phdvirtual E-GUIDE

RECOVERY RULES: Your E-Guide to Best Practices for Data Protection and Recovery in Virtual Environments

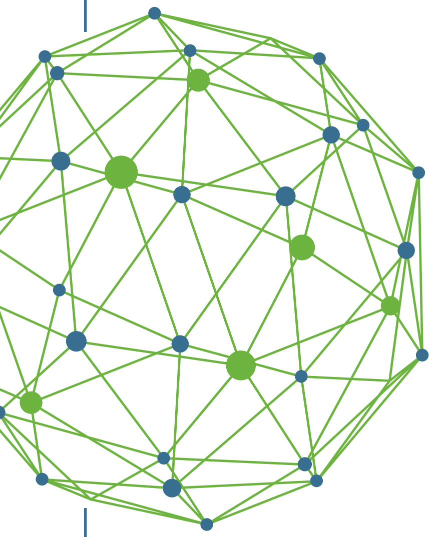

Unmatched Value for Your Virtual World

WWW.PHDVIRTUAL.COM

Table of Contents

RECOVERY RULES:
Your E-guide to Best Practices for Data Protection and Recovery in Virtual Environments

How to ensure your organization can recover valuable data faster, more efficiently and effectively, and improve the value and responsiveness of your IT services.

Published by PHD Virtual © 2013 PHD Virtual Technologies. All rights reserved.

INTRODUCTION:

Beyond backup to the new rules for recovery

This E-Guide was created by PHD Virtual to help IT professionals gain a better understanding of the challenges and benefits involved in adopting best practices for recovering data in virtual environments. While much has been written and presented regarding the value of backing up data, at the end of the day, backup is a fruitless exercise unless you can recover reliably and quickly.

This E-Guide will help you do exactly that.

At PHD Virtual we've been focused on virtual backup, replication and recovery for seven years. During that time we have listened to and worked closely with customers from all sizes of business as they grappled with building, maintaining and improving their virtual environments. This E-Guide will share with you the hard earned lessons and best practices across the entire continuum of data protection and recovery – from granular data, to complex applications, up to entire data centers.

Utilizing a Virtual Backup Appliance (VBA) architecture that is unique in our industry, PHD Virtual has developed solutions that we believe deliver unmatched value for your investment in virtual environments.

To explain those benefits we have incorporated several scenarios following our descriptions of best practices in each chapter of this E-Guide:

Chapter 1: Available Now
Local data protection and recovery

Chapter 2: Available Now
Offsite data protection and recovery

Chapter 3: Available December 2013
Disaster recovery failover, failback and testing

Chapter 4: Available January 2014
Data center migration

In gaining a better understanding of recovery rules, you will be helping your organization recover valuable data and applications in your virtual environments rapidly, efficiently and effectively, thereby enhancing the overall value of your business.

This E-Guide is a living document, and one that we intend to update and improve over time. Your comments and feedback are always welcome at PHD Virtual. If you have questions or comments please contact us at **info@phdvirtual.com.**

CHAPTER 1

Local data protection & recovery

Local backup and recovery in virtual environments has seen significant improvements since its introduction more than 10 years ago.

Today IT professionals can take advantage of several innovations that help to improve the speed and reliability of data backup and recovery that once claimed considerable space and expense in the datacenter.

LOCAL DATA PROTECTION & RECOVERY
BEST PRACTICE # 1:

Streamline recovery and reduce costs by leveraging virtualization APIs

By virtualizing data into a collection of files not tied to a piece of hardware, virtualization backup and recovery has the potential to simplify and speed the process when using the API's of any hypervisor. There are a couple of best practices that can assure you take full advantage of your virtual environment.

No need to add cost and complexity with agents or additional physical servers

The more innovative approach to virtual backup and recovery utilizes what is known as a Virtual Backup Appliance (VBA). Running as a Linux Virtual Machine within the virtual environment, the VBA processes data rapidly and efficiently at the storage level through the hypervisor APIs so there is no need to purchase additional hardware. This contrasts with some approaches that need agents to be placed in virtual machines as well as products that require a physical proxy server. While workable, these approaches entail additional costs for hardware, OS licenses in some cases, and all the maintenance and complexity that is associated with agents and physical servers. Figure 1 illustrates the difference in these contrasting approaches.

