

# 10 mistakes

That Can Undermine Your  
VM Data Recovery Efforts

by Brien M. Posey

**ALTARO**

# CONTENTS

Introduction .....	3
#1: Relying Exclusivly on Hypervisor Snapshots .....	4
#2: Failing to Establish Goals for Restoration Granularity.....	6
#3: Forgetting to Protect Your Backups Through Replication.....	7
#4: Overlooking the Need to Perform Sandboxed Test Restorations .....	8
#5: Assuming that any Backup Application Can Protect Any Resource .....	10
#6: Neglecting the Importance of Storage Planning .....	12
#7: Not Having a Plan for Minimizing Recovery Time .....	13
#8: Ignoring Your Data's Rate of Change .....	15
#9: Not Keeping Track of Backup Health, and Backup Resource Consumption.....	17
#10: Selecting a Backup Product that is Overly Complex .....	18
Conclusion .....	19
About the Author .....	21
About Altaro .....	21

## INTRODUCTION

Although backup hardware and software play an undeniable role in data protection and recovery, an organization's ability to effectively recover data following a data loss event hinges on decisions that were made long before the data was ever lost.

Simply put, making mistakes at this stage can put your data at risk by diminishing the chances that data can be recovered when needed. Such mistakes may stem from cutting corners in an effort to reduce costs, or from not fully understanding the impact of data recovery decisions.

Regardless of the cause, the end result is the same – data restorations become more difficult or more time consuming than they should be, or recovery becomes impossible. As such, this paper discusses ten of the mistakes that are commonly made with regard to planning data protection.

# 1. RELYING EXCLUSIVELY ON HYPERVISOR SNAPSHOTS

One mistake that is sometimes made in SMB environments is that of relying on hypervisor snapshots as a backup substitute. On the surface, hypervisor snapshots would seem to be an excellent alternative to traditional backups. After all, snapshots allow an organization to instantly revert a virtual machine (VM) to a previous state, without the need for a lengthy restoration.

Additionally, snapshots do not require supplementary media such as backup tapes, and because snapshot capabilities are integrated into the hypervisor, snapshots do not require backup software to be used.

In reality, hypervisor snapshots are not a true backup replacement. Using snapshots for data protection can result in problems ranging from poor performance, to data corruption, or even data loss.

The reason why this is the case has to do with the way that hypervisor snapshots work.

Unlike a backup, creating a snapshot does not create a copy of the data that needs to be protected. Instead, creating a hypervisor snapshot results in the creation of a differencing disk. VM level write operations are redirected to the differencing disk, leaving the previously existing virtual hard disk contents in an unaltered state. Snapshots can be used to almost instantly revert a VM to a previous state, because the virtual hard disk contents remain in the exact state that they were in at the time that the snapshot was created.

There are several reasons why hypervisor snapshots should not be used as a backup alternative:

- **First, as previously mentioned, hypervisor snapshots do not create a backup copy of your data.** If the disk containing a virtual hard disk were to fail, a snapshot cannot be used to restore the virtual hard disk. In fact, the snapshot will often be

lost too in this situation, because some hypervisors store snapshots and virtual hard disks on the same physical volume.

- **A second reason why hypervisor snapshots should not be used as a backup alternative is because hypervisor snapshots tend not to be application aware** (although application aware checkpoints were introduced in Windows Server 2016). As such, applying a snapshot of a virtualized application server could result in data corruption.
- **A third reason why hypervisor snapshots should not be used as a backup alternative is because each time a VM snapshot is created, the hypervisor creates an additional differencing disk.** Differencing disk usage impacts the VM's storage I/O performance (especially read I/O). This happens because the hypervisor knows that the most recently created differencing disk contains the most recently written data. Hence, VM read operations read the most recent differencing disk.

One especially common example of an application that can experience problems as a result of the use of hypervisor snapshots is Microsoft Exchange Server. In fact, Microsoft goes so far as to say that “VM snapshots aren’t application aware, and using them can have unintended and unexpected consequences for a server application that maintains state data, such as Exchange. As a result, making VM snapshots of an Exchange guest VM isn’t supported” ([https://technet.microsoft.com/en-us/library/jj619301\(v=exchg.160\).aspx](https://technet.microsoft.com/en-us/library/jj619301(v=exchg.160).aspx)).

