



The 7 Deadly Sins of Backup and Recovery



7 Technology Circle
Suite 100
Columbia, SC 29203

Phone: 866.359.5411
E-Mail: sales@unitrends.com
URL: www.unitrends.com

The 7 Deadly Sins of Backup and Recovery

This paper guides IT executives and system administrators in designing and implementing business backup and disaster recovery systems that enable every organization to rapidly recover entire business systems including business data, server OS, applications, and user data in case of a system failure or disaster, without experiencing undue planning and management pains.

Your organization's servers, systems and data are the lifeblood of your business, and maintaining the continuity of these business systems is a complex, and full time endeavor. There are often hundreds of servers or sub-components that can fail, and many more ways that environmental and human factors can trigger such failures. It is impossible to predict in which manner the next failure might occur. But it can be surprisingly straightforward to implement appropriate protections that will limit the occurrences—and the impact-of the inevitable failures.

This whitepaper outlines the most common missteps that organizations make in setting up disaster recovery or business system continuity plans. Smart IT executives will ensure that none of these mistakes is made in their organizations.

While there can be no guarantee of perfect success, avoiding these common mistakes vastly improves the odds that an organization will successfully sidestep the worst of the possible ramifications from the inevitable failures that will occur in any IT environment.

Any one of these mistakes has the potential to introduce fatal failures into a company's business system continuity plans. Let's explain each mistake in detail, and offer prescriptions for avoiding the traps they represent.

Backing Up Only The Data

It used to be that backing up only user data was the norm for virtually every organization. The operating systems' environments were stable and reliable (and there was also essentially no practical way to restore an operating system unless it was fully rebuilt). Organizations got in the habit of thinking of user data as the only volatile part of the environment, the only part requiring protection.

There's also a practical reason for backing up only the data. With data growing typically 50% per year, backup windows are shrinking to the point where there's only time for the data.

But today, there's significant risk at the OS level. Patches and updates are layered to the point that a full rebuild is close to impractical, and fraught with risk. This assumes that one can find the optical media, locate the license keys, and apply patches in the right sequence.

Today's data protection processes should include not only backups of user data, but also backups of the OS layer and all applications, along with the ability to quickly restore and recover each.

Fortunately, today's technology not only makes this practical, but straight forward. Disk-to-disk backup systems with "bare-metal" capability provide significantly faster backups, and reduce the backup time window enough that all critical information can be protected.

Bare-metal technology allows taking a full snapshot of the operating system and all applications, and allowing it to be fully restored to a new, "bare-metal" server in a matter of one or two hours, instead of several days.

More sophisticated systems provide this bare-metal capability across multiple operating systems, so organizations that are running Windows, Linux, and other OS's can continue to use a single integrated backup solution.

Regardless of the operating environment, it's critical to remember that everything needs protection, not just user data.

Allowing Backups To Go Untested

Organizations often spend an enormous amount of time making backups (weekly masters, nightly incrementals, and so on). If the backup volumes being created cannot be restored on a reliable basis, the process has effectively failed. The basic rule of data protection and disaster recovery systems is that they are unproven until they have been fully tested, and shown to be effective, on a frequent and ongoing basis.

The typical organization checks backup schedules to make sure that they are correct; checks that the backup jobs ran; and checks error logs to be sure they ran to an apparent successful completion. With those confirmations in hand, we assume that data recovery will be possible when necessary.

But media degrades over time; operating environments change in ways that make it unlikely that previous backups can be successfully restored; and the typical error rates associated with the media we use make it at best uncertain that the restoration will be successful.

A solid disaster recovery plan must include redundant backups intended to adequately compensate for normal error rates, and must incorporate time factors that reflect real world data from actual test restorations.

Why do most organizations not fully test their backups? The answer is simple: in an environment where backup windows are already too compressed, with too little time to actually *make* the backups, there is inevitably also too little time to test them.

A surprising number of organizations still use tape as their backup medium. The use of tape requires an enormous amount of time to build, test, and validate a backup. Why? Because simply restoring the data from tape isn't enough.

