Zerto

a Hewlett Packard Enterprise company

Disaster Recovery Guide

DISASTER RECOVERY GUIDE

INTRODUCTION

In today's digital-centric landscape, organizations depend on infrastructure, applications, and data that are up and running 24/7. The combined costs of downtime and data loss can put a company out of business, make critical public services unavailable, and even threaten national security. Disasters that cause IT downtime and data loss range in size and scope from local cyberattacks to regional natural disasters. Thorough security and business resilience strategies are crucial for organizations looking to compete, thrive, and face these threats in the coming decades.

The data center has evolved far beyond a climate-controlled room of servers, storage, and networking devices. These components are now defined by software, expanding data past the four walls of the data center. That data is expanding not just to the cloud but also to the edges of networks, where it's collected and used directly with users, customers, students, patients, and the digital devices that are running our businesses—and our lives.

Digital transformation and data growth have changed at a rate that has outpaced the ability of many disaster recovery (DR) solutions to adapt to modern needs for recovery. Most DR solutions are still modeled around the physical data center and lack the ability to scale with the amount of data modern organizations produce and use.

In this guide, we provide insights into the challenges, needs, strategies, and available solutions for data protection, especially in modern, digital-centric environments. We explain which benefits and efficiencies Zerto, a Hewlett Packard Enterprise company, delivers, and how it compares to other business continuity/disaster recovery (BC/DR) technologies. Within this guide, we provide organizations with the right information to choose the best possible data protection solution for their needs. If you have any questions while reading this guide, please contact us at info@zerto.com.

SECTION 1

Disaster Recovery Overview

A disaster can be lurking behind nearly any planned or unplanned event, threatening an organization's ability to do business. To keep operations running, organizations need to have a sound DR strategy that focuses on two key goals:

React effectively to unplanned events, such as natural disasters, infrastructure failures, power outages, user errors, ransomware, corruption, and more.

Protect data during planned disruptive events such as migrations, cloud adoption, data center consolidations, and more.

A good DR strategy can help with reacting to not only unplanned events like disasters, but also during planned disruptions like business-driving initiatives. The challenge for IT is to stay protected enough to efficiently handle the reactive nature of unplanned disruptions while also innovating through proactive disruptions that deliver business value.

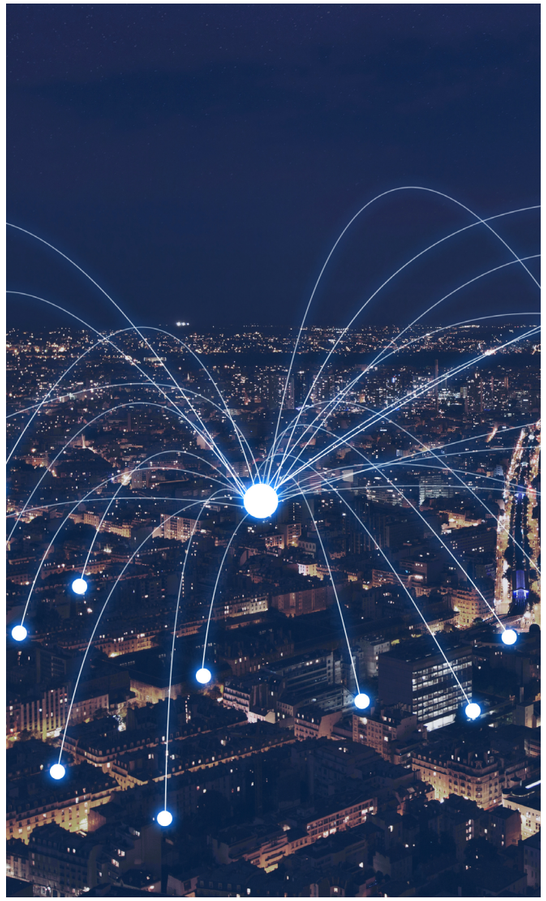

Disasters in All Shapes and Sizes

There are two main categories of disasters: natural and man-made.

Natural disasters include earthquakes, hurricanes, tornados, wildfires, and floods—events that cause both major and minor disruptions. Natural disasters range in scale from those affecting small areas, like tornados, to those destabilizing entire regions, like hurricanes. Natural disasters can be so devastating that they physically destroy entire sites—or even multiple sites—within a region. Even without directly damaging a data center, natural disasters can cause outages in power and communications.

Man-made disasters have an even wider scope. They can include fires and floods, or they can be more subtle, like cyberattacks and accidental/malicious deletion of data. Because man-made disasters can be purposeful or accidental, their type and scope are difficult to predict. Things like simple misconfigurations or unplanned environment changes could affect a single building, while events like power outages can affect entire regions.

Recently, ransomware has taken the stage as one of the most prominent man-made disasters. Ransomware attacks are proven to cause widespread disruption by impacting supply-chains and affecting utilities. These attacks have become a top concern for security specialists, DR specialists, and, to some extent, those in charge of national security.

Whatever the size and shape of the disaster, a DR strategy must be prepared to respond, recover, and continue operations quickly. In the same way, a DR solution should provide options to recover quickly from any disasters—natural or man-made, local or regional.

The Relationship Between Disaster Recovery and Business Continuity

Business continuity has a larger scope than DR alone. Business continuity ensures that business operations are never interrupted by planned or unplanned events, which encompasses DR response strategies for unplanned disruptions. Examples of business continuity for planned events include workload migrations or planned maintenance without downtime. Business continuity can also help avert disaster—consider this example:

A remote data center is in the path of a hurricane that may hit in a few days. The DR team fails over that data center's applications and data to a remote standby site until the danger has passed. Whether a disaster occurs or not, business continuity is guaranteed through preemptive action.

While business continuity goes beyond the scope of DR, there should be no need to have completely different toolsets for each. A good DR solution can also fulfill the needs of business continuity with capabilities like workload migration, nondisruptive patch testing, and nondisruptive malware scanning. Consider how the capabilities of a DR solution extend beyond disaster recovery alone to get the most return on investment from your solution.