It’s fairly obvious how using snapshots could potentially damage a mailbox database, snapshots can even interfere with something as seemingly benign as an Exchange hub transport server. The Hub Transport Server is a legacy Exchange Server role that has since been combined with the Mailbox Server role. Its job was to pass mail through the transport pipeline. If a hypervisor snapshot were taken of a Hub Transport Server, it would capture the state of that server at a particular point in time. If the snapshot were later applied, it would likely result in numerous E-mail messages being resent.

If the requested data is not found, then the hypervisor reads the next most recently created differencing disk. The hypervisor works its way through the differencing disk chain until it either locates the requested data, or arrives at the original virtual hard disk. Hence, the more snapshots that exist for a VM, the longer the chain of differencing disks, and the worse the VM's read performance will potentially be.

## 2. FAILING TO ESTABLISH GOALS FOR RESTORATION GRANULARITY

A second mistake that is sometimes made by backup administrators is failing to establish goals for restoration granularity. It is not enough to simply be able to restore a backup. Administrators must be able to restore a backup in a way that fits the situation at hand.

Suppose for example, that a VM fails and needs to be restored. Restoring a backup of the entire host server would accomplish the goal of restoring the VM, but would not be appropriate to the situation at hand because a blanket host server restoration would also cause other VMs residing on the host to be overwritten.

In this example, granularity makes the difference between being able to recover an individual VM, and being required to restore an entire host server. Although this is an extreme example, the version of Windows Server Backup that was included with Windows Server 2008 could only perform host volume level recovery of Hyper-V servers, and was unable to restore individual VMs (<https://blogs.technet.microsoft.com/hugofe/2011/09/01/backup-hyper-v-virtual-machines-on-windows-server-2008-by-using-windows-server-backup/>).

When planning for recovery granularity, it is important to determine what levels of restoration might be required. In a virtualized environment, it is obviously important to be able to restore host servers and individual VMs, but there are also other degrees

of granularity that should be considered. For example, an organization may need the ability to restore files and folders within a VM, an application that is running on a virtual server, or an individual host server volume, or virtual hard disk.

Although most modern backup applications will allow data to be protected with varying degrees of granularity, it is often the way that backup jobs are configured that determines how data can be restored. For instance, it is relatively common for a backup application to require the backup administrator to explicitly specify that application level protection is required for a VM. Otherwise, the VM may be protected at the file level, but such a backup will not perform a VSS backup of the application that is running on the VM. The resulting backup would be crash consistent, but not application consistent.

Backup administrators should periodically review backup jobs to ensure that they provide the required degrees of granular protection. This is important because VMs tend to evolve over time, as do VM protection requirements.

### 3. FORGETTING TO PROTECT YOUR BACKUPS THROUGH REPLICATION

Another common backup related mistake is the failure to protect backups against accidental data loss. In order for data to truly be protected, there needs to be at least three copies of the data – the original data, a local backup, and an offsite backup. The local backup can be used to restore data more quickly than would be possible with an offsite backup, but the offsite backup guards against the loss of the local backup.

Suppose for instance, that a data center was destroyed by a fire. In such a situation, the production data would be lost, as would the local backup. An offsite backup would provide the only means of recovering the data.

Offsite backups provide redundancy that can also be useful in less extreme situations. Suppose, for instance, that a backup needed to be restored and the local backup media was found to be damaged or corrupt. In that type of situation, the offsite backup would provide an additional layer of redundancy that would allow the backup to be restored in spite of the damage to the local backup.

At one time, it was common for organizations to perform a nightly tape backup and then ship the tape offsite to a tape storage facility. Although this method did achieve off-site backup storage, it did so at the cost of convenience. A large insurance company where I used to work had a policy of shipping backup tapes offsite each morning. One day however, a data loss event occurred after the previous night's tape had already been shipped. It took several hours to get the tape back from the courier, and the courier service charged the company a hefty fee for an unscheduled tape delivery.