FIGURE 1

Virtual Backup Appliance approach versus virtual backup requiring physical proxy server

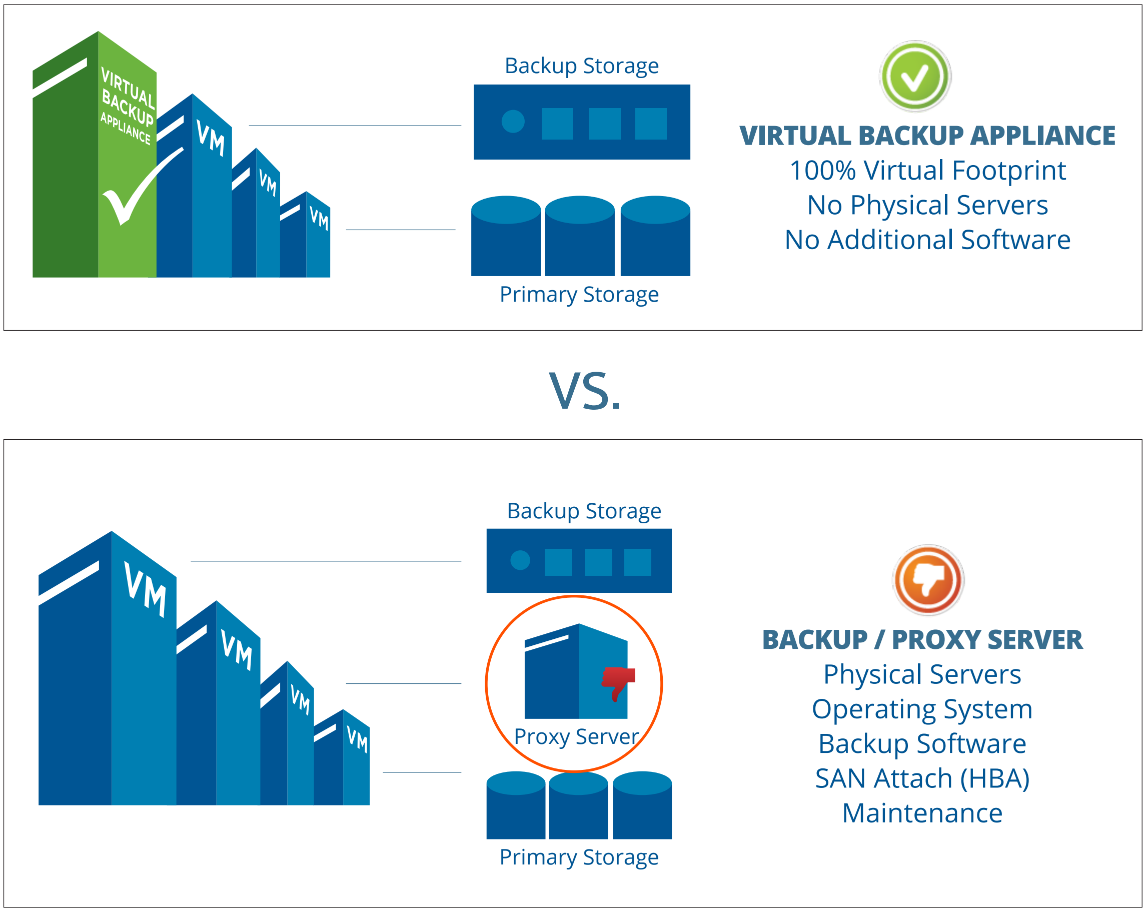

Look for backup and recovery solutions that support all major hypervisors

As individual virtualization vendors have come to dominate specific areas of the virtualization industry, such as VMware's dominance in enterprise production environments, IT professionals have begun to explore ways to avoid vendor lock-in and reduce costs. That means choosing backup and recovery solutions that support multiple hypervisors.

For example, an organization that relies on VMware to run its production environment may not want to duplicate that cost for its dev-test environment. Additionally, some new virtualization projects may not require all of the functionality that VMware provides, so IT staff can begin to look at lower cost hypervisors for specific applications. Citrix XenServer and Microsoft Hyper-V 2012 can offer substantial savings for organizations that have the expertise to manage these hypervisor alternatives.

In the immediate future, IT professionals will have more solution options that will combine the management of multiple hypervisors integrated into a single console for easier and simpler operation and maintenance

LOCAL DATA PROTECTION & RECOVERY

BEST PRACTICE # 2:

Incorporate deduplication to enhance recovery speed and reduce storage costs

In the last few years, deduplication has rapidly become a necessity for even small businesses seeking to enhance their local backup and recovery capabilities. That's because as the popularity of using virtualization has grown, so have storage needs accelerated. This is especially true in the age of "big data" where companies are seeking ways to exploit large amounts of information for insights on a daily basis. And with additional storage needs growing so quickly, costs have increased in proportion to the demand.

