

Backup and recovery considerations

Learn how to properly back up your data and safely restore it



Contents

Recovery misconceptions	3
Plan with the restore in mind	3
Know the requirements of true backup	7
Follow recovery best practices	15
Conclusion: Find true backup and recovery	17
Appendix A: Backup and recovery scorecards	18
Appendix B: True backup at a glance	19

To put together a good plan, you need to know the risks and how to mitigate them. Most backup approaches don't protect you in a real disaster, putting your valuable data at risk. Your last line of defense should be a true backup approach that protects all your data, not just a subset.

Recovery misconceptions

Challenges to data security and business continuity haven't slowed down. Cyberthreats, extreme weather, and war are a reminder of how crucial backup and recovery are. As these threats grow, so do the number of vendors who offer new backup and recovery solutions that make big promises—promises that don't always reflect true backup and recovery practices. Their products and services may sound enticing. But they could result in vendor lock-in, require you to purchase more hardware, and still leave you vulnerable to data loss.

That said, approaches that fall outside true backup have their benefits—they just serve other purposes. They don't provide what you need to meet regulatory requirements or build a last line of defense when catastrophe strikes.

Our proven experience in backup and recovery has taught us that following the fundamentals is your best chance in the event of an attack. True, not every organization faces the same risk. But living in uncertain times means every organization should be cautious.

Here are three elements to consider as you develop your backup and recovery strategy:

1. Plan with the restore in mind.
2. Know the requirements of true backup.
3. Follow recovery best practices.

Each section of this white paper is devoted to one of these considerations. By the end, you'll understand our perspective on how to thoroughly back up your data and safely restore it.

Plan with the restore in mind

To put together a good plan, you need to know the risks and how to mitigate them. Your last line of defense should be a true backup approach that protects all your data, not just a subset. We'll define the requirements for true backup in [section 2](#).



Risk mitigation								
	True Backup*	True Backup* 2 Locations	Data Copy (including snapshot copy management)	Data Mirror	Data Replication	Hope and See	Array Snapshots	
Risk								
Ransomware Attack	✓	✓	⊖	⊖	⊖	⊖	Limited	Data encrypted
Data Corruption	✓	✓	⊖	⊖	⊖	⊖	⊖	Data corrupt
Potential Compromised Data	✓	✓	⊖	⊖	⊖	⊖	⊖	Data corrupt? Something trustworthy there to compare/check?
Inside Job	✓	✓	⊖	⊖	⊖	⊖	Maybe	Data corrupt
Human Error	✓	✓	Limited	Limited	⊖	⊖	Limited	Data loss. This is currently the most common threat to your data!
HW Component Failure	✓	✓	Limited	✓	Limited	⊖	⊖	Data loss on device, some trustworthy source for restore existent?
HW Firmware / SW Code Error	✓	✓	⊖	⊖	⊖	⊖	⊖	Data corrupt on multiple similar devices/locations
Search & Seizure	✓	✓	⊖	Maybe	Limited	⊖	⊖	Loss of control / access over HW from 1 or multiple sites
Legal Proof Required (Prior Art)	✓	✓	⊖	⊖	⊖	⊖	⊖	Trustworthy historic data samples available for restore?
Site Closure / Local Riots Etc.	✓ remote restore	✓	Maybe remote restore	✓	Limited	⊖	⊖	Loss of access to at least 1 site, potential data corruption there
Site Power Outage	✓ with external vaulted backup sets	✓	⊖	✓	Limited	⊖	⊖	Loss of access over HW from 1 site, potential data corruption there
Wide Range Power Outage	✓ with external vaulted backup sets	✓ with external vaulted backup sets	⊖	⊖	⊖	⊖	⊖	Loss of access over HW from multiple sites, potential data corruption there
Site Flooding/ Storm/Ice...	✓ with external vaulted backup sets	✓	⊖	✓	Limited mirror has to be in different location	⊖	⊖	Loss of access over HW from 1 site, potential data corruption there
Migration/ Upgrade Failure	✓	✓	Limited	Limited		⊖	Limited	Data corruption on device, some trustworthy source for restore of previous (older) data set existent?