The bigger problem, however, was the costs incurred in the form of lost revenue and lost productivity as a result of the downtime. Estimates as to the hourly cost of a critical application outage vary widely, with some estimates going as high as half a million dollars per hour (<http://devops.com/2015/02/11/real-cost-downtime/>). In all likelihood, only the largest corporations incur losses at that rate. Regardless, the idea that financial losses are tied to the outage of critical systems, and that those losses are proportional to the duration of the outage are undisputed. Hence, it is critically important for an organization to retain a local backup copy that can be restored quickly, but to also have a redundant offsite backup copy.

## 4. OVERLOOKING THE NEED TO PERFORM SANDBOXED TEST RESTORATIONS

One of the backup best practices that receives the most attention is the need for backup testing. It's one of those things that gets drilled into the IT professional's head from day one. After all, the only way to know for sure that a backup is good is to test the backup to make sure that it can be restored.

Although the importance of backup testing cannot be disputed, it is important to take things a step further and test the backup in a sandboxed environment. For the uninitiated, a sandbox refers to a completely isolated environment in which testing can occur without fear of impacting the production environment.



The reason why it is so important to test backups in a sandboxed environment is because without a sandboxed environment, the ability to test a backup becomes very limited. Suppose for instance that an administrator needs to perform a test restoration of a VM, but does not have a sandboxed environment. In such a situation, the type of testing that could actually be done would depend on the type of VM that is being tested, and on the organization's configuration.

The administrator could likely test the ability to restore individual files and folders to an alternate location. The administrator would probably also be able to restore the VM in its entirety to a different location, but would not be able to boot the VM without interfering with the production environment. If the VM happens to be an application server, the administrator could probably go through the motions of an application restoration (to an alternate location), but would be unable to fully test the success of the recovery because doing so would require the newly restored VM to be booted.

It is important to be able to boot a VM as a part of a recovery test. Otherwise, it is impossible to verify that the backup is good. I once saw for example, a Windows Server on which the Winload.exe file had been deleted. Because this file is used only as a part of the boot process, the server continued to run normally. The server was backed up regularly, and the backup was successfully verified because everything on the server's hard drive had been backed up. The next time that the server was rebooted however, the boot process failed as a result of the missing file. Restoring the backup did not fix the problem because the backup did not contain the missing file. Had the backup been fully tested, the problem would have been found much sooner.

Savvy administrators sometimes attempt to get around the problem of test restorations interfering with the production environment by restoring the VM to an alternate host, and disconnecting the VM's virtual network adapter from the virtual switch. Doing so has the effect of disconnecting the VM from the network, which would allow the VM to be safely started.

Admittedly, this type of testing is far better than simply restoring a VM and not attempting to start it. However, this testing method isn't perfect. The vast majority of the virtual servers that are in use today have dependency services that must be accessible to the VM.

For example, a VM may have a dependency on the Active Directory and on an external SQL Server. Booting a test VM that has been disconnected from the network will allow the VM's operating system to load, but it does not permit the VM to access external dependencies. This will generally result in the failure of the VM's application. Microsoft Exchange Server for example, will fail to start critical system services unless the server can access the Active Directory. Similarly, Microsoft SharePoint cannot function without its SQL Server database, which is usually located on a separate server.

Worse yet, complex applications such as SharePoint often have compound dependencies. For example, SharePoint servers are often a part of a farm consisting of multiple SharePoint servers that must be able to communicate with one another. Similarly, the SQL Server on which SharePoint depends is often clustered, which means that multiple SQL Servers must be online and able to communicate with one another if the cluster is to achieve quorum and make the database available to SharePoint.

The point is that it is impossible to adequately test the ability to recover complex application servers unless a test recovery can be performed in a sandboxed environment.

As previously mentioned, a sandbox is completely isolated from the production network. As such, the VM of interest, and all of the VM's external dependencies can be restored to the sandboxed environment. This allows the recovery process to be fully tested without fear of impacting the production environment.