Deduplication helps to reduce storage demands in virtual environments by removing redundant data during the backup and recovery process, capturing and storing only unique data elements that have been added or changed since the last backup.

Incorporate global deduplication capability to improve efficiency of backup and recovery

To conserve storage it is important that virtual machine data be deduplicated at a global level. Deduplicating at just a virtual machine or job level still exposes you to writing duplicated data to the backup target, thus increasing storage requirements. Regardless of how backup jobs are organized, they should still be deduplicated across the entire dataset ensuring the backup data houses only unique data.

Source vs. target deduplication

Source deduplication is the removal of redundant blocks before transmission to the backup target. Source deduplication products reduce the bandwidth required to transfer information for backup as well as the amount of storage needed in the backup store. No additional hardware is required to back up to a remote site and many source deduplication products also support automation for offsite copies. Thus source deduplication minimizes the impact of backup and recovery on the network and reduces storage requirements.

Target deduplication by comparison is the removal of data redundancy at the backup target. That means that all data has to be transmitted to a hardware target first, before it can be deduped. While this method will reduce the storage required for backup, it does not reduce the amount of data that must be sent across a LAN or WAN during the backup process. However, many target deduplication solutions are powered by hardware appliances that can deduplicate at very granular levels. Therefore, in some cases, target deduplication can provide better storage savings – at the cost of using additional bandwidth.

Not all target deduplication solutions are the same. Some will deduplicate "in-line" prior to writing data to the disk. Others will perform the deduplication in a "post-process" manner, which leaves data untouched on disk for a period of time and then deduplicates that data as time goes by. Some post-process solutions are smart enough to deduplicate only blocks that are not accessed frequently to improve performance. Therefore, it is often the case that post-process deduplication will use more storage than in-line deduplication, but with less overhead and faster access for recovery, thereby showing better performance than in-line deduplication solutions.

Ensure your deduplication strategy supports your backup and recovery strategy

It may seem logical to put the most powerful deduplication device possible into an environment, assuming that it will give you the most efficiency. However, there are several factors to consider when choosing a solution.

First is cost. Software based deduplication is typically built into most backup software for free. As long as it provides global deduplication capabilities, software based deduplication can generally meet most needs. It also uses less bandwidth because it conducts source deduplication, which is becoming critical in many environments. In many cases, a backup product's own deduplication process can also recover faster than target side deduplication solutions.

Second, you will want to ensure your deduplication strategy fits your backup target. Global deduplication requires many random IO intensive transactions to occur as part of the backup process. If you want global deduplication, ensure that you have a stable target that can perform well. It does not need to be expensive. It should just consider that certain protocols and file systems perform much better at scale with these IO patterns. For instance, if your backup product leverages a virtual appliance, then anything attached as a virtual disk to that appliance will typically perform better and more stable with global deduplication than a typical network share.

Finally, if you require the granularity offered by target deduplication appliances, then be sure your backup configuration is optimized for that type of target. Deduplication at the target typically should not be mixed with source-side, global deduplication solutions. This configuration often exhibits significant performance issues over time. Therefore, it is best to configure your backup product to use large backup files when you are writing data to a deduplication device.

Choosing a backup product that can perform global deduplication on its own, as well as be configured to optimize target deduplication when needed, will provide the most flexibility for your strategy.

LOCAL DATA PROTECTION & RECOVERY

BEST PRACTICE # 3:

Use Volume Shadowcopy Service (VSS) for Windows

Volume Shadowcopy Service or VSS is a critical component to data protection, but it does not apply to all systems. It is vital when you are working with systems running Microsoft Exchange, SQL Server, or Active Directory, since these systems require special handling that is assisted by VSS. There are numerous published documents on how VSS works, but in simple terms, it provides you with the ability to perform application consistent backups. This is a Microsoft best practice that improves recoverability of these databases from backup. Failure to not leverage VSS can expose issues where the databases may not mount cleanly, or, even worse, could lead to data corruption.

Another benefit to VSS is the automation of log truncation. Each of these applications handles the truncation piece differently, but with the assistance of VSS it removes the need of having to perform this manually. Once a backup has been completed using VSS the database system would run through its own operations for executing the log truncation. This is especially important for Exchange because those log files could potentially fill up the drives and force the databases to dismount themselves without warning.

