

Zerto

The Future of Backup: From Periodic to Continuous

September 2020


Gijsbert Janssen van Doorn



Introduction

Backup has been an essential component of IT infrastructure since its inception, and that is unlikely to ever change. But with the IT landscape changing rapidly and the number of cyberthreats increasing, can we rely on the backup technology we currently use? In this white paper, we discuss how backup requirements are changing and whether today’s backup technology can meet evolving business demands to drive modernization and digital transformation. We also explain why the future of backup is continuous journal-based protection, and why it’s time to move from recovery to availability and from restore to resume.

Changing Requirements



24/7 – As organizations increasingly focus on digital transformation, IT has become a critical strategic partner to business. The importance of keeping your systems up 24 hours per day, 7 days per week has never been higher, but availability means much more than just having systems “up.” Users expect the same experience every time, which requires IT systems to deliver high performance and stability regardless of the time of day.




Ransomware – The threat of ransomware is increasing quickly, and the impact of an attack is enormous. It’s not a question of “if,” but a question of “when” you will face this challenge. Choosing between paying the ransom or suffering loss of data is both costly and risky. Using traditional backup methods for recovery can result in up to 24 hours of data loss, and it could take days before all applications and systems are up and running again.

Today’s organizations can’t afford to lose data. To avoid productivity, data, and revenue loss, companies need more granularity in recovery, while maintaining the same level of performance. The International Data Corporation (IDC) determined that the average cost of downtime is \$250,000 per hour across all industries and organizational sizes. For more details, see: <https://www.zerto.com/page/idc-the-state-of-it-resilience-report-2019/>

Besides losing data and productivity, damage to an organization’s reputation is at stake. It’s easy for customers to share their frustrations on social media, where it can be seen by other customers or prospects. Research shows that an organization’s reputation is important to its stakeholders.

Importance of Company Reputation to the Board of Directors



Source: *The State of Corporate Reputation in 2020* by Weber Shandwick

Shortcomings of Traditional Backup

When looking at the backup technology currently protecting your data—one of your company’s most valuable assets—not much has changed over the last 35 years. The basic process has remained the same: during off-peak hours, copy the data that changed in your production environments and store that copy in another, secondary location.

Performance Impact

Most backups take place during off-peak hours because copying all that data takes time and impacts performance on production environments. Whether the solution uses agents in the operating system or snapshots on the virtual machines, the data is read directly from production systems and sent across the network. At best, the VMs are sluggish—at worst, they’re temporarily unusable. Every IT support engineer knows exactly what to look for when users complain about “slow” systems on Monday mornings.


Complexity and Cost

Scheduling and managing backups are usually resource-intensive tasks that require ensuring backup jobs don’t interfere with each other, and that database maintenance jobs don’t impact the run-time of your backups. Not to mention the extensive time spent checking and monitoring backup completions.

In an attempt to avoid the impact on performance and keep backup windows as short as possible, traditional backup vendors introduced distributed systems to handle the data being transferred (e.g. backup proxy, media agent). As the environment gets larger, more of these components must be added and configured—and often scaled up as well, typically requiring costly, high-spec physical servers. Managing and sizing the backup infrastructure becomes a complex process that requires dedicated specialists within the IT team. All of these components, complexities, and hidden costs translate into a high total cost of ownership (TCO).

Granularity

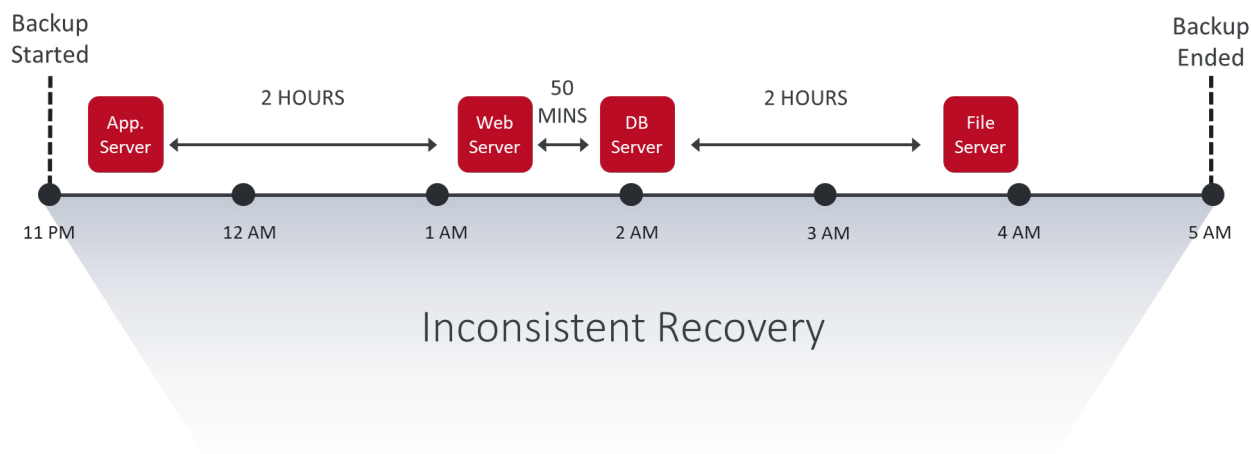
Due to the periodic nature of backups, IT teams are unable to meet the requirements for more granularity. And because of the performance impact on your production systems, backups can’t be made multiple times a day—thus the prevalence of the daily backup. However, this means when data needs to be recovered, the last available copy could be 24+ hours old and any changes since then are entirely lost.



