REQUEST A DEMO

DISASTER RECOVERY | 101

Everything You Always Wanted to Know—But Were Afraid to Ask

REQUEST A DEMO

1 | WHY YOU NEED DISASTER RECOVERY

REQUEST A DEMO

DISASTER RECOVERY 101

Confused about RTOs and RPOs? Fuzzy about failover and failback? Wondering about the advantages of continuous data protection over snapshots? This eBook will help you learn about disaster recovery (DR) from the ground up so you can make informed decisions and implement an effective DR strategy. We'll show you how that DR strategy can be one key part in building a resilient IT infrastructure—the backbone of successful digital businesses and always-on customer experiences.

Why Do You Need DR?

No one is immune to disruptions, malicious attacks, and unrecoverable data

91%

91% have experienced a tech-related business disruption in the past 2 years

91%

91% have experienced some type of business disruption (variety of factors)

84%

84% have experienced a malicious attack in the past 12 months (that they know about)

56%

56% have experienced unrecoverable data within the last 3 years (up from 49% last year)

Source: IDC State of IT Resilience Report

REQUEST A DEMO

THE COST OF DOWNTIME

Modern businesses cannot afford to lose data. Customers and stakeholders, both internal and external, expect seamless 24/7 access to their data and applications. Whatever the cause—natural disaster, human error, or cyberattack—downtime and data loss are costly and can be extremely risky to the life of a business. Every enterprise, no matter the industry, needs a cutting-edge disaster recovery strategy to ensure uptime, minimize data loss, and maximize productivity no matter what kind of disruption or outage comes along.

Disruptions cost a business even when it's not tier 1 or critical applications that have an outage. And it's important to keep in mind that the cost of downtime is not only impacted by revenue-generating VMs, or those directly involved in creating or processing sales. Consider indirect impact as well:

- Brand damage, either for the IT division or the business as a whole

- Loss of productivity; for example, when email, file servers, or the CRM goes down

- Time spent during and after an incident on analysis, communication, or reporting

$250,000/hr

Average cost of downtime per hour across all industries and organizational sizes

$2,000,000/yr

The collective cost of 8 hours of downtime per year to an organization

Source: IDC State of IT Resilience Report

REQUEST A DEMO

IT'S ABOUT MORE THAN JUST DISASTERS

A rock-solid disaster recovery strategy is not limited to traditional disasters either. Organizations regularly contend with the need for operational recovery too. Operational recovery deals with the day-to-day realities of accidental deletions, overwritten files, or corrupted folders.

Straddling the line between operational recovery and disaster recovery are the potentially crippling effects of viruses, worms, and malware of all types. Notably, ransomware has become a particularly thorny problem for enterprises. According to Cybersecurity Ventures, the global cost of ransomware is predicted to reach $11.5 billion in 2019. This malicious software gains access to files, often via unsecured

RDP, then encrypts the data and generates a pair of private-public keys. The hacker holds the private key, and without it the data is impossible to decrypt until the ransom is paid (usually in Bitcoin). Sometimes, even after a company pays the ransom, the attackers never provide the decryption key, leaving victims without their money or their files.

Recent advancements in encryption technologies, coupled with the ease with which hackers can conceal their identities, has resulted in an increased adoption of ransomware strategies.

What Would Downtime Cost You?

Try our Downtime Calculator or contact Zerto for an in-depth consultation.

CALCULATE NOW

$ $ $

$ $ $

Zerto

CALCULATE NOW

REQUEST A DEMO

2 | MEASURING DOWNTIME

[icon] | **REQUEST A DEMO**

INTRODUCING RTO & RPO

Recovery point objective (RPO) is the last point in time IT systems and applications can be recovered to. It indicates the amount of data that will be lost, measured in elapsed time.

- The cost of **ONE HOUR** of lost data for most enterprises can easily hit six figures, which is one reason organizations of all sizes are reconsidering whether nightly backups (with an RPO of 24 hours) are still sufficient in today's demanding business environments.

- Due to the RPO's importance on data loss, it is recommended to agree on an acceptable, achievable RPO on a per-application basis.

- Always aim for the lowest RPO possible, then configure alerts to warn if you are in danger of the actual RPO exceeding your defined SLA. Ensure that your solution enables the prioritization of individual applications should the bandwidth for replication become constrained.

Recovery time objective (RTO) is the time that it takes to recover data and applications, meaning, how long will it be until business operations are back to normal after an outage or interruption.

- The cost of downtime associated with waiting for applications and data to be recovered (RTO) can result in significant loss in revenue and productivity.