In most cases, Exchange and SQL require special functionality outside the scope of what typical hypervisor tools will provide. Be sure your backup product can handle this additional functionality for those applications.

LOCAL DATA PROTECTION & RECOVERY
BEST PRACTICE # 4:

Make sure onsite solutions deliver recovery across a data continuum

It's important to recognize that not all virtual backup products deliver the capability to assure dynamic granular recovery solutions. With IT professionals being asked to do more with less on a daily basis, it's essential that you be able to recover and restore across the full spectrum of situations including restoring a file, an application object, a virtual disk or an entire virtual machine depending on the circumstances and criticality of the business service involved.

Instantly recover a file or folder Your recovery solution should enable you to pull a single file or folder from backup storage without forcing you to recover the entire Virtual Machine. This obviously saves an enormous amount of time and effort.

Instantly recover an application object Your recovery solution needs to handle the most common recovery tasks such as finding and restoring a deleted MS Exchange email, or locating and restoring a missing SharePoint document. In fact, the backup product should be able to expose backup data instantly for any application that may have recovery management tools built-in. SQL Server is an example because its management tools are sufficient for browsing and recovering database objects so long as the backup product can quickly expose that database.

Quickly restore a Virtual Disk Most VM's contain multiple virtual disks. It saves time and effort if you are able to recover only the affected disk or disks on a particular machine.

Rapidly restore a full VM When restoring an entire VM, you want to be able to choose restoration options based on where and how the VM is being used in your environment. Incremental restores or rollback capabilities will allow you to restore a VM to a specific point in time in a matter of minutes by only recovering changes. In addition, if a VM that's part of a critical application is deleted for any reason, instant recovery capabilities would allow you run the VM from the backup copy until such time that you can restore the VM. Bottom line: You want to maximize your choices for recovery options by taking into account the need of that specific application.

Onsite data protection recovery scenarios

PHD Virtual provides leading backup and recovery solutions for virtual environments through our innovative Virtual Backup Appliance (VBA) technology. Supporting the three major hypervisors, VMware, Citrix and soon Hyper-V, PHD Virtual Backup and Replication software gives you the ability to optimize your data backup and recovery. Additionally, it can be extended with PHD ReliableDR to recover full applications, business services, and sites. Here are several recovery scenarios that illustrate PHD Virtual solutions for implementing new recovery rules in your virtual environment.

SCENARIO 1

A user permanently deletes an email and realizes days later that he needs to get it back ASAP

With PHD Virtual backup and recovery, you can quickly recover individual files and folders directly from backup storage without the need to restore the entire virtual machine image. For example, you get seamless search and recovery for Microsoft Exchange (and SharePoint, by the way) making it easy to locate specific items from those applications and recover them to files or directly back to the production app within just a few clicks. Recover emails, attachments, calendar events, contacts, and tasks from your PHD Backups in just a few minutes.

Powered by Kroll Ontrack, the leading vendor in granular application recovery and e-Discovery, you get super-fast search and recovery of specific items with no impact on your applications. So you can respond immediately to internal requests while meeting SLAs, and provide efficient collection of key information for investigations and compliance requests. Plus, you can recover items across different versions of the application to ensure upgrade compatibility.

PHD Virtual also simplifies file recovery for any Operating System, consuming fewer resources and providing unprecedented flexibility with file recovery by eliminating the need to allocate additional virtual appliances dedicated to the process.

SCENARIO 2

VM gets corrupted, and there is no snapshot available

Instead of having to wait for hours or even days to get your data back when a virtual machine is corrupted or crashes for any reason, PHD Virtual enables you to get that virtual machine back in minutes with Rollback Recovery. Rather than going through a lengthy, full Virtual Machine restore process, you can use PHD Virtual to simply roll a VM back to an earlier state by restoring only the virtual disk changes over the top of the existing VM.

Onsite data protection recovery scenarios

This helps to significantly reduce the time to recover a full VM—delivering recovery up to 100 times faster. Figure 2 below illustrates the process.