Backup Is Not Disaster Recovery

Backup of data or an entire system is a concept that has been around since nearly the beginning of IT. Generally, backup means replicating data to another device or location for long-term retention or compliance. However, the days of a backup alone providing sufficient DR capability are long gone. Traditional backups fail to deliver adequate recovery time objectives (RTOs) and recovery point objectives (RPOs) because they typically impact system production and are taken only periodically, every few hours or once a day. But the time it takes to actually recover data from backup can be measured in days or weeks.

Backup services have attempted to improve recovery times and recovery points over the last two decades, but backup still just doesn't meet the recovery requirements for modern organizations that run 24/7 services. Backups are a necessary component of an IT environment, but they should be supplemented with a DR solution that fills the gaps that backup can't.

The Costs of Downtime and Data Loss

During a disaster and in its aftermath, downtime and data loss can become extraordinarily costly to organizations. Many businesses closed their doors forever because disasters irreparably damaged their reputation, productivity, and revenue. Downtime and data loss have their own unique costs, but sometimes data loss is the cause of downtime, which is often the case in a ransomware attack.

Downtime

From online shopping and banking to streaming entertainment, services are now expected to be available 24/7. Even businesses with limited operating hours like restaurants or hair salons often rely on digital services for appointments and reservations, services that customers and owners expect to be available at all times to keep business moving.

Downtime during peak production hours can incur costly disruptions to productivity, services, and transactions. However, any amount of downtime, whether planned or unplanned, can result in loss of productivity, business, loss, and reputation.

Data Loss

Loss of data that can be measured in minutes, hours, or days can have an extreme cost to an organization, especially when that data represents productivity, intellectual property, or key business transactions. Data loss can also cause additional downtime when key systems are unable to function correctly without current data.

Research from various institutes shows that the volume and cost of data loss are increasing year after year. A business continuity strategy—which ensures uptime, diminishes data loss, and maximizes productivity amid any compromising situation—is a necessary digital assurance policy for any company. The question is not if a disaster will strike, but *when*.

CONSEQUENCES OF DATA DISRUPTIONS

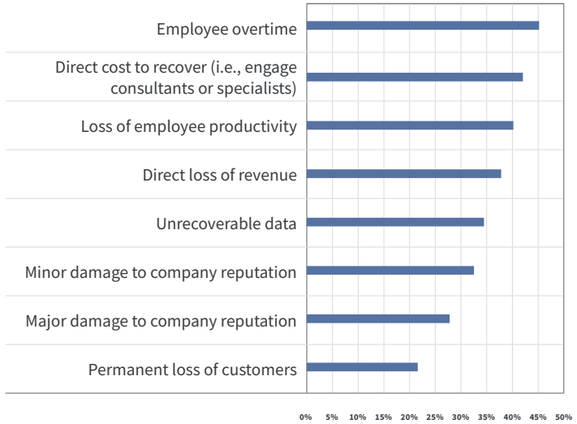

Figure 1. The number one reason survey respondents cited for data loss was the gap between backups. Infrequent, periodic backup solutions are not effective in eliminating data loss.

Source: IDC's Worldwide State of Data Protection and Disaster Recovery Survey, sponsored by Zerto, January 2022

RTO and RPO for Unplanned Downtime

Disaster recovery is expressed in two types of objectives: recovery time objective (RTO) and recovery point objective (RPO).

- The RTO is the amount of time a business can be impacted without significant losses or risks. RTO differs greatly between solutions, with many backup solutions taking days or weeks to recover systems.
- The RPO is the most recent point in time from which data can be recovered. Traditional backup or snapshot technologies have RPOs as low as 15 minutes and as high as 24 hours.

In a modern, digital-centric world, both RTOs and RPOs need to be as low as possible. They can no longer be expressed in hours; they should be minutes or seconds. Though many organizations focus on RTOs to get the business up and running as soon as possible, the inability to reproduce the data loss—RPO—will haunt an organization for a long time after any disaster.

RTO AND RPO USE CASES

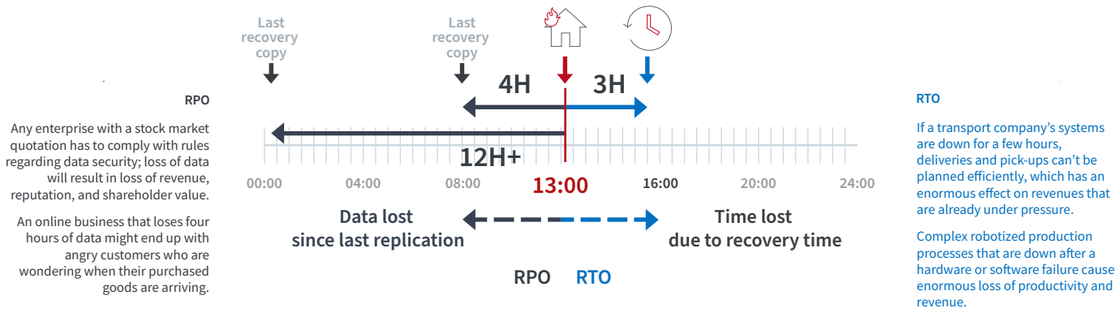

Figure 2. RPOs and RTOs affect different organizations and industries in different ways and can impact both productivity and revenue. Measuring and defining both is very important to assessing and meeting DR needs.

SECTION 2

Key Considerations of a Disaster Recovery Strategy

When creating a DR strategy, you have many factors to consider, especially when choosing the right solution to execute your strategy. Understanding these considerations will help your organization be prepared when disaster strikes.

Achieving the Best RTO and RPO

The two most important factors in DR are quickly restoring operations and preventing data loss. Both RTO (the time it takes to get back up and running) and RPO (the amount of data lost) should be as low as possible. You may have tightly defined SLAs that you need your RTOs and RPOs to achieve. But regardless of your SLAs, the better your RTOs and RPOs, the less time and money you will lose when a disaster strikes.

The best RTOs are measured in minutes. You can achieve these in two ways:

1. In a partial or limited primary system outage, you can instantly recover data and systems from a local replica. Whether the local replica is a backup, snapshot, or CDP journal, this DR strategy instantly brings data back to production and resumes normal system operation with no extra delay. In this method, the original site and systems need to be available for recovery.
2. In scenarios where the primary site and systems are unavailable, failing over to a warm site for recovery is preferable. This can be nearly as fast as local recovery because the warm site can recover from the same or similar replicas that existed at the primary site but that were replicated to the warm site. However, redirecting users across the network to the warm site to access the applications and data may slow down RTOs, potentially making this option slower than the first.