DAILY BACKUPS

RPO = 24 Hours

UP TO $273,972.60[*]

SNAPSHOT-BASED REPLICATION

RPO = Hours

UP TO $45,662.10[*]

CONTINUOUS REPLICATION

RPO = Seconds

UP TO $7,610.35[*]

Example: Organization with a turnover of $100M

REQUEST A DEMO

PROVE YOUR DR WORKS

Recovery Report for Virtual Protection Group Production 3

Report was generated on 09/14/2019 12:02:29

Recovery Operation Details

Initiated by	System
Recovery operation	Failover Test
Point-in-time	09/14/2019 11:46:17
Recovery operation Start Time	09/14/2019 15:46:30
Recovery operation End Time	09/14/2019 16:01:07
RTO	00:07:43
Recovery operation result	Passed by user
User Notes	Stop Test for VPG Production 3

Virtual Protection Group Recovery Settings

Protected Site	Production
Recovery Site	Culpeper Prod
Default recovery host	Prod 1
Default recovery datastore	DSXtremCP6
Journal datastore	DRJOURNAL01
Default test recovery network	Zerto_TestNet
Default recovery folder	DR

Detailed Recovery Steps

#	Step Description	Result	Start	End Time	Executi
1.	Fail-over test VM 'c3putdmo2212d1'	Success	11:46:32	11:46:42	00:00:09
1.1.	Create Recovery VM 'c3putdmo2212db1'- testing recovery'	Success	11:46:33	11:46:38	00:00:05
1.2.	Reconfigure IP for VM 'c3putdmo2212db1'- testing recovery'	Success	11:46:41	11:46:41	00:00:00
16.	Fail-over test VM 'c3putdcts1'	Success	11:46:42	11:46:50	00:00:08
16.1	Create Recovery VM 'c3putdcts1''- testing recovery'	Success	11:46:43	11:46:49	00:00:06
16.2	Reconfigure IP for VM 'c3putdcts1' testing recovery'	Success	11:46:50	11:46:50	00:00:00
19.	Fail-over test VM 'c3pitdga2122ap1'	Success	11:46:42	11:46:50	00:00:07
19.1	Create Recovery VM 'c3pitdga2122ap1- testing recovery'	Success	11:46:43	11:46:47	00:00:03
19.2	Reconfigure IP for VM 'c3pitdga2122ap1'- testing recovery'	Success	11:46:49	11:46:49	00:00:00
20.	Fail-over test CM 'c3pitdoh2004ap1'	Success	11:46:42	11:46:50	00:00:07
25.	Fail-over test VMs' 'c3putdmo2212db1' volumes	Success	11:47:19	11:48:06	00:00:46
25.1	Create scratch volume for VM 'c3putdmo2212db1'	Success	11:47:19	11:47:44	00:00:24
25.2	Detach volume VMs' 'c3putdmo2212db1-0:1:' from	Success	11:47:47	11:47:56	00:00:08
27.1	Attach volume VMs' 'c3putdmo2212db1-0:1:' to	Success	11:47:54	11:48:02	00:00:08

In order to benchmark your RTO and tweak your DR plan to minimize downtime, testing is a must. By testing your plan with a DR technology that allows for no downtime in production or break in the replication, you can perform a test during working hours. This ensures you are able to fully recover, and you can run through the recovery operation multiple times to get your RTO as low as possible.

This is an actual successful failover test from a healthcare organization using Zerto virtual replication. The test was completed during a regular work day, with zero production impact.

This failover test covers the organization's tier one healthcare applications, consisting of 23 VMs with 8.3 TB of data, and took less than 15 min, with no downtime.

Note: Some data points in this report have been redacted to protect customer confidentiality.

REQUEST A DEMO

3 | COMPARING TECHNOLOGIES

NEXT

REQUEST A DEMO

5 REPLICATION TYPES

Array-Based Replication

Sometimes called storage-based replication, these solutions are deployed inside the storage array and replicate the entire LUN, regardless of its utilized capacity. They are designed for physical rather than virtual infrastructures and, as such, eliminate the benefits of virtualization.

There are two types of replication that can be deployed.

RPO = 0

Synchronous

1

Ensures all data is written to the source and target storage simultaneously, waiting for acknowledgment from both arrays before completing the operation. This relies on matching storage arrays and close geographic proximity between sites to achieve low fiber channel latencies and minimize performance impact. With the rise of all-flash arrays (AFA), latency of the connection between arrays becomes a bottleneck. Synchronous replication also runs the risk of quickly propagating malware and thus dramatically extending recovery times.