The IT staff will have to start with a new server, locate the CDs necessary to install the operating system, registry and other environmental components; re-install all necessary applications; and then finally restore user data from tape. Only if each of these steps works correctly can the organization be confident that key data can be successfully restored in a true disaster. This is often not the case.

The cure for this deadly sin is to have both the right technology, and to implement best practices. The right technology means using disk-based backup systems, incorporating modern "bare-metal" backup technology.

It's well understood that disk-based systems are faster, provide random access, and are vastly more reliable as a data storage medium than tape systems. "Bare-metal" technology allows a full snapshot of a system's operating environment—including the operating system, all applications, as well as complete user data—and allows that environment to be restored in a fraction of the time necessary to rebuild it from scratch. Using the right technology sets the stage for successful restorations.

The final element is to implement today's best practices for data capture and restoration testing. These include performing a full bare-metal restoration of each critical server—an Exchange e-mail server, an e-commerce server, and so on—no less frequently than four times each year. Careful organizations also conduct a random test of a file or folder restoration on each major server at least twice each year. In a company with ten major servers, that means a total of sixty tests per year, or five each month. This is only practical with sophisticated, disk-to-disk, bare-metal protection.

Best practices also include the concept of testing *something* each month. Operating systems and environments change rapidly enough that it's possible that the normal flow of changes in a 30 day period may have compromised something in the current backup resources. This suggests test restoration of at least some user data—and preferably an entire server—once each month.

Finally, an important best practice is to capture and analyze the data that the organization gathers about error rates of the test restorations, and the time required to conduct them, and feed this knowledge into its disaster recovery plan. When the plan is properly informed by real world test data, management and IT staff can have a far higher level of confidence that the plan on which they are depending will actually deliver when the chips are down.

Lack of Adequate Recovery Planning

Most IT organizations have some level of a “plan” that they’ll reach for in the event of a serious IT failure or natural disaster. But these plans often merely describe the systems that exist today, as opposed to laying out a roadmap for a successful recovery from a range of different possible emergencies.

At a minimum, a good disaster recovery plan should be based around the minimum number of different systems or technologies that are required to back up the entire company’s IT infrastructure. The more complex and numerous the tools that the organization will need to recover after a natural disaster, the more lengthy the plan will be, the more difficult it will be to update, maintain and test; and the less likely it will be to be followed correctly when needed. For this reason, smart organizations try to standardize around one or two backup recovery systems/technologies. Systems like Unitrends’ support more than 25 operating systems, simplifying both planning and recovery in the event of a disaster.

In a real disaster, it’s enormously helpful to have resolved priorities and the optimum actions and sequence far in advance, put together in an easy checklist. Albert Einstein said “Everything should be made as simple as possible, but no simpler.” This is important advice for disaster planners. Avoiding extraneous information and low-priority issues will increase efficiency and effectiveness of recovery in the event of a disaster.

Plans should be clear and specific about the organization’s recovery sequence and priorities in the event of various *different kinds* of disasters. Those responding to a disaster will obviously take very different steps after the failure of a single critical server than they will after an entire site disaster.

The more these choices can be identified and evaluated in advance, the simpler the execution process will be when rapid action is required. Other best practices associated with the planning part of the process include:

- > One person should be ultimately responsible for the execution of the disaster plan, and that person should have one backup person in the event that he or she is unavailable. Ownership of the different parts of the plan, the roles of the individuals involved, specific accountability, and specific communication requirements should all be spelled out in detail, in advance.
- > Just like basic backup and restoration testing, disaster plans themselves must be tested frequently and completely. Subtle interactions between components, or missing elements, can often only be identified by virtue of a literal simulation of each kind of disaster, with teams following a checklist exactly and tracking what works, what doesn’t, and how long each step takes.
- > Testing should focus not just on whether the plan can be executed successfully, but also on finding ways to simplify the plan for the future.

Finally, remember to keep multiple copies of the plan itself in multiple locations, just as you would data archives or backup tapes. Organizations are occasionally tripped up by having access to everything they need for a full recovery, except the plan itself, which is in a binder back at the office that was hit by the hurricane.