## 5. ASSUMING THAT ANY BACKUP APPLICATION CAN PROTECT ANY RESOURCE

A backup is really nothing more than a copy of a protected resource that can be used to repair that protected resource if necessary. Given this simple definition and the fact that backup applications have been around for decades, it may be tempting to think of

a backup application as a commodity utility, and to assume that aside from a few vendor specific bells and whistles that backup applications are all basically the same.

The idea that any backup application can protect any resource might have held true at one time. Today however, backup applications must be specifically designed to support the resources that need to be protected.

There are several reasons why the backup application must be matched to the resources that it is protecting. For example, a backup application that is designed specifically for use in Windows environments probably would not be able to protect Linux workloads due to the architectural differences between the two file systems.

Similarly, a backup application must be application aware if it is to be used to protect an application server. Without application awareness, a backup application would attempt to protect an application server by simply performing a file copy. The application would be crash consistent, but not application consistent.

Many years ago, I did a consulting job for an organization that had use a file level backup to protect its Exchange Server, rather than using an Exchange aware backup application. Following the restoration, the server's mailbox database would not mount because it was not in a consistent state. Fixing the problem required the use of low level utilities such as ESEUTIL and ISINTEG. Although these utilities were able to eventually put the database into a consistent state so that it could be mounted, the database repair process took an entire day and resulted in at least some permanent data loss.

The reason why application consistency is so important is because the application server is presumably in use at the time that the backup is being created. As such, making a file level copy of the database is not really an option, because the database contents are almost certain to change before the process completes. The end result would most likely be a corrupt backup.

To prevent these sorts of issues from occurring, backup vendors design their products to be application aware.

An application aware backup application knows how to use the Volume Shadow Copy Services to safely back up the application. The backup application acts as a VSS requestor and sends a backup request to the application server. The application-specific VSS writer that exists as a part of the application momentarily “freezes” the application database. This process places the database into a consistent state, and prevents any modifications until a VSS snapshot can be created.

The snapshot creation process occurs very quickly (so as not to be disruptive to the still running application). The VSS provider notifies the VSS writer that the snapshot has been created, and the VSS writer thaws the database to allow normal operations to resume. The backup is then created using the VSS snapshot.

## 6. NEGLECTING THE IMPORTANCE OF STORAGE PLANNING

Yet another mistake that a backup admin can make is neglecting the importance of storage planning. Backups and storage have always gone hand in hand, and storage planning is an important part of the backup planning process.

Storage planning needs differ based on an organization’s backup architecture. In most cases however, storage planning is synonymous with capacity planning. If disk-based backups are being used, for example, then the backup target storage array will need to have sufficient capacity to accommodate the data that is being backed up. Likewise, the storage array must support a sufficient number of IOPS to be able to efficiently support the backup process.

Of course no storage array has an infinite capacity, and given enough time, disk based backups will eventually fill a storage array to capacity. That being the case, an important

part of the storage planning process is determining the frequency with which data must be offloaded from the storage array in an effort to prevent its capacity from being fully consumed. The offload usually involves moving aging backup data to a secondary storage medium such as tape, a cloud storage gateway, or to lower cost higher capacity disks.

Determining a storage array's capacity for storing backups can be somewhat tricky. On the surface, it would at first seem that doing so would simply be a matter of multiplying the newly created (or modified) data by the number of recovery points created each day to determine the amount of data being backed up each day. The storage array's capacity could then be divided by the amount of data being backed up each day to determine the total retention time that the array can provide. In reality however, things are rarely this simple.

The volume of data that an organization creates each day is not necessarily the same from one day to the next. Similarly, the backup target likely uses block deduplication and other data reduction methods to reduce the data footprint. Consequently, standardized formulas for estimating backup storage tend not to work very well.

Administrators should talk to their backup vendor to determine what level of data reduction the software can realistically provide. It is important to keep in mind that data's ability to be reduced depends greatly on the uniqueness of the data. Administrators should also continuously monitor the volume of data being backed up, because data tends to grow over time.

## 7. NOT HAVING A PLAN FOR MINIMIZING RECOVERY TIME

Neglecting the need to plan for minimizing recovery time is an all too common mistake made by admins. If a data loss event or a system outage occurs, then there is usually a sense of urgency around the recovery effort. There is a real and tangible cost associated with an unplanned outage of a mission critical system.