RPO > 1hr

Asynchronous

2

Uses storage snapshots to take a point-in-time copy of the data that has changed and sends it to the recovery site. The frequency is typically set on a schedule of hours, depending on the number and frequency of snapshots that the storage and application can withstand.

REQUEST A DEMO

5 REPLICATION TYPES (CONT'D)

RPO > 1hr
VM Snapshot-based
3

Uses VM-level snapshots to take a point-in-time copy of the data that has changed and sends it to the recovery site. The snapshots are created in the hypervisor and incur performance impact. It is not recommended to create, remove, or leave VM-level snapshots running on production VMs during working hours.

When VM-level snapshots are being used the only type of supported replication is asynchronous. The frequency of replication is typically scheduled to occur every few hours due to the performance impact of this type of technology.

RPO > 1hr
Guest-based
4

Otherwise known as agent- or OS-based replication, these are software components that must be installed on each physical and virtual server. Although more portable than array-based solutions, the requirement to install modules on every server limits scalability and is limited to only certain operating systems.

Guest-based replication typically only supports asynchronous replication. As it runs on the operating system of the production systems itself, it can impact the performance of these systems.

RPO = secs
Hypervisor-based
5

With these solutions, all writes are captured, cloned, and sent to the recovery site at the hypervisor layer, making it more efficient, accurate, and responsive than prior methods.

Hypervisor-based replication uses continuous data protection (CDP) and is constantly replicating only the changed data to the recovery site within seconds—it's always on. This technology thus combines the best of both synchronous and asynchronous replication. It does not need to be scheduled, does not use snapshots, and writes to the source storage without having to wait for acknowledgment from the target storage.

LEARN MORE

PERIODIC VERSUS CONTINUOUS BACKUP

Snapshots

When using any kind of snapshot (whether VM-level or storage-level technology), the replication occurs on a schedule. This results in the potential for more data loss in case of a disaster.

This periodic approach also impacts the granularity in recovery as checkpoints will only be generated every few hours.