Risk mitigation						
True Backup*	True Backup* 2 Locations	Data Copy (including snapshot copy management)	Data Mirror	Data Replication	Hope and See	Array Snapshots
Data copy (including snapshot copy management):		Snapshots or binary copies stored on mounted volumes on similar/identical hardware devices				
Data mirror		Data read and mirrored on schedule to similar device different location				
Data replication		Data is constantly instantly replicated. Once replication is stopped only 1 (or 2 different!) data set exists. Data replication synchronously/asynchronously replicates any mistakes made. While replication to some extent helps you to avoid loss of current data, data protection helps you to recover any data. Data replication synchronously/asynchronously replicates any mistakes made. While replication to some extent helps you to avoid loss of current data, data protection helps you to recover any data.				
True backup: see definition “true backup” in Chapter 2						

Table 1. IT risks and mitigation options

The ability to restore to any point in time depends on how often you perform backups. Recovery point objective (RPO) is especially important for file system backups. For application backup, the ability to back up log files is critical to RPO.

Follow these steps when planning

Here’s how to plan with restore in mind:

- See what is required for regulatory compliance: For example, GDPR, HIPAA, KVKK, PIPEDA, CCPA, notifiable data breaches act and many more require data loss prevention.
- Classify your data according to a limited set of categories.
- Understand the requirements of stakeholders, insurance companies, and so on.
- Accept the fact that all critical data is in scope, no matter if it’s in the cloud, off the cloud, or part of a hybrid environment.
- Define appropriate retention policies for ongoing backup and consider regular long-term backup sets for legal proof and integrity checks.
- Define the number of backup copies to keep and how to distribute them across locations and different media (for external safe storage and integrity checks).
- Keep it simple—recovery must work in panic mode. “Simple” might require you to keep more data sets to bypass some obstacles.
- Know the potential fees (legislation, ransomware, business on hold) if restore fails.
- Prepare for change, reduce hardware and software dependencies, and avoid multistep approaches.

In short, you should define restore service-level agreements (SLAs) based on requirements and specific infrastructure setups. These SLAs should include the recovery point objective and recovery time objective.

Recovery metrics

Your plan should account for two critical recovery metrics:

Recovery point objective

A recovery point objective (RPO) is the maximum acceptable interval during which transactional data is lost from an IT service.

The ability to restore to any point in time depends on how often you perform backups. RPO is especially important for file system backups. For application backup, the ability to back up log files is critical to RPO. A sophisticated RPO definition automatically tells you the recurrence (scheduling) of your backup jobs.

Recovery time objective

The recovery time objective (RTO) is the amount of time a business process must be restored after a disruption to avoid a break in business continuity.

RTO mostly defines the performance needs of recovery. The faster the restore, the sooner you are back in business. Setting your RTO may take some testing because you may not know how fast the infrastructure is. Any backup and restore performance depends on the underlying infrastructure—server, storage, and network. And guess what? If you can recover fast, your backup is good as well.

RPO and RTO help you plan

RPO and RTP are key. They make sure information can be recovered in a timely and accurate fashion. Defining RPO and RTO also highlights prerequisites for backup. Often, data protection designs only consider backup. And when business-critical data needs to be brought back, there is no way to compensate for slow, ineffective recovery. Test your recovery procedures regularly.