What is important to consider here is that if data has been replicated to a second site, but there are not orchestrated and automated failover mechanisms to quickly bring that data and applications online, then RTO will be seriously impacted and will be measured in hours or days to fully restore systems.

The best RPOs are measured in seconds. They rely on replication technologies that are sometimes described as real time, synchronous, near synchronous, or have intervals measured in seconds. Another key factor in replication for RPO is application-centric protection. Enterprise applications contain multiple VMs and dependencies and must be recovered as single, consistent entities with acceptable RTOs and RPOs. Not all replication technologies are able to achieve recovery consistency alone—instead, they must rely on periodic checkpoints, which reduces application RPOs to minutes or hours rather than seconds.

Achieving the best possible RTOs and RPOs simply minimizes the cost of a disaster by minimizing the disruption. The sooner systems are back online with the most recent data, the sooner the organization can be back up and running.

BIGGEST CHALLENGES REGARDING BACKUP AND RECOVERY

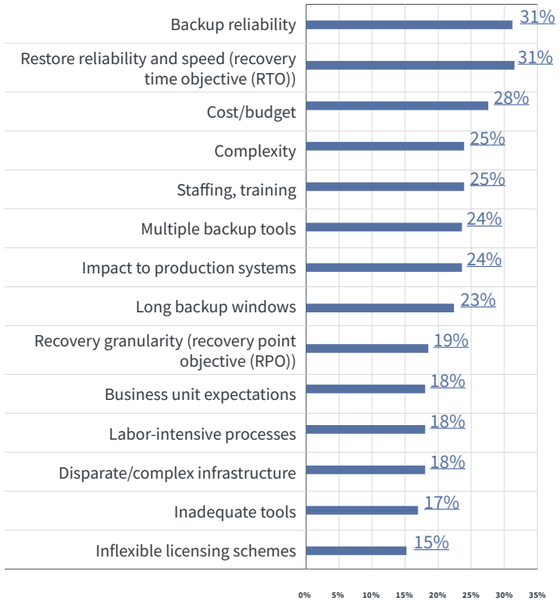

Figure 3. Outside of cost, the top challenges cited by survey respondents were regarding reliability and complexity in their backup and recovery solutions. Both backup reliability and recovery reliability were cited which impact both RPO and RTO.

Source: IDC The State of Ransomware and Disaster Preparedness: 2022

Ransomware Resilience

Cyberattacks used to only steal data, and it was the sole responsibility of cybersecurity experts to stop them. But ransomware has changed the scope and impact of cyberattacks, and now DR experts are at the frontlines of cybersecurity. Historical trends show that when businesses and other organizations like schools, hospitals, and local governments experience disruptions and don't have effective DR measures in place, they find it cheaper shown that without effective recovery solutions, organizations are finding it less costly to pay a ransom—which perpetuates the cycle of ransomware. Ransomware has become so prevalent that it's not a question of if but when your organization is attacked.

Ransomware can attack with a number or encryption methods, targeting anything from individual files to entire systems. Rather than a backup solution, you need a DR solution that can:

- Rewind your systems to the last point in time before the infection struck, to within a matter of seconds.
- Automate recovery of all critical systems within minutes, with only a few clicks of a button.
- Create multiple copies of data across multiple sites in case an entire site is compromised.
- Restore entire applications, databases, and individual files with consistency.
- Perform nondisruptive failover tests at any time to ensure your business can be brought back online straight away.

- Assist in ransomware detection by offering on-demand sandboxes for security scanning.
- Assist in security patching using on-demand sandboxes for patch testing.
- Create offsite data copies for immutable data copies and longer-term data retention.

Workload Prioritization

The more digitized the world becomes, the more critical some of your IT systems are in offering the 24/7 services that users and customers expect. When developing a DR strategy, you should identify which systems, applications, and data need the highest levels of data protection and availability and which might require less. Your SLAs may determine this, or you may have to examine how different systems impact revenue streams or productivity.

For the core applications, a working DR strategy with a remote DR site, low RTO and RPO (low data loss and short recovery time), and a tested recovery plan is essential. For other applications and types of data, higher RPOs and RTOs might be acceptable.

Prioritization is a key element for DR planning. Review with business owners what downtime can be tolerated for each application. It will become clear which applications need to be available fast with minimal data loss.

Total Cost of Ownership

When designing your DR strategy, there are many solutions to choose from. The cost of these solutions will vary, but so will their effectiveness. It is important to look beyond the purchasing or licensing costs of the solution. The total cost will include additional expenses for implementation, management, training, and, most importantly, recovery when disaster strikes, which includes the cost of downtime and data loss. Not all solutions will provide the same RTO or RPO, and the cost of recovery time and data loss must be factored into the overall cost. Sometimes a solution that seems less costly up front can end up costing far more when a disaster hits and you must deal with a difficult recovery.

You should also consider the number of tools that will be included in a DR plan. For example, a DR plan that relies on many different and complex technologies will lead to a complex and difficult recovery process. When the pressure is on, using several different tools can lead to errors in a time where errors are the costliest. Consider a single solution that delivers all the components of your DR plan should be considered.

Even within a single solution, however, additional tools sometimes incur an additional cost. A solution's cost can increase dramatically if the software vendor charges a premium licensing fee for additional enterprise features, such as orchestration and automation.

Replication Technologies

Over the years, many replication technologies have been developed. Many were developed before virtualization existed, and they have yet to become virtualization- or cloud-aware. Understanding these technologies—and which are best suited to the modern digital world of virtualization and cloud platforms—is key to an effective DR strategy.