4+ hours

24+ hours

CORRUPTION
3:43PM

Snapshot   BACKUP   Snapshot   Snapshot   Snapshot   Snapshot   Snapshot
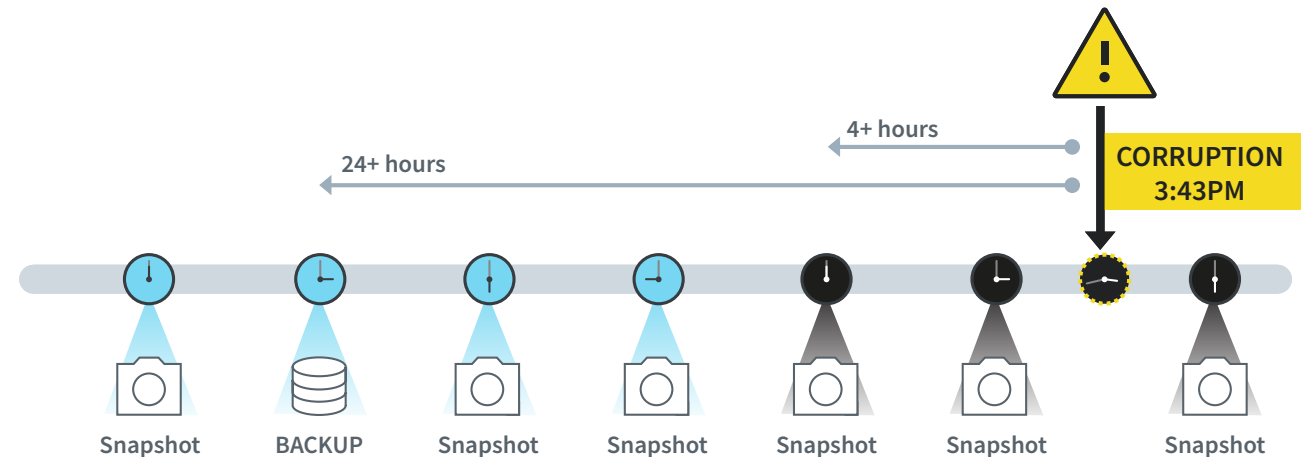
If we take the above example of a data corruption at 15:43, then a VM-level 24-hour snapshot-based replication solution means you are going to potentially have nearly 16 hours of data loss, as you would have to restore a replicated snapshot from last night. The same example with storage-level snapshots would result in data loss of nearly 4 hours.
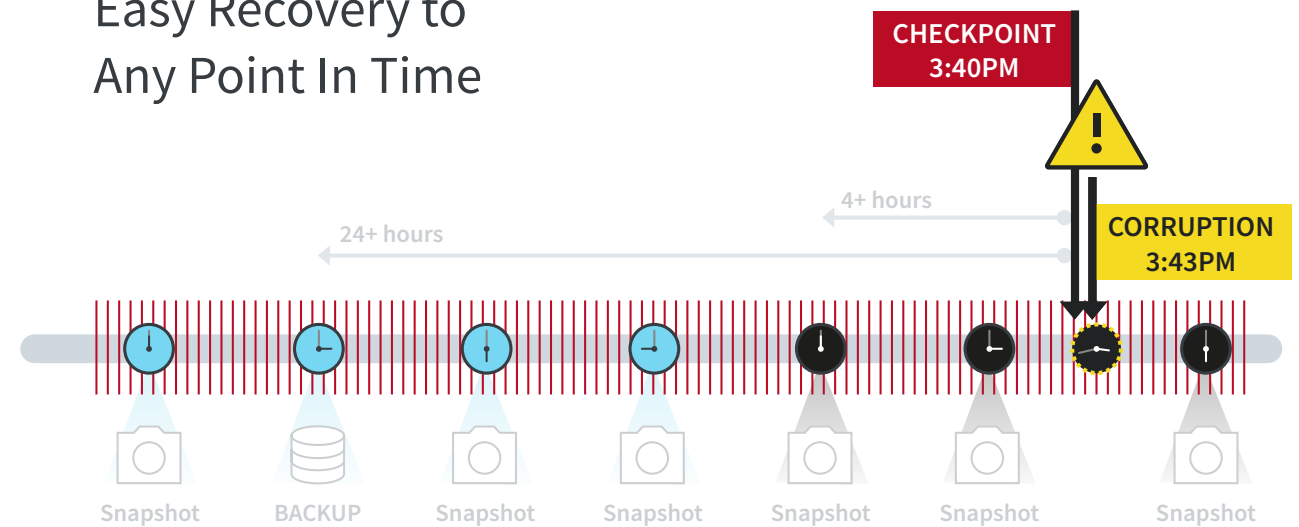
REQUEST A DEMO

NEXT

REQUEST A DEMO

Continuous Data Protection

Continuous data protection (CDP) utilizes changed block tracking at the hypervisor layer to constantly replicate data as it is written to storage. Because CDP replicates only changed information, rather than an image of the entire host or array, there is no impact to the performance of the replicated VM.

Hypervisor-based CDP also utilizes journal technology to keep a log of all the changes occurring in a specified journal timeframe, allowing point-in-time recovery in increments of just seconds for the length of the journal.

Because CDP is always on and always replicating the most recently changed data, it offers considerably lower RPOs than snapshot-based solutions. This results in significantly less data loss to the business and consequently, a far lower cost of impact.

Easy Recovery to Any Point In Time

CHECKPOINT
3:40PM

4+ hours

CORRUPTION
3:43PM

24+ hours

Snapshot   BACKUP   Snapshot   Snapshot   Snapshot   Snapshot   Snapshot

REQUEST A DEMO

THE ROLE OF CDP IN BACKUP

Backup typically uses snapshot-based technology to protect virtual environments. Because the data that is backed up is read directly from production systems, backups are almost always run at night to avoid "VM stun," or any kind of application lag or slowdown.

When an environment is protected by CDP, you can recover data from seconds ago on production-grade systems within only a few minutes. Besides that, the recovery site can also be used as a source for longer-term backups. This completely eliminates the concept of backup windows.

DOWNLOAD THE WHITE PAPER

SNAPSHOTS	VS	CDP
Storing multiple snapshots on replica VMs incurs a significant VM performance penalty if you attempt to power on the replica VM.		With journal-based protection, the journal is only used until you commit to the point-in-time selected, without the performance impact of many snapshots.
Using snapshots on replicated VMs allows no way of controlling the total space used for snapshots or the ability to store the data change on a separate datastore. This makes it unscalable in terms of being able to set SLAs and define maximum limits on the data space used by the snapshots.		With journal-based protection, you can place the journal on any datastore and place maximum size limits and warnings so as not to fill the datastore, which could otherwise break replication and recovery.
With snapshot-based replication, there is often significant overhead on the storage arrays for replication reserves, which can be 20-30% on both source and target storage in many cases.		With journaling technology, no extra space is used in the source storage as no snapshots are created. Only 7-10% of the target storage is typically used for the changed data, freeing up significant amounts of disk space.