Although cost estimates vary widely based on organization size and industry, a recent study of US datacenters estimates that an unplanned outage can cost \$7,900 per

minute. Obviously the costs are much lower in SMB sized organizations, but this figure serves to underscore the point that unplanned outages have a tangible financial impact regardless of the organization size.

Because there is a cost associated with systems outages, it is important to be able to perform restorations as quickly as possible. The age old practice of waiting for several hours for a tape to restore is simply not acceptable in today's world. Recovery operations should be measured in seconds or minutes, not hours or days.

**There are two main things that an organization must do to minimize the duration of a recovery operation:**

- **First, the organization should keep a backup copy of the data on premises.** Restoring from the cloud takes far longer than restoring from a local backup because of network bandwidth limitations. A 50-megabit connection to the cloud for example, can theoretically transfer 6.25 megabytes of data per second, which means that it would take roughly 2.7 minutes to transfer a gigabyte of data.

Several websites contain calculators that administrators can use to determine theoretical transfer rates to or from the cloud, for example: [http://www.convert-me.com/en/convert/data\\_transfer\\_rate/dmegabitps.html](http://www.convert-me.com/en/convert/data_transfer_rate/dmegabitps.html)

In the real world of course there are factors such as provider throttling, packet loss, and bandwidth contention that can result in much slower transfers.

Although speeds vary widely based on hardware, local recoveries occur much faster. Some LTO-3 tape drives for example, can perform compressed backups at a rate of 324 GB per hour (<http://www.quantum.com/serviceandsupport/softwareanddocumentationdownloads/lto-3drives/index.aspx>).

- **The other thing that an organization should do to minimize recovery time is to use disk based backups, and a backup application that includes an instant recovery feature.**

Instant recovery capabilities are based on the idea that because the backup is disk based, it is possible to mount a backup copy of a VM directly from the backup storage array, rather than performing a traditional restoration. This allows the VM to be made immediately available rather than the organization having to wait for a traditional restoration to complete.

In order to preserve the integrity of the backup, the backup software creates a differencing disk as a part of the recovery process. All write operations destined for the VM that is being recovered are redirected to the differencing disk, thereby leaving the backup contents unaltered.

Once the backup copy of the VM has been brought online, a traditional recovery begins running in the background. When this recovery completes, the contents of the differencing disk are merged into the newly restored VM, and the users are redirected from the backup VM to the recently recovered production VM copy. The end result is that the VM remains online and available for use throughout the recovery process.

## 8. IGNORING YOUR DATA'S RATE OF CHANGE

An easy mistake that backup admins sometimes make (or simply overlook) is that of neglecting the protected data's rate of change. The rate at which data changes can impact backups in several different ways.

First, the rate of change must be a primary consideration when establishing a Recovery Point Objective (RPO). The RPO refers to the frequency with which backup jobs are run. For example, legacy nightly backups have an RPO of about 24 hours, while modern continuous data protection solutions often have an RPO of five to ten minutes.

The greater the rate of change, the more data could potentially be lost between backup jobs. If for example, an organization's rate of change was 1 GB of data per hour, then a 24 hour RPO would risk the loss of up to 24 GB of data. Similarly, if the RPO were five minutes then the organization could potentially lose up to 85 MB of data in the event of a storage failure.



The rate of change must also be considered with regard to network and storage bandwidth and backup hardware performance. Suppose for example, that an organization decides to back its data up to the cloud. Let's also suppose that the organization's Internet connection can accommodate 1 GB of upstream data per hour. If the organization's change rate is 1.5 GB per hour (and no data reduction techniques are being used), then the cloud backup will not be able to keep pace with the demand due to an Internet bandwidth bottleneck.

The point is that under the right circumstances, data with a high rate of change can completely overwhelm a backup system unless the backup admin has anticipated the high rate of change and planned the backup accordingly.