FIGURE 1

PHD Virtual Rollback Recovery

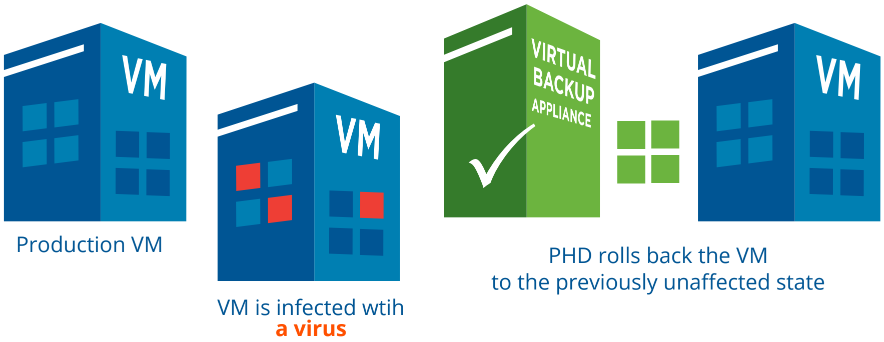

SCENARIO 3

Someone deletes a VM and now an entire application is not functioning

PHD Virtual Instant Recovery can replace a VM as fast as the time it takes to boot it from the backup target, getting any dependent applications back up and running in minutes. Your applications are instantly available without the need for additional infrastructure or a lengthy restore process.

Thus, PHD Virtual Instant Recovery enables you to eliminate costly downtime and meet SLA's by making an application available as quickly as possible. In the event of a failure, PHD Instant Recovery will simply:

- Turn on a VM using the data that resides directly on the backup target, enabling users to experience very little downtime.
- Then, once ready, users can either leverage VMware Storage vMotion or use PHD Motion to move the VM's data to production storage, and the VM remains operational the entire time.

To learn more, watch this PHD Virtual video demonstration:
http://www.phdvirtual.com/videos/instant-vm-recovery-phd-virtual-backup-vmware

CHAPTER 2

Offsite data protection & recovery

In Chapter 1 of Recovery Rules, we reviewed several best practices and use case scenarios for onsite data protection. Chapter 2 focuses on offsite data protection and recovery which is rapidly becoming a key component in data protection and recovery strategies for businesses of all sizes. The reason is simple. Data is growing at an accelerated rate making it essential for companies to manage their backup, recovery and storage much more effectively and efficiently. Advancements in offsite data protection for virtual environments help achieve that in a number of ways. In Chapter 2 of our E-Guide we describe the value of offsite data protection and recovery in the context of several best practices.

OFFSITE DATA PROTECTION & RECOVERY
BEST PRACTICE #1:

Always keep a copy of your backups offsite

While backing up data is a commonly accepted best practice among nearly all organizations, it's important to recognize the value of always keeping a copy of backups offsite as well. That's because hardware or software failures can corrupt data and you could easily lose all of your local backups in an outage or failure with no chance of recovery---a potentially devastating event.

Benefits of offsite data protection and recovery:

1. More reliable – Offsite backup provides for an automated way to protect your local backup data on a daily basis, or according to a desired schedule. Because backup copies are stored offsite in another physical location or in the cloud, you have "insurance" that critical data will be available for recovery if and when needed.

2. Reduces workload - Traditional offsite backup solutions require a substantial amount of time to manually back up or copy files to tape or removable disk and then transport them to another location. This is both time consuming and highly inefficient. Many offsite backup solutions can be initiated and maintained with a few clicks and automatically scheduled once the initial transfer of backup data is completed.

3. Secure – Storing data offsite is getting more secure, especially among cloud providers. To ensure that backup data is transmitted securely, most solutions use advanced encryption tools, and offsite cloud facilities are protected with the latest in network and host security technologies.

4. Saves money – Using tape or removable media storage solutions can be very costly as data storage grows, especially if you have numerous computers with large amounts of data that you back up regularly. Offsite backup solutions, in the cloud for example, are billed on a monthly-basis and can support a near unlimited amount of backup storage at relatively low cost.

Those IT shops that only perform nightly incremental backups along with weekly full backups are increasingly finding that their data — and the recovery requirements for that data — are forcing them to look at offsite backup alternatives.

1. Tape or rotational media - One of the most popular ways to store a copy of backups among small businesses has been tape or rotational media. It's the lowest out of pocket cost approach in many cases but typically requires lengthy manual processes for recovery and becomes cumbersome to manage over time.

2. Cloud storage – Copying backups to a public cloud, private cloud or a hybrid mix of the two, is growing in acceptance as smaller companies outgrow the traditional offsite tape storage model for copying backups.

3. Hosted/secondary site – Backup files can be copied and sent offsite to a secondary site. Host-based replication is the practice of using servers to copy data from one site to another.

Creating and managing a secondary site for storing backups can be relatively expensive for many smaller organizations. If you don't have a secondary site to store backups, choosing a cloud or removable media offsite backup approach can serve as your disaster recovery strategy. If for any reason a disaster strikes your primary location (a fire or flood) your data is protected and available for recovery. However, recovery times from offsite backup can be much slower than more robust strategies that involve VM replication, which will be covered in this E-Guide in a later chapter.