REQUEST A DEMO

4 DISASTER RECOVERY AT SCALE

NEXT

REQUEST A DEMO

BUILT-IN AUTOMATION & ORCHESTRATION

The top hypervisor-based replication solutions include replication, recovery automation, and orchestration all in one. The VMs that form each application are recovered together in consistency groups from the same point in time. Boot-ordering is then applied to ensure that the VMs come online in the correct order, and re-IP or MAC addressing can be utilized if needed to ensure there is no break in communication. This ensures an RTO of just minutes with no manual operations required since the application is automatically recovered in a working and consistent state.

No-impact failover testing also enables this automated process to be tested during working hours in minutes, with no shutdown in production or break in

replication. Reports can be generated to show the testing outcomes and prove the recovery capability. This enables organizations to increase the frequency of DR testing, mitigate risk, and satisfy compliance initiatives.

ANALYTICS

Orchestration & Automation				
DISASTER RECOVERY	OPERATIONAL RECOVERY	LONG-TERM RETENTION	HYBRID, MULTI-CLOUD	OPERATIONAL SERVICES
Continuous Data Protection				

REQUEST A DEMO

APIs AND SCRIPTING

Although you'll want to avoid solutions that are overly dependent on scripting—which will inhibit efficient scaling, even with professional services—it's inevitable that enterprise-scale deployments will need open, REST APIs. The platforms that have an API-first approach (potentially using something like Swagger or Postman) will allow a business virtually unlimited ways of integrating systems together for effective, automated DR.