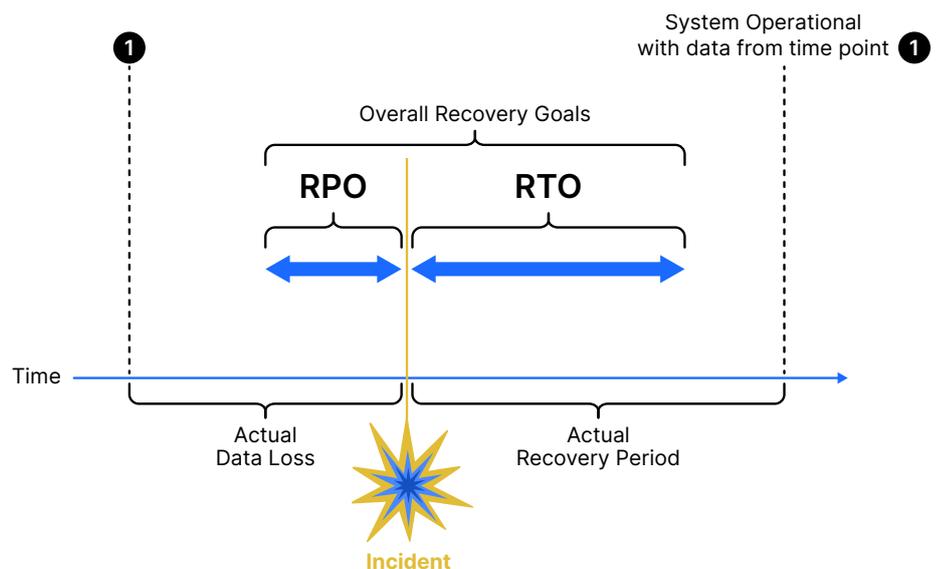


Figure 1. Illustration on RPO and RTO

Before we focus on how to design your backup environment it is important to reiterate that there is no shortcut to restorability. In case of a disaster, you need proper backups.

Know the requirements of true backup

Before we focus on how to design your backup environment it is important to reiterate that there is no shortcut to restorability. In case of a disaster, you need proper backups.

True backup

The word backup can mean different things to different people—and vendors. But a true backup must:

- Follow the 3-2-1 backup rule (more on that later).
- Obey the policy for different media in two categories:
 - **Logical:** Use of backup software with its own data format (transportable).
 - **Physical:** Backup to dissimilar media—for instance, tape or to the cloud.
- Ensure that copies are distributed to at least two secure, separate locations. Storing tapes externally is also a good choice.

Snapshots don't count

A snapshot describes data that is frozen at a certain point in time. In fact, snapshots are often implemented by tracking the changes on the primary storage. However, a snapshot is not a backup—only an additional component in a data protection strategy.

A snapshot only works for restore if the original volume still exists and blocks from the snapshot can be recovered to a consistent volume. That's why you need an independent backup. It will help you if a storage system is gone, or you need to restore to a different infrastructure.

For virtual machines (VMs), a VM-snapshot is a binary clone of the VM at best. So it is missing the logical transformation to make the data independent from the specific architecture.

Neither do data copy and snapshot copy management tools

Be wary of tools that 1) create data copies that stay on mounted disk volumes and 2) store data copies or snapshots without proper format change. Any such approach cannot be considered a backup. Even if analysts endorse these tools, they don't provide reliable restore options nor the required data integrity guarantee. They also can't protect your data from being repurposed, stolen, or used against you.

Be wary of tools that
1) create data copies that stay on mounted disk volumes and
2) store data copies or snapshots without proper format change.
Any such approach cannot be considered a backup. Even if analysts endorse these tools, they don't provide reliable restore options nor the required data integrity guarantee.

Don't rely on your system's backup capabilities

It is a common misconception that some primary storage and system architectures don't need additional backup. Here's why that's not the case:

- Any binary clone on similar hardware will always leave you in doubt if the data is actually readable and uncompromised.
- Any replication or mirroring will require you to trust the implementation and risk losing the data entirely if something goes wrong or the setup is compromised.
- Any array or platform-specific replication or mirroring will require you to trust the implementation, stay with specific building blocks, and may even lock you into pricing, shipment, replacement, repair, and so on. You also risk losing the data entirely if something goes wrong or the setup gets compromised. For instance, if a firmware update fails, you will need a backup of the data.

Select the backup target architecture

The most common approach to backup is the 3-2-1 rule. It requires you to have at least three backups of a given data set. These three copies shall be distributed across at least two separate locations, with one backup being offline. To follow this rule, backup target technologies must be able to handle volume and performance requirements, with budget constraints considered.

