery on the machine using the Rapid Recovery boot media. Once the organizations installs the Rapid Recovery boot image on the machine, the Rapid Recovery Wizards may connect directly to the appropriate cloud provider that hosts the cloud archive. The organization then only needs to select the appropriate cloud archive and the desired recovery point to begin the BMR.

Table 2: Backup Solution Integration

Public Cloud Storage Provider	Dell Data Protection | Rapid Recovery 6.0	CommVault Data Platform	StorageCraft ShadowProtect	Unitrends Enterprise Backup	Veritas Backup Exec 15
Amazon S3	✓	✓	✗	✓	✓
Microsoft Azure	✓	✓	✗	✗	✗
RackSpace (OpenStack)	✓	✓	✗	✓	✗

Growing Organizational Expectations for Recovery Play Directly into the Strengths of Dell Data Protection | Rapid Recovery 6.0

Organizational expectations around backup and recovery have become easier than ever to define. They expect backups to work in both their virtual and physical environments and they expect more options than ever when it comes to recovery. These growing enterprise expectations for recovery play directly into the strengths of Dell Data Protection | Rapid Recovery 6.0.

The introduction of agent-less backup into Rapid Recovery 6.0 gives organizations the flexibility they need to backup each machine in their virtual environments in the manner that best meets its respective needs. Whether it is an agent-less backup with less intrusive deployments or an agent-based backup that provides more granular backup and multiple recovery options, Rapid Recovery now satisfies these multiple requirements that organizations often have when looking to implement a single product to protect and recover their entire environment.

By then more tightly integrating with public cloud storage providers such as Amazon, Microsoft and RackSpace, organizations may also more confidently proceed with storing their backup data with cloud services providers knowing that they can directly recover from them. Whether it is restores of individual files or bare metal recoveries, the new flexibility that Rapid Recovery offers to directly mount archives residing with cloud services providers as recovery points makes storing data in the cloud a more viable option for organizations going forward.

Organizations want to simplify and centralize the backup of their data across their environment. But along with that, they want more options than ever to recover their data in the manner that best meets the needs of each of their applications. Rapid Recovery aligns with this new mix of organizational expectations. Its introduction of agent-less backup and heightened abilities to perform recoveries from the cloud coupled with its pre-existing strengths in near real-time recovery combine to provide organizations with the features they need to backup and recover their applications however and in whatever manner they need. ■

About DCIG
DCIG empowers the IT industry with actionable analysis that equips individuals within organizations to conduct technology assessments. DCIG delivers informed, insightful, third party analysis and commentary on IT technology. DCIG independently develops and licenses access to DCIG Buyer's Guides and the DCIG Analysis Suite. It also develops sponsored content in the form of blog entries, competitive advantage reports, customer validations, executive white papers, special reports and white papers. More information is available at **www.dcig.com.**

This special report was commissioned by Dell. • May 2016