VISIT ZERTO ON GITHUB

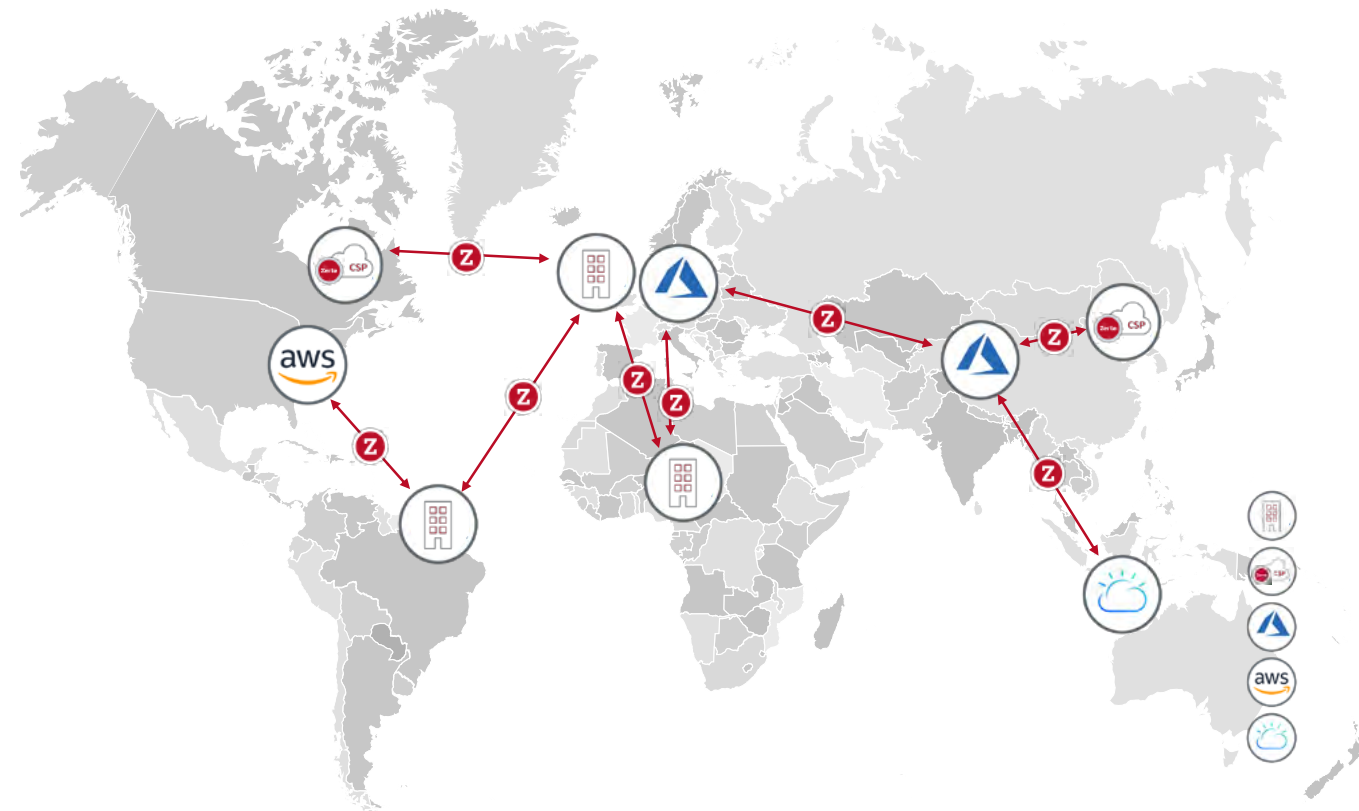REQUEST A DEMO

MULTI-SITE MANAGEMENT

As a business grows, typically so does its need to run multiple datacenters, and they're often geographically diverse and potentially using disparate hardware and vendors. A modern DR strategy must account for all the various sites and clouds being utilized, no matter where they're located, what hypervisor they're on, or how big the environments.

Part of effective management will include role-based access control (RBAC), but it's also important to have a single pane of glass that provides a window into everything the organization is running. A multi-site management platform can show IT teams at a glance what the state of their protection coverage is, as well as highlight any critical issues that need to be addressed. Dynamic analytics, especially if it includes forecasting or modeling, can be

especially useful in seeing exactly how your infrastructure is performing (e.g., in terms of storage or bandwidth consumed) as well as plan for future growth.

REQUEST A DEMO

5 | TOTAL COST OF OWNERSHIP

NEXT

REQUEST A DEMO

SNAPSHOT SPACE UTILIZATION & COST

Challenge: Snapshot solutions typically consume more than 20% of both the source and target storage, plus an additional 5% of replication reserve space. The cost and overheads of utilizing array-based replication must therefore include the cost per TB multiplied by the storage usage of the snapshots and replication reservations.

Solution: The ability to recover to previous points in time is enabled by keeping a journal on the recovery site storage, which dynamically grows and shrinks to the size of the changes for the time it is configured to keep.

Forrester's Total Economic Impact of Zerto

FORRESTER®

📄 | **DOWNLOAD REPORT**

RECOVERY ORCHESTRATION & AUTOMATION

Challenge: Utilizing storage replication simply creates a copy of the data in the recovery site. To recover the data during testing or a DR event, it needs to either be done manually using scripts, or by utilizing an orchestration and automation solution. Due to the time it takes to recover manually and the difficulty in conducting tests, an additional orchestration and automation solution is recommended. The cost of purchasing the licensing of the additional solution and managing multiple solutions should therefore be factored in.

Solution: Hypervisor-based solutions include recovery automation and orchestration features such as application-consistent VM boot-ordering, re-IP/MAC addressing and custom pre-/post-scripting, in addition to the continuous replication technology. This significantly reduces the RTO as well as the cost and complexity of managing multiple solutions.

STORAGE LOCK-IN

Challenge: Array-based replication solutions are vendor-specific and require matching storage arrays in both the source and target sites. This can significantly increase the TCO of the next storage refresh by having to buy new and matching storage arrays, just to configure replication. There is no ability to mix storage vendors and technologies to get the best price-to-performance ratio in a recovery site, or to introduce new storage vendors to improve performance.

Solution: Hypervisor-based replication operates at the virtual, not physical layer, meaning it is inherently storage-agnostic. This allows you to buy or use any storage in any site, reducing the TCO of your next storage refresh and enabling the seamless adoption of new technology. Even if the same storage is used in both source and target sites, replicating from the hypervisor removes complexity to save on the cost of management overheads.

NEXT
〉

REQUEST A DEMO

6 THE FUTURE OF DISASTER RECOVERY

REQUEST A DEMO

DISASTER RECOVERY AS A SERVICE (DRaaS)

Since a dedicated DR site can be expensive to maintain and scale, many organizations are looking to outsource DR to specialized cloud providers who can manage it for them. Replacing the direct and indirect costs of secondary sites—including hardware, software, facilities, people—with a predictable monthly expense and high burst capacity can be a very attractive option. Look for service providers with a deep background in data protection and, if required, proven expertise in your specific industry or with your specific infrastructure and workload requirements. Consider both niche providers in your region as well major players, such as those on Gartner's Magic Quadrant for DRaaS.