Array-Based Replication

Storage vendors provide array-based replication products, which are deployed as modules inside the storage array. They are single-vendor solutions, compatible only with the specific storage solution already in use. The relationship between the VM and storage is fixed, and the entire LUN is replicated regardless of whether it's fully or partially utilized. Array-based replication is limited by some of the following characteristics:

- **Hardware-defined.** Array-based replication is designed to replicate physical entities. It doesn't "see" VMs and is oblivious to configuration changes.
- **Not independent.** Though optimized to work with the existing storage array, array-based replication locks the organization in to a single vendor.
- **More management points.** In addition to the physical storage array's management console, IT needs to manage virtual assets from a virtualization management console as well.
- **Growth and change.** The relationship between the VM and storage is fixed, eliminating the flexibility of virtualization, and removing the ability to respond to evolving business needs.
- **Granularity.** Because it needs to replicate the entire LUN, array-based replication lacks the granularity needed in a virtual environment.
- **Costs.** Replicating the entire LUN increases power, cooling, networking, and storage costs, even if only 40% of the LUN is utilized.
- **Single point for recovery.** Many array-based solutions can't store a history of LUN performance. So, if the last data point was corrupted, businesses must use it for recovery, rendering this DR solution useless.
- **Time requirement.** Without automation, recovery is very time-consuming and complicated; VMs and applications must be built from scratch.

Appliance-Based Replication

Appliance-based replication solutions are hardware-based and specific to a single platform. The main difference between appliance-based and array-based replication is that appliance-based replication runs on an external, physical appliance instead of inside the storage array itself. This makes it more flexible and less resource intensive than array-based replication. But the disadvantages are more or less the same as for array-based replication. Appliance-based replication is limited by some of the following characteristics:

- **Hardware-defined.** Designed to replicate physical entities rather than virtual entities.
- **Not independent.** Though more flexible than array-based replication, appliance-based replication is still specific to a single platform.
- **More management points.** Appliance-based replication requires dual points of management: the physical management console and the virtualization management console.
- **Growth and change.** Appliance-based replication doesn't "see" configuration changes. As a result, BC/DR plans will be out of sync with the current production environment, eliminating the flexibility of virtualization and removing the ability to respond to evolving business needs.
- **Granularity.** Appliance-based replication focuses on the logical unit rather than the VM. This lack of granularity conflicts with the requirements and promise of virtualization.
- **Costs.** Appliance-based replication also replicates the entire LUN, increasing power, cooling, storage, and networking costs.

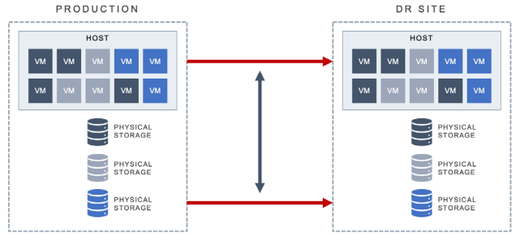

Figure 4. Array- and appliance-based replication methods require coordinating two replication products, one for the physical environment and another for the virtualized environment. This increases management complexity and undermines the investment made in virtualization.

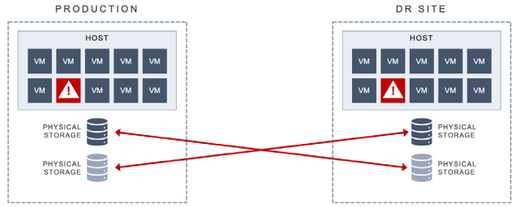

Figure 5. Mirroring systems over a fast network enables a very high availability, but corrupted software components are replicated as well.

Synchronous Replication

Synchronous replication makes a complete copy of an infrastructure on a secondary location and copies or strips every write to that location. Should a disaster occur, then an automatic failover is initiated, and the remote infrastructure takes over. This synchronous replication option, found, for example, in NetApp's MetroCluster, sounds like a perfect, though expensive, solution. However, it is completely based on hardware and is a more highly available solution than a disaster recovery solution. Failover will work in case of a hardware failure, power outage, or a natural disaster, but if the disruption is software based—like a corrupted database or virus—then it will be replicated to the remote site as well. This renders the replication useless, and the team will have to look to their nightly backup for recovery. Synchronous replication is limited by some of the following characteristics:

- **Lock-in.** The secondary location requires an exact copy of the primary location hardware, from the same vendor
- **Cost.** This is an expensive solution, literally doubling hardware costs and in need of a networking solution with a great amount of bandwidth.
- **Incomplete.** Completely hardware-based; in case of a software-based disaster, it falls back on snapshots.

Guest/OS-Based Replication

In an in-guest/OS-based replication solution, software components need to be installed on each individual physical and virtual server. Although this is much more portable than array-based solutions, guest/OS-based replication solutions are not fit for enterprises.

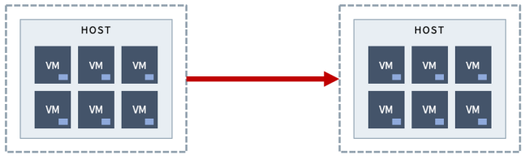

Figure 6. Host-based replication requires an agent on each VM, greatly increasing complexity.

Guest/OS-based replication is limited by some of the following characteristics:

- **Growth and change.** The requirement to install a module on every server limits scalability and makes it impossible to implement and manage in high-scale enterprise environments. Additionally, the overhead of each agent on the VM could present performance issues for applications that the business relies on.
- **Complexity.** Shadow VMs are often part of guest/OS-based replication, increasing complexity and management burden.
- **No application consistency.** Each VM is protected individually, which makes it impossible to manage groups of VMs for one application and replicate it consistently.
- **Management overhead.** All the agents must be managed and maintained. While not too much of an issue in smaller environments, as the environment grows to over 20 VMs, management and maintenance become a much greater problem. Maintenance and updates for the disaster recovery strategy are now a weekend task, often requiring downtime.

Snapshots

Many solutions use snapshots to enable a quick restore. A snapshot is a way to "freeze" a live storage system or VM at a moment in time while changes can still be made beyond the snapshot capture. If changes are made beyond the snapshot capture and the VM or storage system encounters an issue, there is a chance to reject those changes by reverting the VM or storage system back to the snapshot state. A snapshot is especially useful when making changes to a single VM where a rollback may be necessary.

Snapshot-based replication copies snapshots from the production location to a target location, where they can be used to recreate the VM to the nearest viable point in time. Like an incremental backup, a snapshot only contains the data that has changed since the last snapshot, so snapshots can be efficient for both storage and bandwidth efficiency. But compared to continuous, real-time replication, are often performed infrequently, only for backup rather than disaster recovery.