Not Planning For a Dissimilar Recovery Environment

We think of laptops, desktops and servers as nearly universal devices, and often don’t consider the risks and complexities of moving software and data between different brands or models. Also, we often assume that if we have a failure and need to shift to new hardware, we’ll be able to quickly acquire whatever we need.

But those who have tried know that incompatibilities often arise when moving to a new platform—even one that’s nominally fully compatible with whatever came before.

System driver, application and patch complexities and incompatibilities can undermine the best laid plans. It is critical for IT professionals to think about these issues, and plan for realistic restoration choices in the event of a disaster.

There are an increasing number of software systems today that advertise “baremetal” system recovery; some assert that backups taken from one hardware manufacturer can be restored to a server from a different manufacturer. This is called “dissimilar restoration,” and many who have tried it now adhere to this best practice.

But marketing claims aside, it is often the case that what worked in the lab, fails to work in the real world. Each component needs to be tested, in your own operating environment, to ensure success in the event of a disaster.

Another of today's choices is what is called "dislike to known" restoration. This refers to the ability to do a bare-metal restoration to a different kind of hardware than the source was taken from, but only to a pre-defined set of destination hardware devices. This works reliably because drivers for those devices are already embedded in the backups themselves.

Another of today's best practices is to have the restoration software provide the user with the opportunity to load drivers early in the bare-metal restoration process. This can often expand the range of destination hardware devices that can be used, by allowing the IT professional to supply the necessary drivers, instead of having to have selected them months in advance.

Without the ability to perform a successful bare-metal restoration, the user must build the entire new system first, by loading the operating system, applications, and then data—just as in the case of a data-only tape restoration. This is clearly not desirable. In a true disaster, time is of the essence, and a full rebuild like this wastes precious time...in many cases, days.

Today's careful planners will catalog all their hardware and operating system environments in advance, including:

- > What current machines are available from their preferred manufacturers
- > How today's generations of machines are different from the generation currently in use
- > Advanced identification of replacement sources for the servers that they use, and have documentation of what will be ordered in the event of a disaster.
- > Confirmation that their "dislike to known" hardware models are supported by their bare-metal software technology
- > Duplicate CDs of all necessary drivers, both stored onsite and at an offsite disaster recovery location
- > Testing, to be sure that all of their company's software technology—and all of the hours that were spent making backups—lead to a successful restoration in the event of a disaster.

Not Having Offsite Copies

Software security threats—viruses, spyware, malware, zero day exploits—often grab headlines and the attention of IT execs. But what is often forgotten is physical security, which can also disrupt a company's operations and undermine its backup and recovery plans. Surprisingly, many companies do not have a formal process for regularly taking backup copies to a remote location. Those that do may use an assistant's car trunk for tape storage and tape rotation...obviously not a suitably rigorous and reliable approach.

An increasingly large number of firms must store information offsite to comply with legislative requirements. Whether rooted in Sarbanes-Oxley, the Gramm-Leach-Bliley Act, or other recent legislation or regulations, there's an increasingly complex and comprehensive set of requirements for formal (and often offsite) retention of key data.

Organizations think of offsite storage as protection against true natural disasters—hurricanes or other events that can physically destroy a building. While there's ample recent evidence of this risk, the benefits of offsite storage cover a multitude of less dramatic but equally damaging potential problems.

For example, simple water damage from an air-conditioning drain or a leaking supply line can destroy a backup set in minutes. Burglary is a less-frequent but equally powerful risk. Loss of one's only backup disks due to a burglary can destroy an organization's information safety net.

Finally, we hear all too frequently of disgruntled employees who gain enormous power over an employer by destroying or holding hostage a crucial set of backups. Again, this risk can be mitigated easily by maintaining current copies offsite.

The logistics of keeping data offsite are much simpler today. Using removable disk drives as an archiving medium reduces the physical space required for offsite storage; today's systems typically provide fully integrated support for archiving as a part of the overall backup process. Removable disks have a small "form factor" and are highly stable.

In fact, today's hard drives are thought by most experts to be far better than burning DVDs as a backup medium. CDs and DVDs have well-documented problems with surface oxidation that can render them unreadable much more quickly than once thought. For many firms, removable disks are now the archive medium of choice.