A typical storage tier strategy will put low-cost media offsite for backup copies and archived files, allowing you to keep long-term retention offsite in a more cost-effective way.

Offsite copies allows you to feel more confidence in using excess primary storage space for local backups and very small retention for fast local recovery and automated long-term retention

OFFSITE DATA PROTECTION & RECOVERY
BEST PRACTICE #2:

Leverage disk based options when possible

Despite advances in offsite backup and recovery technologies, many organizations continue to rely on tape or removable media storage for backup files. This approach has traditionally been the most "affordable" in terms of out of pocket expense for smaller firms to assure backup files are available for recovery. And changing your offsite backup strategy can seem like a hassle since existing habits require change. But as data growth continues to accelerate, maintaining tape and/or removable media for storing backups becomes more costly for companies in the long run due to reliability issues, manual steps to manage the process, and costs associated with physical data transport and storage of the offsite media.

In most cases, you can't verify that the data you've backed up is reliable and has not been corrupted. It's simply too difficult to test your backups. Another major concern is that recovery from an outage or disaster from tape or removable media is typically very slow, requiring days and even weeks to fully recover. Many businesses could not survive this kind of a "recovery."

Offsite disk based backup and recovery solutions for virtual environments, on the other hand, offer many more options that are more efficient and less costly.

Disk based backup and recovery processes are completely automated. After initial setup, IT administrators can quickly and easily backup and recover everything from full Virtual Machines to granular data such as emails. Disk based backups can be more readily tested to verify proper copying of files using features such as Instant VM Recovery and VM replication from backup data.

Options for disk based backup include hardware appliances that replicate backup data to a similar appliance in another datacenter. In the case of PHD Virtual Backup and Replication, you can use the built-in software archiving feature, which can copy backups and automate long-term retention enabled by its global deduplication structure. Direct cloud storage access is an advanced option to consider as well, even if it is just for a small portion of the environment. Today, many solutions are writing backup data to the cloud, including PHD Virtual. Just be sure to understand the recovery process. If you have to bring all backups back to local disk before recovering anything from the cloud, then you haven't improved on one of the main drawbacks of tape – slow recovery times.

Regardless of how you choose to use disk based backup, you should make sure your offsite solution supports automated long-term retention with deduplication across all backup sets. This will help you avoid excess requirements for offsite storage and reduce the bandwidth necessary to transmit backup copies offsite.   (See Chapter 1: Best Practice # 2 for an explanation of deduplication options)

OFFSITE DATA PROTECTION & RECOVERY
BEST PRACTICE #3:

Choose an offsite location for your backups that fits your recovery strategy

In general, offsite backups are viewed as a protection policy. However, it is essential that you consider how you intend to use those backups for recovery purposes.

As offsite data protection options become more accessible, more automated, and less costly, you should consider using them to enable granular search and recovery or full disaster recovery. When selecting an offsite location, you should consider whether you might need to search through your backups a couple years from now for specific email threads. Or, whether they need to be available to recover VMs at another site in the event of an outage. You might also want to ensure that if you lose the original backups, you can return the copies back to local storage for future recovery needs. Just make sure you've thought through how your offsite strategy matches your potential recovery needs. Here are some tips to help. explanation of deduplication options).

Online Backup Providers

Online backup providers can offer a complete service that includes the backup software, an offsite target, and sometimes a local target in the form of a physical or virtual storage appliance. This type of solution can help those users that want to minimize management of the backup process and use a fully integrated solution. In most cases, these types of services are not always robust in their virtualization functionality and recovery times can suffer as a result. That being said, the tradeoff is minimal backup management by the end customer. This type of service is usually good for very small companies with few IT resources that are focused on managing production applications rather than data protection. Some online providers are starting to offer recovery into their cloud for DR purposes as well. Disaster recovery options are discussed in Chapter 3 of this guide.

Public Cloud Storage

For smaller organizations with less than a few terabytes of data who cannot afford a secondary site to copy backups, public cloud providers such as Amazon, Google, Rackspace and many more boutique operations now offer secure offsite storage. With these services, you can select your own backup software and have it send offsite backups directly to the cloud storage. You do not need to manage the storage, but you do need to manage the cloud account and the backup process itself.