The pros and cons of snapshot-based replication

Pros	Cons
Efficient	**Infrequent**
Because snapshots only contain data that changed since the last snapshot was taken, they have a relatively low storage footprint and are bandwidth efficient for replication. The size of the snapshot will depend on the rate of data change and the frequency of the snapshot schedule. In some environments, such as public clouds, snapshot replication can be an efficient way of reducing data egress charges.	Snapshot frequency is often limited to increments no smaller than 15 minutes because of performance impact. Taking snapshots of many VMs concurrently or time required to complete batches of snapshots and replication between intervals. Compared to continuous data protection or real-time replication technologies, these intervals create a greater gap between recovery points. But, this could also be considered a trade-off for efficiency at scale.
Agentless	**Performance Impact**
VM snapshots do not require an agent to be installed on the VM but instead are performed directly from the underlying virtualization or storage platform. Agent-free replication is a key management benefit for scaling out protection and reducing performance impact on VMs.	Depending on the virtualization platform or storage system, snapshot technologies have varying impact on system performance while taking a snapshot or batch of snapshots. Snapshots can directly impact VM operations with a brief, sub-second freeze, affect storage subsystem performance, or affect CPU performance of the host system.

Disaster Recovery Offered by Hypervisors

Hypervisor vendors like VMware often offer their own software- based replication solutions limited to their own hypervisor. A solution like VMware vSphere Replication (VR) offers limited replication functionality and does not include all the orchestration, testing, reporting, and enterprise-class DR functions that are needed for a complete DR solution. Even combined with VMware Site Recovery Manager (SRM), VR's recovery time and scalability may not be enough to satisfy your organization's needs. While SRM adds capabilities around the planning, testing, and execution of a DR plan, it can't overcome the replication limitations of VR, as VMware vSphere Replication utilizes virtual machine snapshot technology.

Software-Defined Replication

All the aforementioned categories of replication technologies have critical limitations for modern virtualized and cloud-based architectures. They undermine the promises of virtualization and the cloud and limit the capabilities of a multiplatform, hybrid IT infrastructure. To fully realize DR across your modern IT infrastructure without compromising on simplicity and cost, a new approach is required: software-defined replication. Zerto moved replication up the stack—from the storage layer, above the resource-abstraction layer, into the virtualization/hypervisor/cloud layer.

Disaster Recovery: On-Premises or Cloud

One common question about DR is whether an on-premises DR site, a hosted site, or the public cloud is best. Each option can be a good choice depending on several factors.

- **Availability of a suitable on-premises site.** A DR site should be geographically distant from your production site to protect against regional/metropolitan disasters. The site will also require sufficient power, cooling, and connectivity (bandwidth) to be a suitable DR site. Using the cloud quickly eliminates these issues.

- **CAPEX vs. OPEX.** Because of the significant IT resources required to build and maintain its infrastructure, establishing an on-premises DR site can be a large upfront capital expense. A hosted site or public cloud, on the other hand, provides more cash flow flexibility and represents an operational expense in which the cost can be equivalent but also more predictable. There are several key differences between the two models that you should consider before committing to either.

- **Expertise and staffing.** Building and maintaining an on-premises DR site will likely require more staffing and, in some cases, more expertise. With a hosted or public cloud option, the site's maintenance cost is included in the recurring fees, which may eliminate the need for additional staff (or at least minimize it).

On-premises can be a successful and cost-effective choice with the right site and the right staff to execute it, as well as the investment in capital expense. On the other hand, hosted sites and the public cloud offer a ready-to-go site at a recurring expense.

Disaster Recovery as a Service Solutions

Not every organization has a suitable site for DR or the expertise to implement a solution themselves, whether on-premises or cloud. For those organizations, disaster recovery as a service (DRaaS) can be an ideal solution. DRaaS provides a resilient, remote, cloud-based location for DR; DR specialists who manage the solution; and a predictable recurring cost.

But not all DRaaS solutions are created equal. Different DRaaS solutions have different degrees of service management. With some services, you may self-manage part of the DR solution. Depending on your organization's needs, you may benefit more or less from different levels of management.

As with any DR solution, the ultimate effectiveness of DRaaS will depend on the underlying technology. Unfortunately, many DRaaS solutions still use legacy backup and replication technology that provides poor RPOs and RTOs. Evaluating recovery times and points is crucial to not only meeting your SLAs, but also reducing the cost of the solution when a disaster strikes. Downtime costs and lost data—which, with legacy technologies like snapshot-based replication or slow manual recoveries, is a greater risk—add to the cost of the solution.

To help you evaluate DR solutions, including DRaaS, we set up a requirements checklist, found on the following page.

Disaster Recovery Requirements Checklist

For both in-house and DRaaS solutions

Performance

1. Does the DR solution offer continuous replication? What impact does the particular technology (e.g., snapshots) have on the production site?
2. What RTOs and RPOs does the solution offer? Are they measured in seconds, minutes, or hours? Can this be proven, and do you have continuous insight into them?
3. Do these RPO/RTO numbers realistically meet your business requirements, and at what sacrifices or costs?
4. **DRaaS**—Does the Cloud Service Provider offer a reliable and fast networking solution, and does the DRaaS solution offer networking efficiencies like compression?

Support of your systems

5. Is the DR solution storage- and hypervisor-agnostic? In other words: can you replicate from any environment to the DR solution?
6. Is it app-centric? Does it offer recovery from a consistent point in time even for multi-VM apps relying on disparate hosts and datastores?
7. How scalable is the solution (up and also down in a DRaaS environment)?
8. What does the installation look like? Will you need to reconfigure applications, LUNs, VMs?
9. Does it support change, like when VMs are moved to other storage locations or when you want to do a migration?
10. Does it provide flexibility and choice to change which public cloud to use as a target DR site?
11. **DRaaS**—Does it support multiple sites, and is it multitenant? Does it offer securely isolated data streams for business-critical applications and compliance?

Functionality

12. Is it a complete off-site protection solution, offering both DR and archival (backup) storage, with very limited impact on the production site?
13. Is it suited for both hardware and logical failures?
14. Does it offer sufficient failover and failback functionality, including recovery automation and orchestration, pre- and post-recovery scripts, automatic IP adjustment, etc.?
15. In case of a failover or failback, how does it impact production? And what does the failback process look like? Is it similar to the failover process?