Inconsistent recovery

In today’s IT environment, applications don’t reside on a single virtual machine (VM), but are spread across different VMs with different roles. Usually, the applications also have dependencies on other applications, which creates complex application chains. Successful recovery of entire application chains depends on how consistently you can recover the individual VMs.

For example, with traditional backup technology, jobs start at 11 pm and finish at 4 am—this means there could be up to 5 hours of difference between individual VMs. This inconsistency makes application recovery troublesome, complex, and time-consuming. It’s also the reason backup recovery time objectives (RTOs)—how long it takes to get back up and running—are so lengthy.



Backup Trends

Because data protection is such a vital component of a datacenter, the list of products that support any datacenter strategy can be endless. Let's focus on one of today's biggest trends: hyperconverged backup.

Hyperconverged backup consolidates compute resources, storage, and backup software into a purpose-built hardware appliance that enables scale-out architecture. Combining all of these resources and features into a single solution—and adding an easy-to-use interface to manage and schedule backups—solves many of the complexities involved in running more traditional build-your-own backup solutions.

But, does the hyperconverged backup model address the need for more granular recovery? It successfully reduces complexity in the backup architecture, but uses the same technology to protect the data: periodically copying data from the production systems to a secondary storage target.

The Future: Continuous Backup

To ensure granularity without impacting production performance, the future of backup is moving from periodic backup to continuous backup.

Continuous Replication


By using continuous data replication, you can deliver recovery point objectives (RPOs) of seconds by replicating every change that is generated in real-time. Backup should also rely on scale-out architecture for replication that allows protection of environments with thousands of VMs. All operations should be performed with zero performance impact on the production environment to deliver an uninterrupted user experience.

Granularity of Seconds

All of the replicated changes need to be stored in a journal, which not only allows you to go to the latest point in time, but also offers granularity of seconds so you can safely rewind back to any point in the past—even up to 30 days ago. Recover files, applications, VMs, and entire datacenters by simply pressing a virtual “rewind” button. Most recovery use cases that require granular recovery—such as file deletions, database corruption, or ransomware—only require short-term retention.

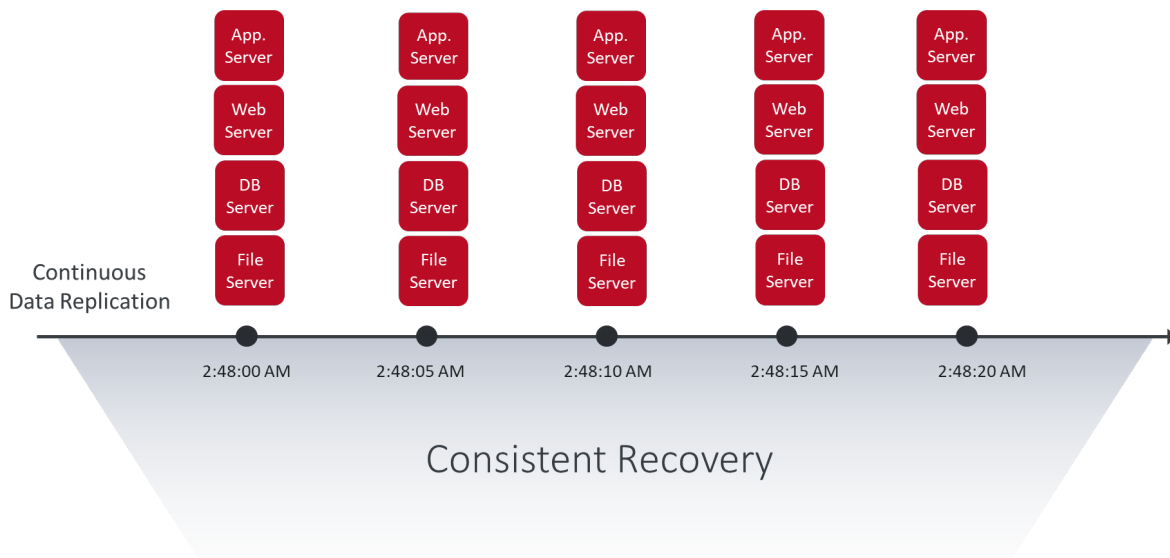
Continuous Data Protection (CDP)

Combining always-on replication and granular recovery enables continuous data protection that allows you to move away from the periodic point-in-time copies used in traditional backup technology. If an outage occurs at 17:26, CDP can restore data from 17:25, rather than from a backup that is probably at least 4 hours old—with all the data written since the 12:30 snapshot permanently lost.