Although the 3-2-1 rule is the most common approach, it has critical shortcomings that this paper will address later.

Disk-based targets

Modern backup architectures usually incorporate disk-based targets, where multiple backups of a given data set are stored. If the target device fails, the backup versions or copies stored there are affected, if not lost. So it is mandatory to keep another backup copy on a different target. This could be another disk-based backup appliance that either uses some form of replication with the initial target or receives a backup mirror or a dedicated backup copy.

The disk-based targets usually come with the benefit of easy random access for fast single-object restore. And high deduplication rates allow for storing a lot more recovery points without wasting a lot of additional resources.

From a risk and availability perspective, keep backup copies in different locations that are protected from fire, unwanted access, and weather-related threats.

Whenever a disk-based backup target is selected, it should differ from a mounted volume. Preferably, it can hide behind an API so that backups cannot be compromised easily. Backup sets on mounted volumes are the primary target for many ransomware and malware attacks. But they are also prone to be compromised by other IT or human mischief.

Backup targets with a solid immutability feature provide enhanced security from unwanted altering.

While some data backup might require a backup proxy server or gateway server for security or connectivity reasons, the most resource-efficient, secure backup will be a direct transfer from the original data source to the backup target.

Does the following sound familiar? “Backup software needs integration with XYZ and only works with third-party component ABC.” It could be a mandatory database you need to procure, manage, and troubleshoot or a selected backup target technology for best results.

Direct backup, dissimilar targets, and format change

While some data backup might require a backup proxy server or gateway server for security or connectivity reasons, the most resource-efficient, secure backup will be a direct transfer from the original data source to the backup target.

Often, solutions incapable of direct backup require “staging” or “parking” of the backup data on mounted volumes, which puts your data and your RPO SLA at risk.

One standard approach is to store backups on dissimilar media from the original data source. In case there is an issue with the primary media—such as firmware error, malware, controller error, and so on—you still have uncompromised copies. Another advantage of storing backups on dissimilar hardware is that the data is checked for readability and consistency, not just binary cloned.

Readable data allows you to add integrity checksums, change the data format to either target, and—even better—change to a specific backup format that is independent of both source and target technologies. Backup and recovery software that uses a specific backup data format instantly translates data that it reads, transports it in that specific format, and stores it on target in the specific format.

The most important advantages of having a specific backup data format are:

- Independence from source and target platforms architectures
- Independence from past, current, and future technologies
- Enhanced integrity protection
- Protection of backup data from repurposing by other IT procedures (mandatory for GDPR)
- Protection from unallowed access and data theft
- Transportability independent from the data path, connection technology, and intermediate systems
- Additional features of the backup software itself or integrated APIs
- Ability to search, browse, and restore backups after everything but the backup media is gone.
- Protection from misuse—only the backup software can read the data

The right choice of backup devices should also consider performance. That's both individual device performance and the way it handles parallel data streams. This is relevant to RTO since your SLAs are useless if you can't restore them in the required timeframe.

Dependencies

After selecting a backup and recovery software with a specific backup data format, the next step is to avoid remaining dependencies.

These dependencies include:

Limited availability

Backup software is available only on one operating system and therefore is bound in its core processes to the limitations and critical issues of that operating system.

Vendor Lock-in

Most backup software is designed to work with specific building blocks that lock you into one vendor or specific architecture. Lock-in will limit your ability to work around issues or walk away without paying a steep price.

Given the sheer amount of backup data to hold, vendor independence and freedom of choice for backup targets is key to keeping costs down, evolving with technology improvements, and selecting multiple target architectures for their specific advantages.

Integration issues

Does the following sound familiar? "Backup software needs integration with XYZ and only works with thirdparty component ABC." It could be a mandatory database you need to procure, manage, and troubleshoot or a selected backup target technology for best results.

Scale

Backup software has limited scalability in its internal structure, processes, and database. Solving this issue requires multiple instances with substantial resource requirements, just for handling the data volume.