Compliance

16. Can it be tested easily, and are testing reports available? What is the impact of the test? Is this something that can be done during business hours, or is this a weekend activity? Does production need to be taken down? Is replication paused or broken during testing, impacting the DR solution during every test?
17. **DRaaS**—Are there any license issues or other investments upfront?
18. **DRaaS**—Where is the data being kept? Does the service provider comply with EU regulations?

Usability

19. Is it easy to learn and use? Does it add more management control points to your environment, or does it integrate seamlessly?
20. Does it offer the right recovery granularity? Can you recover a file, single VM, single application, a few applications, or the entire site?
21. **DRaaS**—Does the DRaaS solution offer both self-service and managed services?

SECTION 3

Zerto

Zerto is based on a foundation of continuous data protection (CDP). It brings together DR, ransomware resilience, and multi-cloud mobility. Zerto users benefit from a unified and automated recovery and data management experience across virtualized and cloud workloads.

Zerto enables an always-on user experience by simplifying the protection, recovery, and mobility of applications and data across private, public, and hybrid clouds.

Benefits of Zerto Disaster Recovery

Zerto CDP technology delivers the best RPOs and RTOs in the industry to mitigate any disruption, including natural disasters, hardware failures, ransomware, and other planned or unplanned outages. This provides you with the confidence that data loss and downtime won't interrupt your continuous business. Additionally, Zerto CDP works across clouds and platforms to help you unlock the promise and potential of a hybrid and multi-cloud world.

Continuous Data Protection

- **RPOs of seconds.** Unlock RPOs of seconds via always-on replication and continuous backup locally, remotely, and across platforms.
- **RTOs of minutes.** Fail over an entire site or select workloads to a remote site in just minutes with only a few clicks.
- **Near-synchronous replication.** Deliver the best of both synchronous and asynchronous approaches. Sitting at the hypervisor level, Zerto is a software-only solution totally independent of the underlying hardware and infrastructure, including storage.

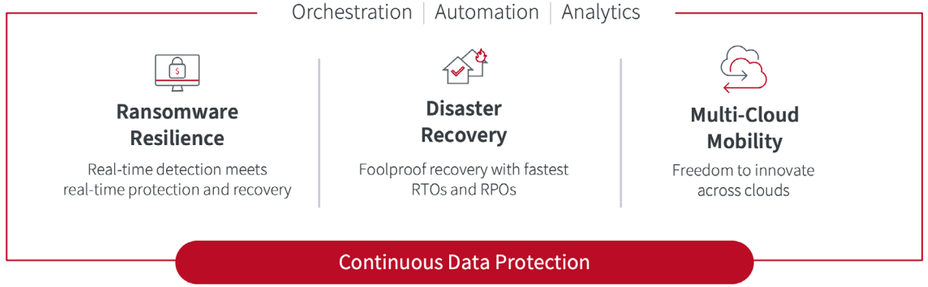

- **Unique journaling capabilities.** Zerto provides continuous block-level replication with zero impact on application performance. All checkpoints, seconds apart, are stored in the journal for up to 30 days to deliver RPOs in seconds and restore files, folders, and VMs—even entire applications and sites, with virtually no data loss.

- **Application-centric protection.** Enterprise applications contain multiple VMs and dependencies. Traditional methods like incremental backup jobs bring significant challenges to recover applications consistently and with an acceptable RTO. Zerto resolves this with Virtual Protection Group (VPG) capability. VPGs allow you to protect and recover an entire application (and all its VMs) in one click, at one consistent point in time, with write-order fidelity.

- **Ransomware recovery down to the second.** CDP delivers a continuous stream of recovery checkpoints. In the event of ransomware or other malicious attacks, data can be recovered to just seconds before the corruption took place, minimizing impact to the business and the brand.

- **Real-time encryption detection.** Detection and alerting capabilities warn users of encryption-based anomalies to pinpoint and mitigate the earliest stages of a ransomware attack.

Multiplatform, Multi-Cloud Support

Zerto supports multiple platforms and multiple cloud configurations. With Zerto, you get:

- **Hardware and hypervisor agnosticism.** Remove barriers to innovation with a replication solution that has no hardware or hypervisor dependencies and no vendor lock-in.

- **Simple and seamless installation.** Install seamlessly within minutes into the existing infrastructure with no downtime or configuration changes required.

- **Scalable and granular.** Scale to thousands of VMs. Use cloud instances for the granularity to manage each workload individually.

- **One-to-many, multi-platform.** Easily replicates to the public cloud, a service provider, or a secondary site without disruption—even simultaneously for additional recovery options.

- **SaaS backup for business-critical data.** In addition to providing protection for virtualized workloads across on-premises and cloud platforms, Zerto Backup for SaaS, powered by Keepit, provides backup for your data on platforms such as Microsoft 365, Salesforce, Google Workspace, Microsoft Azure AD, and Microsoft Dynamics 365. This data can be just as valuable as any on-premises data you may be protecting. Zerto provides peace of mind, so you know your data from these platforms is protected in the event of any type of disaster or disruption.

Management and Orchestration

Automation and orchestration make things simple with Zerto, giving you:

- **Automated VM protection.** Automatically protect VMs using vSphere tags to ensure complete, flexible data protection across your environment, even when you add new VMs.

- **Simple and centralized management.** Manage centrally across on-premises and cloud. Use complete APIs for integration and automation.

- **Disaster avoidance.** Proactively failover or migrate to other sites before an incident occurs, such as an incoming hurricane or other foreseeable potential disaster.

- **On-demand sandboxes.** Create sandbox versions of your production environment for nondisruptive DR testing, update and patch testing, malware scanning, and more. Test the full recovery process without impacting production environments or ongoing replication at any time, giving your team confidence in the event of any disaster.

Compliance

Zerto's capabilities facilitate various regulatory X, from SLAs to governmental requirements. Stay compliant with:

- **Reporting and analytics.** Includes automated disaster recovery reporting for compliance and health checking.

- **Ransomware resilience.** Restore entire applications, databases, or individual files with consistency and granularity. Perform nondisruptive failover tests at any time, boosting confidence that you can bring your business back online immediately.

