

WHITE PAPER

Top 10 Best Practices for Windows Server Backup on Virtual and Physical Servers

by Don Jones, Microsoft MVP



Backup. Disaster recovery. Business continuity. Almost nobody *likes* to talk about them, but they're like car insurance: You *have* to deal with it. Given the absolute necessity of a backup plan, you might as well make the time to make yours a *good* backup plan – or even a *great* one. The best plans minimize downtime, minimize data loss, and minimize overhead to help backup become a robust, and even *appreciated*, part of the IT environment.

With that in mind, here are ten of the things you should be doing to make your backup plan the best it can be.

1. Backup to Your Cloud

The idea is to get your backups into a safe location, so that they'll be readily available when you need them – even if your primary storage fails.

Today, smart organizations are making sure that *their* backups include the cloud. More specifically, “their cloud”, whether that's their on-site server room, an off-site location, a hosted data center, or even public cloud-based storage. The idea is to get your backups into a safe location, so that they'll be readily available when you need them – even if your primary storage fails.

To make this practice work, you'll need a backup solution that's capable of replicating its own storage to your cloud – wherever that might be, and however it might be built. That replication will commonly be asynchronous, enabling it to minimize its bandwidth utilization while gently streaming your backup data to a safe place.

Your backup data gets replicated offsite, stored on fault-tolerant, highly-available, highly-scalable services.

2. Get Backups Offsite More Easily

The whole idea of “the cloud” is simple: Service providers build out an infrastructure larger than you could afford to do on your own, and distribute the cost of that infrastructure across its users. A backup solution that replicates to your cloud helps to protect your backup data by making it highly-available, fault-tolerant, and massively-scalable – all with less effort and risk on your part. A good backup solution will not only leverage whatever cloud you’ve built for yourself, it should also come with relationships with great cloud providers. That way, if you *haven’t* built yourself a “backup cloud” that’s as solid as you might want, you can get someone to help you do so.

That means building a backup plan that includes solutions which have deep relationships with providers. Get yourself some SAS 70-compliant offsite backup hosting, for example, to provide rock-solid, never-fail offsite storage – all without the hassle of rotating backup tapes. Your backup data gets replicated offsite, stored on fault-tolerant, highly-available, highly-scalable services. The guy from the backup tape storage company stops dropping by every Monday morning. And when you need to recover, your offsite partner either ships you your data, lets you start streaming recoveries directly from the cloud, or more likely a combination of both that will get your most critical services back up and running in a flash.

3. Recovery as Fast as Possible

IT often tends to focus on the side of backup that we see the most: The *making* of the backup. Fitting it into our maintenance window. Figuring out how long it will take. Finding faster tape drives.

But the most important part of backup is the *recovery*, which is when you actually use your backup investment. Recovery doesn’t wait for maintenance windows: It needs to happen *now*, and it needs to happen *pretty darn quickly*. Instantly, if possible. If you can get an entire server back up and running in the time it takes the CEO to make a K-Cup of coffee, you’re a hero. Tell the CEO that he’ll have time for a whole four-course meal, and you’re a zero.

So as you look at your backup plan, start with the recovery side, and work to make it as instantaneous as possible. Here's a hint: Spinning through backup tapes is probably not going to be the right answer. Instead, you'll likely be relying on modern disk-to-disk solutions that can stream data from faster online storage right to your servers.

In fact, the best backup solutions will actually let you resume a server *right from the backup file*. Sounds crazy, but why not? If your storage is already living on faster disk storage instead of slow, 1970s-technology magnetic tape, you should be able to resume critical services literally *instantly*, while data streams back to your production storage in the background. Tell the CEO to hold the cream and sugar, because he's going to be getting back to work faster than he thought.

4. Work While Recovering

Sounds like science fiction – or even fantasy – but the technology exists.

The problem with recovering data has always been its synchronous nature. Server goes down, everybody yells, and then everybody sits and waits while you restore the *entire* server. You're playing the backup tape shuffle, finding the latest full backup, loading on the last differential, applying incrementals. You're tapping your foot while the tapes spool, and your entire organization is waiting on you. Most IT pros just hide in the data center while it's all happening, rather than facing the out-of-work user population.

Your backup plan needs to deal with this situation by simply eliminating it. Ditch the tapes, and rely on a modern recovery solution that can provide *full user access* to servers *while they're recovering*. Sounds like science fiction – or even fantasy – but the technology exists. The right tools can actually detect the disk blocks that are needed to fulfill user requests, and restore those blocks first, handling user requests and getting servers back online in just *minutes*. Sure, the full restore might take longer to complete – but you'll be safe to leave the data center knowing that users are back to work.