Why Consider DRaaS?

DRaaS providers do this every day, so they are knowledgeable about getting environments online quickly and can help you avoid common mistakes. They also serve as additional resources that are focused on your datacenter recovery when you need it most.

$$\$\$$$

Control Costs

Gain greater predictability of storage costs and choose the DR strategy that is right for you.

Diversify Data Protection

Gain confidence with target site diversification. Take advantage of the extended global network of DR sites afforded by managed service providers.

Take DR to the Cloud

Leverage a DRaaS provider to be your guide to the cloud.

REQUEST A DEMO

CONVERGED PLATFORMS OVER POINT SOLUTIONS

In recent years, both IT vendors and IT customers have increasingly found consensus that continuous data protection provides the highest level of protection and DR readiness. The emergence of robust CDP platforms has led to a shift away from point solutions that, for example, only provide DR, only do migrations, or only provide backup. The future of DR lies with converged solutions that offer DR, backup, cloud mobility, and on-demand operational services that cover the full breadth of what today's modern IT is asked to support.
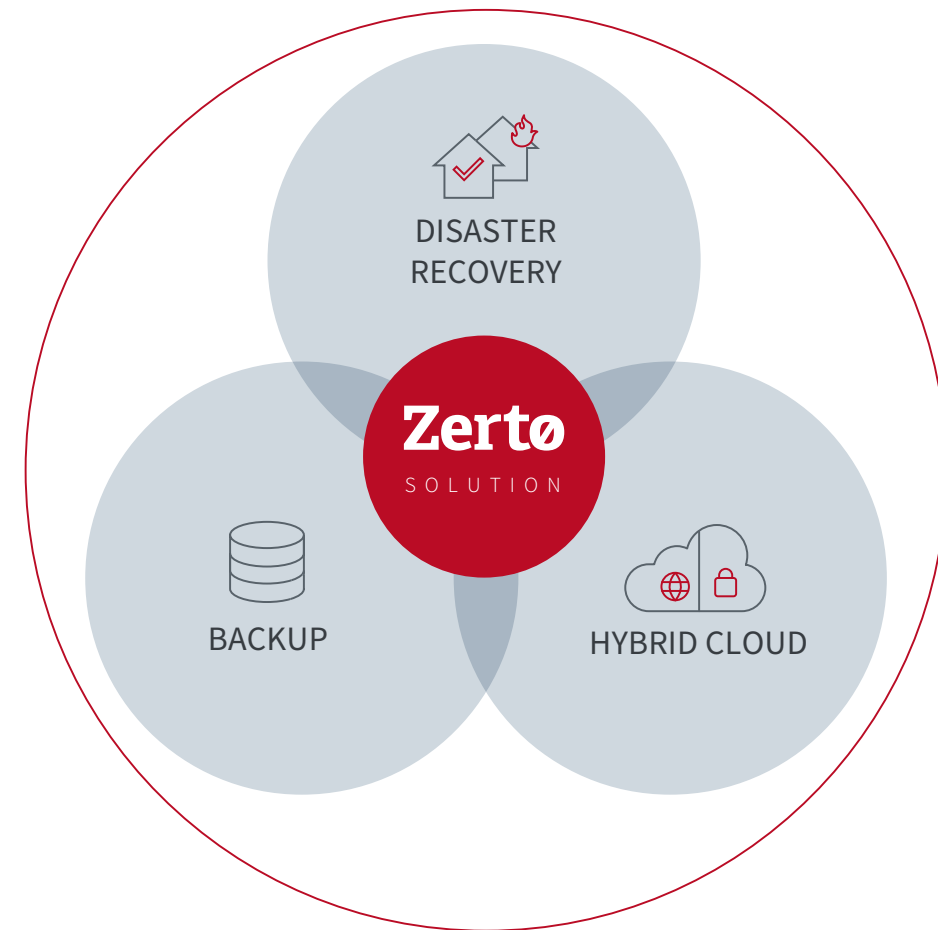
DISASTER RECOVERY