Hardware

In their marketing and advertisements, some vendors leave out one detail of their solutions: They require a huge pile of hardware to be fully functional. The worst scenario is the requirement to keep backups on mounted volumes. Even if it was for only 10 percent of your backup data, the cost to operate this environment will be huge, as will the carbon footprint.

Plus, any backup data parked solely on a mounted disk—even if only temporarily—is not truly backed up and can't ever satisfy an RPO. In other words, there is no reliable restore source, which leads to a high risk of being manipulated, deleted, or corrupted.

Backing up data on edge devices often comes with specific requirements that need expert management. As long as these devices are always on, have a stable connection, have sufficient bandwidth, can be trusted, and can segregate relevant data from irrelevant data with few surprises and bearable redundancy—you can rely on the guidelines in this paper.

The German Bundesamt für Sicherheit in der Informationstechnik (BSI) describes the following mandatory tasks within the IT organization to protect from ransomware disasters: "The following measures MUST be implemented within an IT infrastructure from the BSI perspective... Regularly perform multilevel data backups, especially offline backups. A backup always includes the planning of the restart and a test of the restore of data."

Source: ACS - Maßnahmen zum Schutz vor Emotet und anderen E-Mailangriffen (allianz-fuer-cybersicherheit.de)

Performance

The right choice of backup devices should also consider performance. That's both individual device performance and the way it handles parallel data streams. This is relevant to RTO since your SLAs are useless if you can't restore them in the required timeframe. Network-based backup and recovery are easy, but on many occasions, they're not fast enough. You may want to consider FC SAN since it's usually available in larger storage infrastructure and is more dedicated to high I/O transfers without many side effects. LANfree backup and recovery are more than just options. They may be the factor that determines a successful recovery—or going out of business.

Online backup, offline backup—so what?

Offline and online can have vastly different meanings in backup and restore. Here is some further clarification by definition:

Online backup

Online backup usually means that a backup can be performed while the system is still online and working, with little impact. In a situation where you're backing up applications within a running VM or container as a snapshot, you should be aware that some data is in the app server's RAM and not written to disk. This portion of data might not get protected properly, and a restore may be partially unsuccessful.

Therefore, an online backup solution must keep applications consistently running and ensure application users are not disrupted. An application integration agent would know the current state of the application, initiate backup mode and redirect ongoing application changes to dedicated areas.

Once the backup is done, data is merged back into the usual data sets. In a recovery scenario, data is first placed back into the app server. Then log files are restored and used to recover the app to a given point in time.

Online backup integrations for a list of common applications or databases can include:

- The options to trigger both backup and recovery from within the application's administration.
- An easy way for the backup administrator to set up the backup
- An easy way for the backup administrator to recover the application back to a very specific transaction or time stamp without requiring the application expert at all.

A different meaning for online backup is a backup of a given data set that is available for instant restore or an instant mount. After all, any backup on a mounted volume can be considered online.

Expert knowledge is required to perform the common tasks around safeguarding your data for the worst of all circumstances. Limit user privileges to the very few experts that fully understand the implications and dependencies of I/O intense backup and restore operations in environments of limited capacity, concurrency, and bandwidth.

Offline Backup

Offline backup technically describes a backup that is performed while the service holding the data is not available, like an application being shut down and all outstanding I/O has been written to disk. In that case, a standard file system backup can take away that data and store it on whatever device is available or selected.

In current discussions, the term refers to a backup of a given data set that is unavailable, at least from outside the backup application. In other words, it can't be accessed as easily as a mounted volume.

In common sense, any backup behind an API—like a deduplication API—qualifies for an initial level of “offline-ness,” as long as the actual bits on disks are protected from manipulation. A disk media that gets unmounted and powered off would be a further improvement. But it's seldom used, due to the many disadvantages of such technology.

Any backup copy in the cloud is not considered offline. Such cloud backups are a common target for malware and ransomware.

A backup on tape is a good example of a high level of “offline-ness,” as tapes usually get unmounted fast to make way for writing and reading other tapes by the tape drive. The levels of “offline-ness” for tapes further increase by features of the tape robot to vault such media or, even better, to entirely eject the tapes and allow for storing them in an actual vault in a different location.