5. Ready-to-Use Spare Servers

Users might even believe you when you tell them, “no, nothing happened, must have just been a temporary glitch in the network.

Virtualization is wonderful stuff. The ability to inexpensively keep an entire set of servers running on one physical host has already changed the way we build our infrastructures and host our services. Backup and recovery can benefit from this emphasis on virtualization, too.

Here’s the scenario: You figure out what your most critical production servers are. If that’s “all of them,” fine. Your backup solution – because you’ve taken time to select one that has this capability – continuously backs up those production servers, grabbing each disk block as it’s changed on disk. Those disk blocks are, of course, immediately streamed to a backup server for storage and use in recovery. But at the same time, those disk blocks are used to update a virtual machine image of those production servers. In other words, you’ve got a clone of your most important machines, and it’s being updated in near-real-time as the production servers themselves change.

When disaster strikes, it’s no disaster. Just unplug the production server and start its clone. Within a few minutes – however long it takes the virtual machine to start – the entire server is back online and running, with little or no lost work. Users might even believe you when you tell them, “no, nothing happened, must have just been a temporary glitch in the network.” Then you can sneak into the datacenter and start fixing the failed server, without anyone being any the wiser. That’s a recovery plan!

6. Eliminate Restore Errors

About the worst feeling at IT pro can have is that sinking one, in the pit of the stomach, when – while you’re trying to restore a file or an entire server – the dreaded error appears. “CRC error.” “Bad tape.” “Cannot complete restore.” Uh-oh.

Of course, we’re supposed to practice recovery from time to time, to make sure our tapes are good. And, of course, the tapes we *practice with* are just fine. It’s always the tape that’s at the end of a big, critical, this-is-not-a-drill recovery that has a problem.

And let's not just beat up on tapes, either: Plenty of disk-to-disk backup solutions can have the same problem. "Sorry, I missed that file during the backup – someone had it open." Grit teeth, take bad news to the boss. Practice all you want, you just can't have guaranteed recovery success.

Or can you? This is really just a software problem, right? So with the right software, you *should* be able to really guarantee that your backups will work when the chips are down. All it takes is an extra step to *verify* your backups, make sure they include *everything*, and make sure they *work*. That should actually happen continuously, so that you're assured of a solid backup when the time comes to use it. You should be able to run tests to verify the integrity of the backup data, *and* rely on the backup solution itself to proactively verify its own health by checking for corruption in the backup archive files or database.

7. Recover Anywhere, Anyhow

In each scenario, the original server keeps running in production while you perform the restore, giving you plenty of time to test, repeat, and anything else you might need.

A physical server dies. A virtual host just can't handle another virtual machine's workload. The data center got hit by a meteor. Whatever the cause, sometimes you simply don't have the option to restore a server to its original location, whether that was physical or virtual.

Your recovery solution needs to support this scenario by allowing you to restore any server to anyplace you want. Maybe you're recovering to a spare machine to get services back up and running quickly. Maybe you're restoring physical machines to virtual ones. Maybe you're restoring to a completely different, off-site datacenter due to a natural disaster. It shouldn't matter. Your recovery solution should act as a kind of middleman, making any potential recovery target a sort of "universal receiver," letting you get back up and running with whatever's at hand.

There's a side benefit to this capability, too: The ability to easily migrate between physical and virtual machines with almost no risk. Need to scale up that SQL Server computer? Just restore it to a bigger, faster machine. Consolidating servers?

Just restore them to a virtual machine. Have a virtual machine that's gotten too big for its britches? Restore it to a standalone physical machine to give it more resources. In each scenario, the original server *keeps running in production* while you perform the restore, giving you plenty of time to test, repeat, and anything else you might need. When everything's working fine – perhaps during a short maintenance window – you can "cut over" to the restored server.

Simply attach to the operating system's file system and capture every single low-level disk block write operation as it happens.

you point it out without any regard for whether or not some of that data might *already have been backed up*.

In other words, it's time to get your hands on some de-duplication magic. After all, the goal isn't to just back up every bit of data, right? It's to make sure that *every bit of data can be recovered*. Well, to recover something, you only need to back it up once, and by only backing it up once you'll save yourself a ton of space in the backup system.

A modern backup solution can do de-duplication on the fly. "Hey, you asked me to back up this file – and I've already got an exact copy of it in the archive. I'll just remember where this additional copy came from, and we'll call it even." A *really* good solution will compress that de-duplicated data, as well, so that it takes up less space in the backup than it did in production, saving your even more space.