- **Enterprise-class support.** Get enterprise-class support services that are built into all Zerto products. Support services include real-time alerts when RPO/RTO targets are not being met, network degradation alarms, and reminders to check configurations and VPGs. Zerto solutions are backed by global support service centers that provide on-demand access to an expert team of support engineers.

Zerto Architecture

The heart of Zerto CDP and replication technology is formed by two components:

- **Zerto Virtual Manager.** The Zerto Virtual Manager (ZVM) is a security-hardened virtual appliance that integrates with the underlying platform to manage replication, paired sites, VM protection, and performance. If any problem occurs, the ZVM represents it visually and sends alerts as well. The ZVM is also responsible for orchestrating and automating failback and recovery processes like boot order, re-IP, scripts, test, and validation options

- **Zerto Virtual Replication Appliance.** The Virtual Replication Appliance (VRA) is a VM that runs on each virtual host or cloud environment. It provides always-on, block-level replication to one or more remote sites. Unlike agent-based replication, the VRA offloads replication from within VMs, eliminating the performance impact on those VMs. And unlike snapshot-based replication, the VRA can replicate data every five seconds, delivering near-instant RPOs.

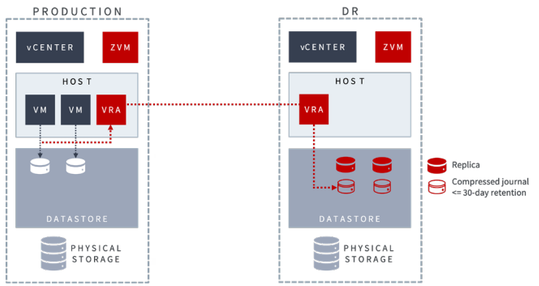

Application-Centric Protection: Virtual Protection Groups

Many enterprise applications consist of more than one virtual server—for example, web, application, and database servers—which are interdependent. When recovery is needed, all servers must be recovered from a single, consistent point in time. To do that, Zerto developed VPGs, which ensure consistency across the disks in a group of VMs. VPGs replicate and recover enterprise applications with consistency, regardless of the underlying infrastructure. Zerto recognizes and preserves these relationships while enabling critical VMware features such as DR, vMotion, and Storage vMotion.

Zerto VPGs are designed to allow replication and recovery that is:

- **Consistent.** Replicates and recovers complete multi-VM applications consistently.
- **Flexible.** Enables organizations to deploy an application across different physical devices to maximize performance or capacity or to reduce the complexity of the infrastructure.
- **Granular.** Delivers the right granularity to recover single VMs, as well as groups of VMs, through many types of disasters.
- **Prioritized.** Prioritizes VPGs for replication and recovery.
- **Supported.** Supports virtualization features like vMotion, svMotion, HA, etc.

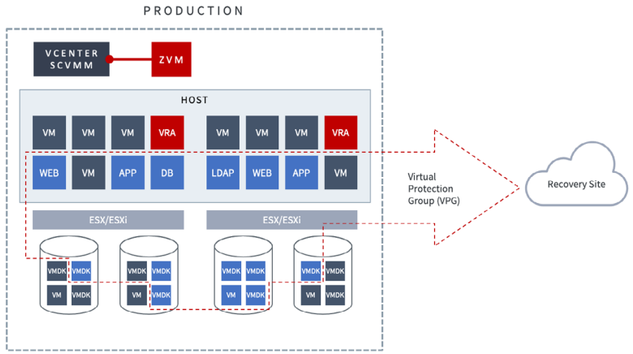

Figure 7. The various VMs comprising an application are in a VPG and replicated consistently even if they are spread over various hosts and datastores.

Automatic Virtual Machine Protection

In today's data centers, new VM workloads are created and deployed quickly and frequently, especially in virtualized environments. But going through a management interface workflow to protect each newly created VM can be overly burdensome. By utilizing VM tags within vSphere, Zerto allows the automatic creation of VPGs and protects VMs without having to open the ZVM. System administrators can ensure VMs are protected upon creation and assigned to a current or new VPG without needing to open another management interface.

Fully Configured Failover Process

One step of VPG configuration is setting up the failover process. As part of this configuration boot order, re-IP on failover, length of journal, and other parameters are calibrated. With all this up-front work complete, the recovery process is simplified, reducing it to just a few clicks.

Fully Automated and Orchestrated

In DR, replicating data to the recovery site is only half the issue. The other half is quickly and easily using that information to protect a business during a disaster. Zerto recognized this issue and built automated and orchestrated processes you can execute in just a few clicks when IT is in the middle of a high-pressure situation.

Failover as a Business Decision

Since every disaster is different, Zerto believes failover needs to be a business decision and not an automated process. Because you can pick a moment in time—the point in time just before a database corruption occurred—this decision phase is essential for a correct failover. After clicking the failover button, an automated and orchestrated process will start to bring services back online.

QUICK 4-STEP FAILOVER PROCESS

The failover process consists of four simple steps. After an incident is visible in the management console:

1.Click Failover ⟶ Live.

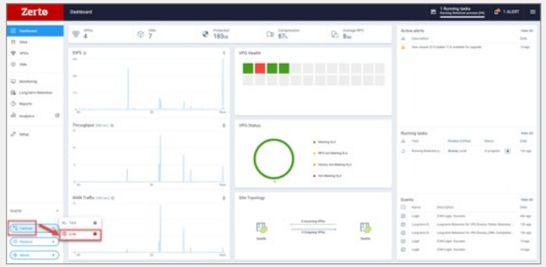

2. Select the applications (VPGs) that need recovery from the list.

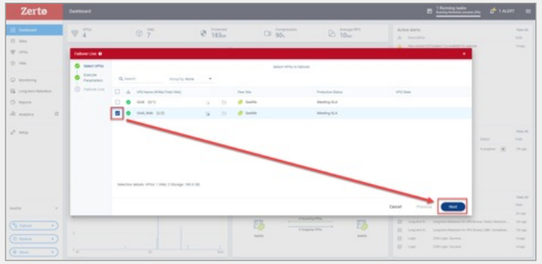

3. Verify the point in time to which the apps need to be recovered. Options are to recover to:

- Latest checkpoint (default)
- Latest tagged checkpoint
- Latest VSS checkpoint
- Specific point in time from the list of all available checkpoints

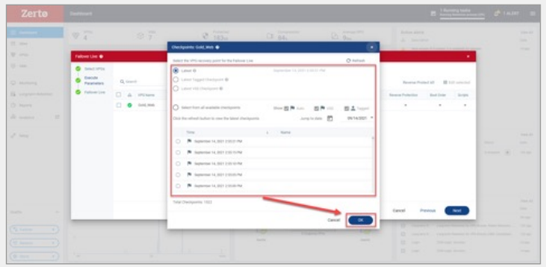

4. Start failover process. The recovery process begins and VMs are booted in the recovery environment and reconfigured as needed.