In the past, a tape target or WORM media provided many means of dissimilarity. But this caused other inconveniences: e.g., tape providing very limited random read access. Today these technologies have further evolved and provide benefits such as speed, cost-effectiveness, and advanced offline features.

More backup considerations

There's more to think about than backup targets. You'll also want to take the following into account:

Licensing model

Your backup software will have great insight into the characteristics of your environment and data. Sooner or later, the backup software vendor might want to capitalize on this insight by imposing costly licensing changes.

While several vendors provide a simple licensing model according to the volume of protected data, others have not. Some backup software vendors even change their licensing model every few years, requiring you to pay extra.

Choose a vendor with a stable, simple licensing model. Find one that offers multiple licensing models and the freedom to choose the less expensive option for your setup.

Beyond encrypted control communications, you don't want anyone to read your backup or restore data being transferred. Some backup tools repackage the data stream into their specific backup data format to make it unreadable, but others do not. Either way, you need the option to securely encrypt the backup data at the source before transmission or en route before it leaves the protected wires.

User management

Effective user management for the backup and recovery tool, as well as the backup target devices, is mandatory.

Expert knowledge is required to perform the common tasks around safeguarding your data for the worst of all circumstances. Limit user privileges to the very few experts that fully understand the implications and dependencies of I/O intense backup and restore operations in environments of limited capacity, concurrency, and bandwidth.

As restores have to work in panic mode, with parts of your IT being down, you should avoid additional dependencies that might get in your way before a restore operation can be triggered.

Don't be surprised by the privileges required on systems for the actual backup and restore operations. Reading and writing data, along with access rights and metadata, including OS restores or application recovery, require extended levels of privilege.

Relationship between components

All components of the backup software must have a trusted relationship. They need to know each other and instantly identify rogue components that might interfere with your backups and restores. This is usually done by exchanging and checking certificates before granting any access. Hackers know that backups are a honey pot for them. So, you need to keep them out.

Encryption

You don't want anyone listening to what's going on in the backup and recovery environment, to avoid sharing any sensitive data like credentials, backup targets configured, and so on. Therefore, any control communication within the backup environment has to be secure and encrypted.

Beyond encrypted control communications, you don't want anyone to read your backup or restore data being transferred. Some backup tools repackage the data stream into their specific backup data format to make it unreadable, but others do not. Either way, you need the option to securely encrypt the backup data at the source before transmission or en route before it leaves the protected wires.

Encrypting backup data comes at the price of poor deduplication and compression as well as compute power overhead. Use it very carefully, only where it is unconditionally required.

Most backup target devices allow for encryption at rest—which the backup application should support. It comes in handy for external safe storage of backup data, as long as you keep full control over the encryption keys and prepare them to be available after disaster strikes.

Spend your budget wisely. You will need every dollar for proper core backup and recovery architecture and operations. Only pay for true backup with secure restorability, and nothing else.

Account for architecture changes

One of the most important considerations is that your IT architecture is and will be in constant change. Local as well as widespread outages, or at least service level derogations, will impact IT processes that require lots of I/O bandwidth on different paths, along with CPU intense operations—namely, your backup and recovery operations. So your backup application must allow for temporary or persistent changes again and again, without extensive readjustments and without impacting restore capabilities for previous data sets.

Similarly, the backup application should allow for easy troubleshooting or tuning options to adhere to RPO and RTO goals. It should also avoid the need to throw more and more hardware on a bottleneck.

Documentation

Documentation is key. Recovery plans define what services need to be alive (and how to achieve that) before restores can be initiated. You need to preserve configuration details to rebuild the backup infrastructure after a substantial disaster.

Operating centralized backup

By now it is clear that a central backup for all your valuable data is one of the most complex, I/O- and compute-intensive processes in your IT. Even still, it's typically the best detector for any planned or unplanned changes and service interruptions.