Application Consistency

For consistent recovery of multi-VM applications, they must be protected as a cohesive, logical entity. All the VMs should share the exact same recovery point so that when the application is recovered, every VM that contains the application spins up from that same cross-application recovery point—no matter where the VMs are located within the infrastructure.



Long Term Retention

In addition to flexible options for short-term (up to 30 days) recovery scenarios, organizations often have compliance requirements to store data longer than 30 days. Long-term retention data requires different storage and recovery times, but must be an integral part of your data protection platform. As with short-term backups, copies should not come directly from production systems to avoid impacting performance and disrupting the user experience. A technology that benefits from the data already protected by CDP (and stored in a journal) allows you to offload point-in-time copies to secondary storage targets as often as you want.

Cost Efficiency

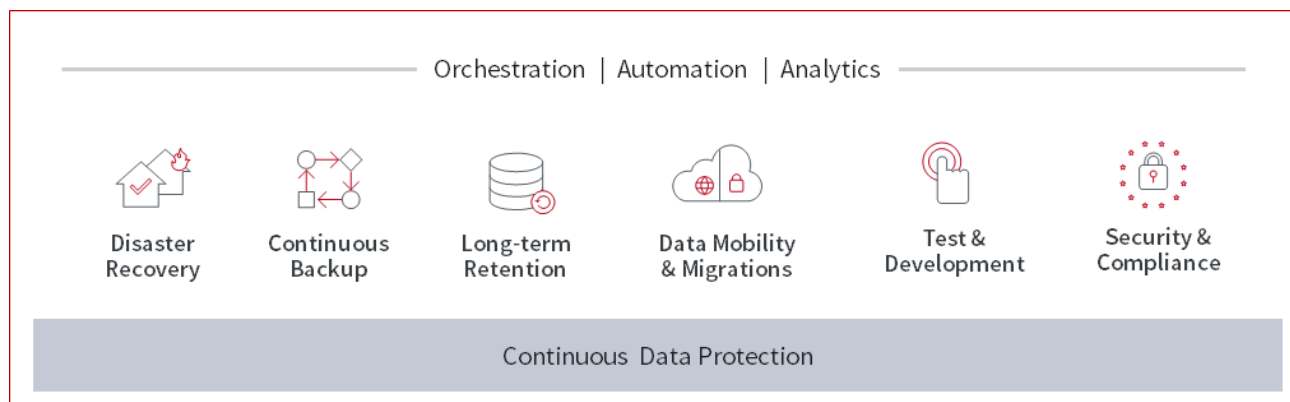
A software-only, scale-out backup solution that tightly integrates with the infrastructure of your choice ensures optimal protection without requiring dedicated hardware. To reduce your TCO, look for solutions with the following characteristics:

- **Technology-agnostic:** Flexible options for running the solution in a variety of virtualized environments, with the storage of your choice, helps future-proof the investment. In particular, look for myriad backup targets, including purpose-built backup appliances (PBBAs) running on-premises as well as native backup to hot or cool storage in the cloud.
- **Multi-purpose platforms:** A converged solution means that one investment serves a variety of needs across the data management and protection spectrum. An efficient platform can eliminate niche point solutions and save resources by consolidating down to one vendor for backup, disaster recovery, migrations, on-demand test/dev sandboxes, ransomware mitigation, and more.
- **Small footprint, low overhead components:** Significant savings result from avoiding dedicated hardware, high-end servers, and multiple pieces just to get a backup job done. Scaling out small virtual appliances is dramatically easier than adding physical servers as business grows and SLAs tighten.
- **Simple to use, easy to manage:** Today's backup vendors know it's critical to bring elegant, consumer-level usability standards to the enterprise world. Gone are the days of needing hefty professional service budgets and weeks of time-intensive training. A single, simple software experience, regardless of cloud or backup requirements, helps reduce hours spent managing the solution and babysitting backup windows.

The Zerto Platform

A Single Platform for Cloud Data Management and Protection

Zerto's software platform utilizes continuous data protection for Backup, Disaster Recovery and Data Mobility across on premises and cloud.



 [SCHEDULE A DEMO](#)

 [START A FREE TRIAL](#)

About Zerto

Zerto helps customers accelerate IT transformation by eliminating the risk and complexity of modernization and cloud adoption. By replacing multiple legacy solutions with a single IT Resilience Platform, Zerto is changing the way disaster recovery, data protection and cloud are managed. With enterprise scale, Zerto's software platform delivers continuous availability for an always-on customer experience while simplifying workload mobility to protect, recover and move applications freely across hybrid and multi-clouds. www.zerto.com

Copyright 2020 Zerto. All information may be subject to change.