Zerto
SOLUTION

BACKUP

HYBRID CLOUD

REQUEST A DEMO

HYBRID AND MULTI-CLOUD DR

The proliferation of hypervisor and cloud options has not been without its challenges, but IT teams are increasingly poised to fully leverage the benefits for DR. New architectures and new models are being unlocked as businesses look to seamlessly move to, from, and between clouds of all types. The ability to place the right workload on the right cloud means greater cost savings and greater efficiencies. The DR solution of the future will need to be more than simply cross-hypervisor, but also allow IT organizations to quickly configure and refigure their deployments as the business evolves without being locked into a single technology or provider. When comparing vendors, ensure you're looking for those with built-in analytics and multi-site management tools that provide visibility across all your clouds and environments.
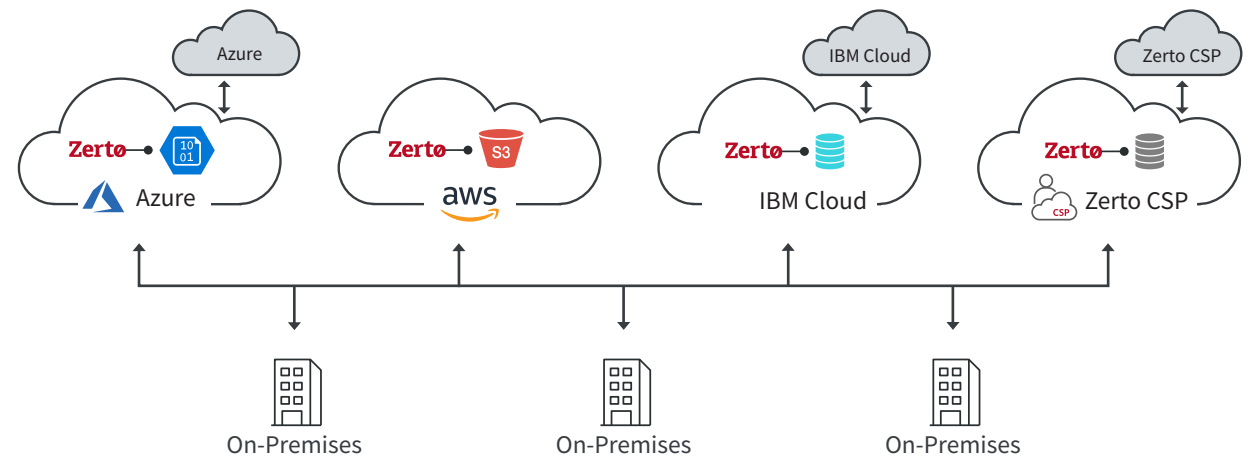
NEXT

REQUEST A DEMO

IT RESILIENCE

Many businesses have recently embraced strategic plans that move beyond DR as a reactive, standalone function within IT. Instead, DR needs can be met with enterprise-class platforms that unlock proactive, integrated IT functions. By leveraging the future-forward features described above—such as converged DR and backup alongside multi-cloud mobility—organizations can be well on their way to becoming IT resilient. IT resilience is the ability for a business to protect what exists today while simultaneously innovating for the needs of tomorrow. It encompasses both the unplanned nature of DR and its unwanted disruptions, as well as planned, beneficial disruptions such as mergers and acquisitions, datacenter consolidations, and hardware refreshes. In that way, the future of DR is not the narrow role it played in yesterday's transactional IT, but as an expansive, strategic part of every resilient IT organization.

With IT Resilience, organizations can withstand any disaster, confidently embrace change and focus on business.

24/7 BUSINESS

IT RESILIENCE

72%

Of companies have experienced an IT outage in the last year [1]

3 STEPS TO IT RESILIENCE

Minimize Service Disruption | Ensure Application Mobility | Create Infrastructure Flexibility

70%

Of enterprises that will have a hyprid cloud strategy by 2019 [2]

1 in 3
firms has had at least one declared disaster or major disruption during the past 5 years

Sources:
[1] *Gartner BCM survey* (https://www.gartner.com/doc/3200321/survey-analysis--bcm-survey)

[2] *Gartner "The future of the Datacenter in the Cloud Era"* (https://www.gartner.com/document/3079122?ref=unauthreader&srcId=1-3478922254)

NEXT

REQUEST A DEMO

CONCLUSION

SO, WHERE DO YOU GO FROM HERE?

Zerto, a Hewlett Packard Enterprise company, empowers customers to run an always-on business by simplifying the protection, recovery, and mobility of on-premises and cloud applications. Zerto's cloud data management and protection solution eliminates the risks and complexity of modernization and cloud adoption across private, public, and hybrid deployments. The simple, software-only solution uses continuous data protection at scale to converge disaster recovery, backup, and data mobility. Zerto is trusted by over 9,500 customers globally and is powering offerings for Microsoft Azure, IBM Cloud, AWS, Google Cloud, Oracle Cloud, and more than 350 managed service providers.

 Learn more at www.zerto.com