Here are three tips for running a centralized backup:

- 1. Implement real-time operational intelligence**, which shows you what's going on and at what success rate. Use tools such as customizable dashboards and reports, root-cause analysis, scenario-based modeling, and predictive analytics for balancing resources and identifying and resolving potential conflicts.
- 2. Don't rely on your backup application to detect threats.** While a backup solution will most likely notice suspicious changes in your data and your IT landscape, this detection would come too late. By relying on your backup application to detect any manipulations, you would make your last line of defense your only line of defense, all while putting it at risk. Detecting intrusions or manipulations should be your second line of defense, integrating pattern analysis from both core and edge before data is altered. Educating employees and business partners along with proper identity and access management should be the first line.
- 3. Spend your budget wisely.** You will need every dollar for proper core backup and recovery architecture and operations. Only pay for true backup with secure restorability, and nothing else.

Backup and Recovery Maturity Level	Description
Worst	<ul style="list-style-type: none"> • Backup and recovery are left to third parties. • There is no true backup, only mirroring, data copies, and replications. • An island approach to data protection introduces vulnerabilities. • The backup tool runs on the Windows platform, which is the primary target for hackers and ransomware gangs.
Baseline	<ul style="list-style-type: none"> • Business-critical data is regularly backed up. • Testing ensures the RTO and RPO of the restore. • Backup copies exist in two or more locations, with at least one backup copy being offline.
Better	<ul style="list-style-type: none"> • One set of backups is stored in a disaster-proof environment yet is accessible by experts. • There are secure communications and a trusted relationship between critical employees and all backup infrastructure components. • Full backups and one-stop restores exist.
Best	<ul style="list-style-type: none"> • Backups of all relevant data assets are done as often as possible with hardened core components. • Backups are stored in a secure, immutable, and transportable format. • Restores can be performed without restoring the malware too.

Table 2. Backup and recovery maturity levels

Follow recovery best practices

Now with secure, reliable backups established, the last chapter will focus on the even more unpleasant part: the recovery.

Imagine that everything is gone. Or even worse, you can't trust critical parts of your infrastructure—your data integrity is compromised.

Should you restore your backup application from scratch or not? Keep in mind any virtualization layer you rely on for core components, or any authentication tool required to log on again for restore operations.

Let's assume your application is up and running, with recent backup catalog information providing access to at least some backup copies. Any required de/encryption keys are available.

There will be high pressure to restore and recover as soon as possible while avoiding reinfection and selecting the most recent and correct backup.

How are restore and recovery different?

Restore refers to the process of returning a saved data set to an accessible data store. This should be independent of backup media or platforms involved in the backup process.

Recovery restores lost data by following processes up to the stage where it can be used again in its intended way—usually in its original state. This is the moment when data is placed back into applications, and they can start with the given point in time desired.

Recovery operations

When disaster strikes, you need to live with what you have and act fast. Regular testing should have validated the steps of restore and recovery. And thorough documentation should be on hand to guide your team.

The recovery operations will most likely have to be performed with limited people resources. Even with everyone on deck and external assistance, there will be many tasks to perform at once. With so many decisions to make, you can't afford unnecessary complexity.

Be sure to avoid:

- Multi-step approaches that waste time and add dependencies.
- Tools that require more components, more people skills, time, and money.
- Search-and-guess approaches that risk resetting your progress.

Central console

A central console offers all potential restore options—from storage hardware snapshots, recent primary backup, replicas, and copies of backups to any older backup set that is available—along with proper filtering and search options. The console should include detailed records on eventually vaulted and offsite-stored backups.

Enterprise-level backup

An enterprise backup and recovery software suite should provide a vast list of application backup and hypervisor integrations and comprehensive capabilities for understanding the applications, hypervisor resources, clusters, and database setups. Enterprise-level software enables simplified restores when every minute counts. It should offer down-to-the-second recovery by automated transaction log backup and truncation, reducing the impact of the disaster to an absolute minimum. Ideally, application owners and backup administrators can initiate such restores with proper guidance in the GUI so that no application-specific expert is required.