The data rate of change is sometimes expressed as a percentage over time, but can be expressed in terms of data size over time. For example, suppose that an organization performs an initial full backup, followed by incremental forever backups. Let's also assume that the initial full backup was 10 TB in size, and that the incremental backups occur hourly and average 20 GB in size. The change rate in this example could be expressed as 20 GB per hour, or as 0.2% per hour (based on the initial backup).

Of course it is this is a simplified example. Factors such as compression and deduplication can impact the changed data's storage footprint. There is a good calculator at <http://wintelguy.com/backupcalc.pl> that can help you to estimate your change rate.

The other factor that must be determined is the amount of data that can be backed up over a given period of time

Once again, there is a calculator available at: [http://www.convert-me.com/en/convert/data\\_transfer\\_rate/dmegabitps.html](http://www.convert-me.com/en/convert/data_transfer_rate/dmegabitps.html). Using this calculator, you can see that a transfer rate of 45 megabits per second can transfer 20.25 GB of data per hour, thereby meeting the change rate requirement in this example. Keep in mind however, that the calculator reports an ideal result. The actual data transfer rate can be substantially lower due to resource contention, provider throttling, and other factors.



## 9. NOT KEEPING TRACK OF BACKUP HEALTH, AND BACKUP RESOURCE CONSUMPTION

One of the biggest mistakes that a backup administrator can make is to take a “set it and forget it” approach to backups. When a backup system is first put into place, the administrative staff typically subjects those backups to a lot of scrutiny. Backups are usually tested, and the administrative staff usually also reviews backup logs, looking for any hint of backup problems.

Over time however, it is easy to grow complacent. Reviewing backup logs becomes a boring and repetitive task, especially when the backups are always successful. Eventually, an administrator might stop reviewing the logs altogether.

The problem with complacency of course, is that it becomes far too easy to assume that backups are working perfectly, just as they always have. An administrator might not even realize that there is a problem with the backups until a restoration is required and the problem reveals itself. For example, I once saw a situation in which an organization’s backup administrator did not realize that someone had changed a service account password, thereby causing some of the backup jobs to fail. Nearly two weeks passed before the admin noticed the problem, because he had gotten out of the habit of checking the backup reports on a daily basis.

The best way of avoiding an unpleasant surprise during a recovery operation is to be aware of what’s going on with your backups. Review backup logs on a daily basis, and watch for things like failed backup jobs, access denied errors, and backup media that is being filled to near its maximum capacity.

Most backup applications include alerting mechanisms that can send E-mail messages to the administrative staff in an effort to make them aware of failed backups, low storage space, or other problems. However, it would be a mistake to rely exclusively on alerting mechanisms alone, because the alert messages might be mistaken for spam, or might arrive on a holiday weekend when no one is in the office to see the alert.

Alerts are an important part of backup monitoring, but they should be combined with other monitoring techniques. There are a number of third party tools available that can track resource consumption, and forecast storage and bandwidth depletion. Such tools are sometimes also able to parse event logs and provide insight to problematic conditions that might otherwise go unnoticed.

## 10. SELECTING A BACKUP PRODUCT THAT IS OVERLY COMPLEX

Finally, an all too common mistake admins make is purchasing a backup application that is excessively complex. After all, your backup software should help you to protect your network servers with a minimal learning curve. If a restoration should ever be necessary, the backup software should do everything possible to help you to get your data back.

A backup application that is overly complex may get in the way of the recovery process, for three major reasons:

- **Firstly, you really don't want to have to put off a much needed recovery operation so that you can figure out how the backup software works.** I have actually seen admins have to call a backup vendor's tech support to have them to walk the admin through a relatively simple restoration process, because the backup application's interface was so complicated.
- **Secondly, backup software that is overly complex can make it difficult to know for sure if you are protecting all of the resources that need to be protected.** Instead, a backup application should be completely intuitive and easy to use. A backup application should never leave an administrator guessing as to whether their network resources are being properly protected. **Finally, the need to recover data from backup is rarely planned** and a backup application needs to be easy enough for different members of an IT team to use and access, rather than getting stuck because an "expert user" of said application is unavailable.