Cloud storage provides a great option for companies that do not have an offsite strategy in place today and can't afford the expense of a co-location facility or secondary site of their own. It gives you control over the backup and recovery infrastructure, and storage costs are relatively low compared with tape and removable media options. Major cloud storage providers such as Amazon and Google continue to lower their prices as they compete for business, and as a result, the public cloud option can significantly reduce complexity, management time and cost.

When to rely on your own secondary site

If data security or regulations prevent you from using online backup or cloud storage providers, then you need and probably already have a secondary storage site---hopefully in another geographic location. In most cases, you can also control bandwidth scale and usage to your own site, which allows for much greater scalability in data transport than the cloud options usually offer. If you don't own the datacenter yourself, there are a number of hosting providers that will rent storage and bandwidth as needed, or allow you to co-locate servers on their premises.

A secondary site is an ideal option for a number of situations:

- If you have a DR failover plan in place with your own site
- If you have regulatory requirements that prevent you from using cloud storage
- If you prefer to control and manage the entire process including the offsite storage

Details about setting up and using a secondary site are beyond the scope of this E-Guide, but keep in mind that solutions such as PHD Virtual Backup can allow you to use those offsite backup copies in ways that enable your application failover strategy. This is explained in Chapter 3 of this guide.

Onsite data protection recovery scenarios

PHD Virtual offers a powerful yet affordable offsite backup and recovery solution based on our unique Virtual Backup Appliance (VBA) architecture.

Our VBA is a lightweight virtual machine that installs in a few minutes and provides complete virtual and physical server data protection without dedicated backup hardware, software or requiring a complex integration project. Running on VMware, Citrix and soon Hyper-V the VBA is controlled from within the hypervisor providing fast and easy VM backups and file- or application-level restores. PHD Virtual can also provide an "Instant Recovery" which enables you to start a VM directly from the backup data store, and scaling involves simply deploying another VBA. The PHD Virtual VBA includes source-side deduplication and changed block tracking to reduce the amount of data handled and stored and supports concurrent backup and restore operations so that data recovery won't impact your backup window.

In 2013, PHD Virtual introduced its PHD Virtual CloudHook™ module that enables connectivity to a public cloud provider, currently supporting Amazon S3, Google Cloud Storage, Rackspace Cloud Files and providers using OpenStack/Swift or S3 Compliant storage. While some backup software providers have added "cloud support" they have done so by installing a separate appliance and writing to it as if it were local disk letting the gateway manage the cloud translation. This adds to your costs and complexity by creating another silo of storage to be managed.

Onsite data protection recovery scenarios

SCENARIO 1

User's backup storage fails and needs to replace all backups

In this scenario, the local backup storage has failed and all data has been lost and therefore cannot be recovered. Fortunately, you have copied backups to an offsite location.

Lost local backup data can occur any time.   And while you may not need the backups right away, you want to recover those lost backups for faster VM recovery in case of a future outage. So, in this scenario you want to copy backups from your offsite cloud location, back to your primary site as quickly and as easily as possible.

PHD Virtual makes this much easier with our Archiving feature. Simply use a VBA to backup to storage on-site and let an Archiving VBA replicate the backup data to the offsite location. If you lose your local backups, you can choose to reverse the archive process to bring back all or some of the backups that were kept offsite for faster recovery in the primary site if a future incident occurs.

SCENARIO 2

User needs to recover a granular bit of data from long-term retention offsite

In this scenario you need to access employee data for someone who has left your organization nearly three years ago, and this granular piece of data is no longer kept on local backup storage. Fortunately you use PHD Virtual Backup and Replication with its CloudHook module to archive backup data to a cloud storage provider.

With PHD Virtual Backup and Replication archiving you can easily retrieve the exact piece of data using the same easy to use process you would leverage for recovery from local backups.

In addition, the Virtual Full backup mode in PHD Virtual provides global deduplication and automated GFS retention, excellent capabilities for managing backup data over long periods of time. There is no need to run periodic full backups or manually delete backups, drastically reducing backup windows, storage usage and bandwidth.

With PHD Virtual you can leverage different retention policies for local and archived backups, making the most of their local storage capacity by keeping only one or two local backups for faster recovery. You can then store long-term archives on less expensive remote storage, including the cloud and readily retrieve data as needed. This archiving capability is very helpful for medical, insurance, financial, and cloud provider verticals because they have varying needs for long-term retention and recovery

Onsite data protection recovery scenarios

WHAT OUR CUSTOMERS SAY

"When it took our old backup product almost 24 hours to restore the server, we knew it was time for a change. With the help of our consultant, we found PHD Virtual Backup---a product that met all our needs for local and virtual backup and quick restores to and from the cloud," according to Brent Adams, VP, Finance & Technology, Pacific Mechanical Supply. "We are so confident of the results we've seen thus far that we back up our critical systems (Exchange, SQL and File Server) to Amazon S3 via PHD's CloudHook to ensure adequate protection in the event of an unplanned disaster."