Best practices involve keeping five successive weekly master backups offsite at all times, as well as end of month archives of the full system. Careful organizations keep bare-metal images of all their critical servers offsite as well, on similar rotation schedules.

Confusing Replication and Vaulting

The idea of replicated disk storage originated in the fail-safe systems of 20 years ago. In this technology, two disks or disk arrays simultaneously stored exactly the same data, such that if one had a failure, no data was lost.

Along the way, it occurred to disaster recovery professionals that replicated disk systems could be used for off site storage by inserting a Wide Area Network between the two disk arrays. Today, replication is typically used by SANs and "data mirror" products. Replication uses near real-time, block-level communications to constantly synchronize data between two different locations.

The advantages to this method are near real-time data protection, and ease of configuration and maintenance. Once established, a replication environment requires little ongoing attention.

The disadvantages are that the synchronization occurs at a very low level, below the ability to see data and file system errors or corruption. This means that the inevitable errors (some of which are potentially fatal to a restoration) that occur naturally over time will be immediately replicated to the offsite location...and the hoped-for protection is thereby immediately lost.

Replication is also highly resource-intensive. These systems create a substantial burden on the local CPU, since most replication technologies utilize a type of RAID engine to capture the block-level changes. These are implemented through software, imposing significant CPU overhead.

This sort of replication is also extremely bandwidth-intensive. As every data change for every transaction is moved across the network to the second location, this approach can impose a huge burden on a corporation's network infrastructure.

An alternative, preferable approach to offsite data storage is called "vaulting," which uses a less real-time technology, but allows a level of file system integrity not found in a replication environment. Vaulting operates at the file level, where the built-in operating system checks and balances on file system integrity are allowed to operate. This ensures that data stored in the offsite environment is accurate, complete, and able to be used for recovery in the event of a disaster. Additionally, the resource requirements are minimized, since the vaulting is provided by a secondary dedicated system, not the active servers themselves.

Today's best practices suggest that IT professionals should look for a vaulting system that moves only *changed* data, not the entire backups or original data source; this minimizes bandwidth requirements and maximizes backup windows. Look for dedicated, special purpose hardware to support your vaulting solution, preferably hardware and software that are fully integrated by the manufacturer, with optimized communication layers to keep the vaulting efficient.

In the case of either replication or purpose-built offsite vaulting, be sure to check compatibility with your local network infrastructure, including switches, routers, NICs and other components—there are hidden complexities in virtually all networks, and these can cause failures with either approach.

Finally, as with virtually every other disaster recovery plan component, test relentlessly. Only successive and consistently successful tests can provide the necessary level of confidence that the backup systems will be there in the event they're needed.

Adopting Inflexible Solutions

The only certainty in backup and disaster planning is that the post-disaster environment will be at least somewhat different than the one we expected and planned for.

This puts a premium on solutions that incorporate natural flexibility. For example, the ability to restore to a virtual machine environment gives the IT leader intrinsic choices. Having a range of known hardware configurations that a bare-metal backup can be restored to (as opposed to needing the exact hardware type of the failed server) does the same.

Pick a business system continuity solution that accommodates the majority of the needs you've already identified, plus others that may not be in today's plan but could still be useful. The more options you have in a real disaster, the better the chances that you'll be back in business quickly.

Conclusion

Business systems have evolved into critical elements of conducting normal, day-to-day business. Backup and recovery systems for an organization should not only be flexible and robust enough to restore user data quickly and accurately, but also entire systems and applications to prevent minimal downtime. Avoiding common pitfalls along with careful planning, testing and redundancy joined with the correct technology should play a critical role in your recovery from the inevitable system disaster.

All trademarks are the property of their respective owners.

About Unitrends

Unitrends offers a family of affordable, all-in-one on-premise backup appliances that support virtual and physical system backup and disaster recovery via disk-based archiving as well as electronic vaulting to private- and public-clouds. Unitrends is customer-obsessed, not technology-obsessed, and is focused on enabling its customers to focus on their business rather than on backup.

For more information, please visit www.unitrends.com or email us at sales@unitrends.com.