For critical data with a short RTO, it helps if the backup and recovery application supports hardware- or array-snapshot-based restores for situations where these systems are not compromised and would allow for reliable restore.

While it is useful to restore subsets of data only, the option to skip certain files has become an even more important feature. By skipping objects on restore, you can massively reduce the risk of reinfection instead of hoping for post-restore scanning processes, saving precious time and resources and improving RTO.

If you set RPO SLAs, restore operations will require backups to be performed on a schedule (the recovery point being actively targeted). The last thing you want is a bottleneck that causes delays and misses the target regularly, which then requires more resources.

If you search the word backup, you will come across many different ideas and solutions from vendors. However, several of these tools don't live up to the definition of true backup. And many will try and lock you in.

If you need to improve your backup and recovery approach, please contact us. We will help you evolve with as little overhead work as possible when elevating your data security to the next level.

Conclusion: Find true backup and recovery

If you search the word backup, you will come across many different ideas and solutions from vendors. However, several of these tools don't live up to the definition of true backup. And many will try and lock you in.

If a backup and recovery product or service catches your eye, consider the following:

- How would it help and in what scenario?
- Is there a way to achieve similar or better results without vendor or solution lock-in?
- Are any of the features too good to be true?
- Will it cause more problems than it's worth?

To conclude, we leave you with a list of questions based on these guidelines:

- Are the backups stored securely, or are they at risk on mounted volumes?
- Is the backup stored in a protected format so that it can't be read or manipulated by someone who is unauthorized?
- Does your backup and recovery solution consume lots of disk and compute power beyond today's requirements?
- Do restore plans exist? Are they tested and revised regularly?
- Is restore panic-mode ready?

If you need to improve your backup and recovery approach, please contact us. We will help you evolve with as little overhead work as possible when elevating your data security to the next level.

Appendix A: Backup and recovery scorecards

Backup Maturity Levels

Organization:

Data protection regulations apply: yes/no

Estimated volume of data to be protected: TB on premise /..TB at service provider /..TB in cloud Which backup and recovery tools are in use:

What further tools are in use for imaging, disaster recovery:

Backup and Recovery Maturity Level	Description
Worst	<ul style="list-style-type: none">• Backup and recovery are left to third parties.• There is no true backup, only mirroring, data copies, and replications.• An island approach to data protection introduces vulnerabilities.• The backup tool runs on the Windows platform, which is the primary target for hackers and ransomware gangs
Baseline	<ul style="list-style-type: none">• Business-critical data is regularly backed up.• Testing ensures the RTO and RPO of the restore.• Backup copies exist in two or more locations, with at least one backup copy being offline.
Better	<ul style="list-style-type: none">• One set of backups is stored in a disaster-proof environment yet is accessible by experts.• There are secure communications and a trusted relationship between critical employees and all backup infrastructure components.• Full backups and one-stop restores exist.
Best	<ul style="list-style-type: none">• Backups of all relevant data assets are done as often as possible with hardened core components.• Backups are stored in a secure, immutable, and transportable format.• Restores can be performed without restoring the malware too.

Table 3. Selecting the backup target architecture

Be wary of tools that:

1. create data copies that stay on mounted disk volumes and
2. store data copies or snapshots without proper format change.

Appendix B: True backup at a glance

True Backup

The word backup can mean different things to different people—and vendors. But a true backup must:

- Follow the 3-2-1 backup rule
- Obey the policy for different media in two categories:
 - **Logical:** Use of backup software with its own transportable, uncompromisable data format
 - **Physical:** Backup to dissimilar media—for instance, tape or to the cloud.
- Ensure that copies are distributed across at least two secure, separate locations. Storing tapes externally is also an option.

Snapshots don't count!

Neither do data copy and snapshot copy management tools!

Be wary of tools that:

1. Create data copies that stay on mounted disk volumes
2. Store data copies or snapshots without proper format change.

Any such approach cannot be considered a backup. Even if analysts endorse these tools, they don't provide reliable restore options nor the required data integrity guarantee. They also can't protect your data from being repurposed, stolen, or used against you.

Learn more at www.opentext.com