SCENARIO 3

User needs to recover a full VM from cloud storage

In this scenario, a user has kept only one day of backups in local storage with the rest automatically archived to a cloud storage provider. Unfortunately, one of her VMs has failed and the system administrator needs to recover the entire machine data from five days ago to fully restore the application it helps to run.

This type of cloud-based recovery requires a software application that supports a "recovery-in-place" process. PHD Virtual accomplishes cloud recovery through a function we call "Rollback Recovery". With this capability, the user requests a restore which automatically identifies which changes have occurred between the requested recovery point and the current copy of data.

With this information, only the blocks that have changed need to be sent across the internet connection, not the entire VM, significantly reducing the amount of data transmitted. The VM can be restored in a matter of minutes versus hours or more with other approaches.

About the Authors

This E-Guide has been written by a panel of PHD Virtual experts based on extensive experience on the frontlines working with our customers in all types of businesses and virtual environments.

Joseph Noonan is a Sr. Product Manager at PHD Virtual. He has been with the company for nearly four years and led a strategic transformation in 2010 to retrofit its Virtual Backup Appliance to protect multiple hypervisor environments. Prior to PHD Virtual, Joe spent almost 10 years with Unisys leading the test engineering organization for its core mainframe business before transitioning to marketing enterprise servers. Joe also held various partner management responsibilities in his tenure at Unisys, including its OEM relationship with Microsoft. Joe has a bachelor's of science degree in Electrical Engineering and an MBA from Villanova University. He has spoken at various tradeshows and conferences including VMUG.

James West is the Sales Engineering Manager for PHD Virtual responsible for managing all of the SE's in the sales organization. Before joining PHD, James spent his time operating as a systems administration/engineer in various industries but started his IT career with the US Air Force. James provides expertise in various technologies such as VMware, Citrix, Windows Server, Active Directory and Exchange just to name a few.

As Sales Engineer at PHD Virtual, **Jim Noonan** has extensive engineering experience in both the hardware and software industries, gaining much of his experience at General Dynamics Information Technology as a Senior Design Engineer. Jim is a graduate of the Navy Nuclear Power Program and spent six years as a Nuclear Electronics Technician/Reactor Operator on a nuclear fast attack submarine.

Ian Jones is a Senior Product Specialist with PHD Virtual with over 16 years of experience working with Microsoft, VMware, Citrix and associated products, specializing in systems installation, management, migrations, consolidations, conversions, integration and architecture. His certifications include MCSE, CCEA, CCSP, VTSP, and VSP. Ian has worked as a systems engineer, pre-sales consultant, and manager for Motorola, ScriptLogic, and VirtualSharp until it was acquired by PHD Virtual.

As Sales Engineer at PHD Virtual, **Sarah Doyle** has extensive experience in VMware and Citrix technologies. Prior to joining PHD, she was an IT Project Manager for a manufacturing company. She is VCP and A+ certified for both VMware and Citrix.

As Sales Engineer at PHD Virtual, **Greg Potocki** is responsible for pre-sales demonstrations and working with prospects to ensure successful trial validations of all products. Greg has more than 10 years' experience working as a System Engineer specializing in backup and Disaster Recovery as a Service. Along with extensive virtualization expertise with VMware, Citrix and Hyper-V, Greg has also worked with many other technologies involving, Windows, Linux, Active Directory, Exchange, SQL and more.

To learn more, view the On-Demand Webcast Series:

Chapter 1: View Here

Chapter 2: View Here

Coming Soon in December!

CHAPTER 3: Failover, Failback & Recovery Testing

About PHD Virtual

PHD Virtual provides the best value in data protection and recovery assurance for virtual and cloud environments. More than 6,500 customers worldwide rely on its solutions that reduce the risks and costs of recovery, are easier to use and far more affordable than competitive alternatives. PHD Virtual has been transforming data protection and recovery assurance since 2006. For more information, please visit: http://www.phdvirtual.com/

PHD Virtual Technologies

MAIN	866-710-1882
INT'L	+1-267-298-5320
WEB	www.phdvirtual.com
EMAIL	info@phdvirtual.com

North America Headquarters
1880 JFK Boulevard, Suite 1301
Philadelphia, PA 19103