## CONCLUSION

An organization's ability to recover from a disaster is often dependent on decisions that were made long before the disaster occurred. It is therefore in an organization's best interest to review its disaster recovery plans and plan and monitor its backup architecture to make sure that the organization is being adequately protected at all times.



If you're looking for a backup solution to protect your Hyper-V and/or VMware VMs that ticks the boxes for all of these best practices, have a look at Altaro VM Backup. It's designed to be easy to use, install and configure and you'll be up and running in a matter of minutes.

You get all the flexibility and functionality you need for a rock solid backup and recovery strategy, at an affordable price. Benefit from unbeatable value, outstanding support and a backup solution that protects Hyper-V and VMware VMs from a single console, across all your hosts.

## Altaro VM Backup - Trusted by over 30,000 SMBs

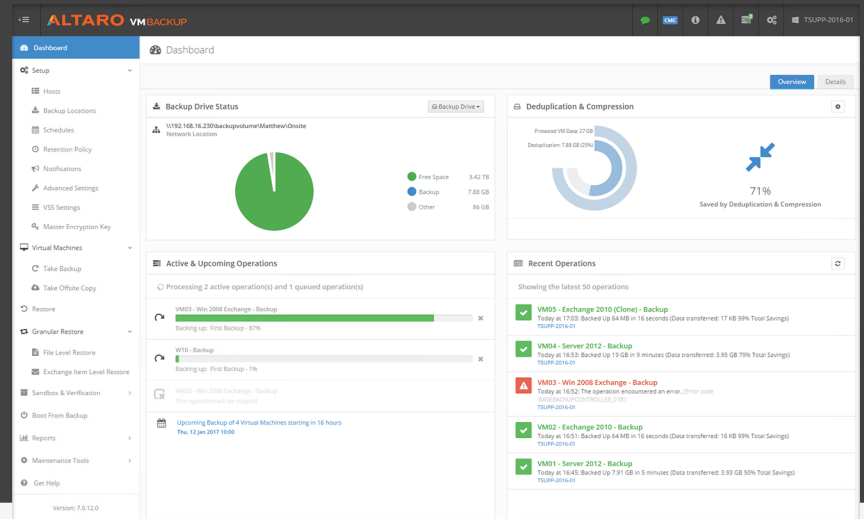
New v7! Altaro VM Backup for Hyper-V & VMware. Hassle-free and affordable VM backup software. Grab your free copy for 2 VMs now!

- ✓ Hassle-free and effective
- ✓ Unbeatable Value
- ✓ Outstanding Support

**Free for 2 VMs, forever.**

Back up unlimited VMs for 30 -days. After 30-days you get 2 VMs for free, forever. Download now!

**Download your  
30-day trial**



### Up and running quickly, without the need for complex configurations!

With Altaro VM Backup, you can install and run your first virtual machine (VM) backup in less than 15 minutes. Get up and running quickly, without the need for complex configurations or software dependencies.

Altaro VM Backup is designed to give you the power you need, without the hassle and steep learning curve.


- **Easy to use, intuitive UI** - making it easy to implement a rock solid backup strategy
- **Managing and configuring backup/restore jobs across multiple hosts** has never been simpler
- **Full control & scalability** - Monitor and manage all your Hyper-V and VMware hosts from a single console



**Virtual machine backup software packed with powerful features for Hyper-V and VMware.**

[View All Features](#)

## ABOUT THE AUTHOR



Brien Posey is a 14 time Microsoft MVP with over two decades of IT experience. Prior to going freelance, Brien worked as CIO for a national chain of hospitals and healthcare facilities. He has also served as a network engineer for the United States Department of Defense at Fort Knox and as a network administrator for some of the country's largest insurance companies. In addition to his work in IT, Brien is currently training to be a civilian astronaut. You can access Brien's Web site at <http://www.brienposey.com>

## ABOUT ALTARO

Altaro is a fast-growing developer of easy to use and affordable backup solutions for small- to medium-sized businesses, specializing in backup for virtualized environments.

Our aim is to delight our customers with full-featured, affordable backup software backed by an outstanding, personal Support team who are determined to help you succeed in protecting your environment.