9. Applications, Not Just Files

Simply attach to the operating system's file system and capture every single low-level disk block write operation as it happens.

Face it, when legacy backup technologies were invented – in the 1960s and 1970s – all we really cared about were files. Files are easy. Make a copy of it on tape, and you're good to go.

That unfortunately hasn't been the case for a couple of decades now. While we still have lots of files floating around the environment, we're also focused on *applications*. SQL Server. SharePoint Server. Exchange Server. Whatever. With those modern server applications, it isn't so much about the files as it is about the *entire server*. The program files. The databases – all of which have to remain synchronized with each other if they're going to work. Most of these applications keep their files open continuously, meaning you can't just grab the file.

It's time to leave that legacy approach where it came from: In the past. Today, you need to image the *entire server*, including every bit of data that's stored to disk. You need the executables. The DLL files. The operating system files. The databases. And yes, whatever simple files might be lying around as well. The best way to do that? Simply attach to the operating system's file system and capture every single low-level disk block write operation *as it happens*. With that approach, it doesn't matter if a file is open or an executable is running, because you're not looking at files! You're looking at the changes being committed to disk, capturing those, and storing them in a backup archive. You gain the ability not only to recover the entire server quickly, but to also rebuild any file, database, or piece of data – to *any specific point in time*.

It's a pretty amazing concept – but it's the smarter way to build a backup plan.

10. Eliminate Backup Windows

Let your backups run all day, and all night, continuously capturing changes to data, to files, to databases, and everything else.

The whole idea of backup windows is archaic, harkening back to a time when people actually went home at night and forgot about work. Those days are long gone, and today's organizations are 24x7x365 – and they'd work longer if they could somehow wedge a few extra hours into each day. Taking servers offline just so you can back them up is beyond unacceptable.

Backup windows also provide a false sense of security. After all, a backup window means you're not backing up your servers *until the window opens*. What about all the stuff that happened during the day, before the backup window started? What if a server crashes, or a file gets modified, or a database gets corrupted *just before* the backup window? Guess you'll just have to write off all that lost work.

Or, you could just backup smarter. *Continuously*. Close the backup windows. That's right, just forget about them. Let your backups run *all day, and all night*, continuously capturing changes to data, to files, to databases, and everything else. As changes happen, in nearly real time, you capture those changes and save them in a backup archive. When something goes wrong, you're not missing everything since the last backup window – you're missing *nothing*. You can bring a server, a file, a database, or whatever back to life to any given point in time – whether it was last week, or just a second ago.

Are You Ready for a Smarter Recovery Plan?

Imagine a recovery plan that enables near-instantaneous, live recovery of entire servers – or nearly-instant recovery of a single piece of data. Imagine pushing your backups to the cloud, enabling an array of recovery scenarios up to and including a complete datacenter failure. Imagine being able to spin up virtual clones of your servers at any time, restoring functionality without a byte of lost data. Imagine a recovery system that detects and repairs errors on its own, so that you're always assured of a successful restore.

Now stop imagining, and drop by AppAssure.com to download a free trial of a modern backup solution that provides exactly these capabilities – and much more. Forget the backup windows. Lose the tape drives. Leverage virtualization technologies and recover from almost any disaster in just minutes. It's all real, and it's waiting for you.



Don Jones

Microsoft MVP

Don Jones is one of the world's leading experts on the Microsoft business technology platform. He's the author of more than 35 books, and is a top-rated speaker at technology conferences worldwide including Microsoft TechEd and TechMentor, and writes the monthly for Microsoft TechNet Magazine.

He also writes features and monthly columns for numerous print and online publications, is a multiple-year recipient of Microsoft's prestigious Most Valuable Professional (MVP) Award, and serves as Editor-in-Chief for Realtime Publishers.

AppAssure®

#1 BACKUP AND REPLICATION TO YOUR CLOUD

About AppAssure Software

AppAssure, the global leader in complete server, data and application protection for virtual, physical and cloud infrastructures, delivers customer-proven backup and replication software that assures the recovery of applications in minutes. AppAssure's groundbreaking technology uniquely guarantees instant and 100% reliable application recovery from your server to your datacenter – to your cloud. AppAssure goes beyond protecting data, to protecting entire applications enabling service providers and enterprises to adopt a cloud model to deliver fast, reliable and secure data protection. With 3,450% growth over the past three years and more than 5,000 customers, partners and service providers in over 50 countries, AppAssure is the world's fastest growing backup and replication software company, ranked by Inc. Magazine.

© 2012. AppAssure Software. All Right Reserved.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

www.appassure.com