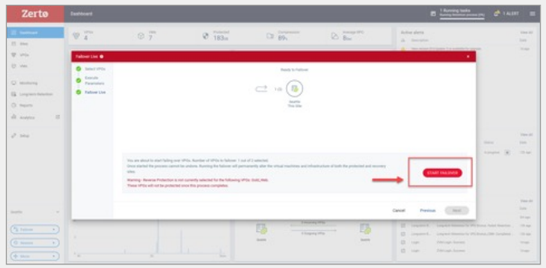

Automated Failover and Failback

Upon configuration of the VPGs, the recovery plan is now in place. Pre- and post-recovery scripts can also be configured on a per-VPG basis. Now, failover and failback are executed in just a few clicks. Even when the DR process is initiated, there is an opportunity to roll back the failover should there be issues at the recovery site unrelated to Zerto, like a network being down. Upon a successful failover, reverse protection makes the failback process even easier. When the production site is ready for use, reverse protection begins syncing the additional work done at the recovery site to the production site. After the applications have been updated to the original production site, failback happens in just a few clicks. Many organizations will not fail over because failback is so cumbersome—but with Zerto, everything is easy.

File and Folder Recovery

The most common disasters that administrators need to recover from are not natural disasters or site outages but lost or accidentally deleted files or folders. Zerto solves this common problem with single-file and -folder recovery from up to 30 days in the past using the journal. CDP delivers recovery points just a few seconds apart, enabling IT to go to the point before the file was deleted or corrupted and recover it. This is executed in just a few clicks, and all work lost is minimized.

Zerto CDP:

- **Minimizes risk.** The ability to recover at any level at any point in time minimizes data loss across files, folders, VMs, applications, and sites.
- **Increases simplicity.** Automated workflows for file, application, and data recovery reduce mean time to recovery.
- **Protects productivity.** Preserves productivity and employee morale since end-users no longer need to recreate hours or a day of lost work when a file or folder is accidentally deleted.

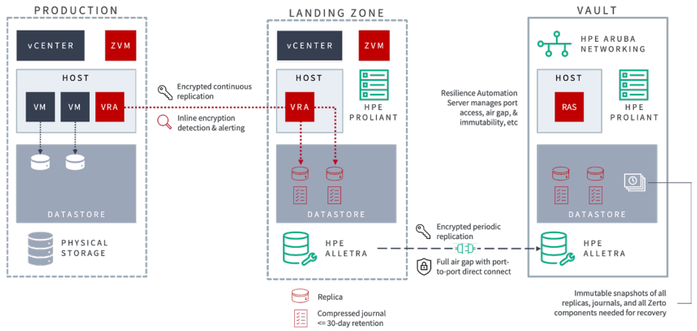

Cyber Resilience Vault Architecture

Ransomware Resilience

Ransomware attacks have become increasingly common in recent years, targeting organizations of all sizes and industries. The consequences of a ransomware attack can be devastating, with organizations losing access to critical data and systems, even facing financial losses and reputational damage. To combat the ever-growing threat of ransomware, Zerto uses a combination of advanced technologies and best practices.

- **Real-time encryption detection.** The Zerto encryption analyzer instantly detects and alerts about suspicious write activity on protected workloads. Get the earliest warning sign of malicious anomalies with the in-line encryption and alerting system. It is no longer necessary to wait to detect ransomware after backing up—detection occurs within seconds at the first moment of impact.

- **Enhanced detection.** A Zerto on-demand sandbox can be created quickly and nondisruptively, copying your entire environment to perform malware scanning that will not impact production. Security scans can be resource intensive, slowing down production machines, but in an isolated sandbox that is a duplicate of production, malware scanning can be performed quickly and with no production impact.

- **Minimal data loss.** With Zerto journaling technology, data can usually be recovered seconds before the ransomware attack happens so data loss is minimized.

- **Multiple options for rapid recovery.** Zerto provides multiple recovery options from local, warm-site, and cold-site data. Recovery options include rapid failover, continuous point-in-time journaling, file and folder recovery, full VM recovery, and orchestrated recovery.

- **Data copies across multiple platforms.** With Zerto, workloads can be protected across multiple platforms, including on-premises and in cloud, which mitigates the likelihood of a ransomware attack reaching backup copies.

- **Isolate and lock with a Cyber Resilience Vault.** The Zerto Cyber Resilience Vault provides recoverability after even the worst attacks. Its completely isolated, air-gapped recovery environment stores immutable copies on secure, high-performance hardware. The Zerto Cyber Resilience Vault uses zero trust architecture and a combination of best-in-class software and hardware to provide a highly secure clean room—all while enabling rapid recovery in minutes or hours, not days or weeks.

- **Recovery in an isolated network.** Ransomware recovery can be difficult because malware may exist even in recovered data. Zerto deploys isolated recovery environments to validate recovered data and systems, which enables you to scan for malware and remove it before recovering fully to production.

- **Immutable replicas.** In addition to multiple copies across multiple platforms, Zerto supports immutable replicas in the cloud as a last resort against ransomware attacks to ensure the replicas cannot be compromised.

Analytics and Reporting

The Zerto SaaS-based analytics built into the solution includes out-of-the-box dashboards and reports. These provide complete visibility across multisite, multi-cloud environments to meet your SLAs and deliver easy, hands-off compliance reporting. With the Zerto mobile app, you can monitor SLAs from anywhere. But Zerto goes beyond visibility alone: it gives you tools to do intelligent, predictive infrastructure planning, deepening your ability to optimize and be proactive about your data protection strategy.

Zerto In-Cloud for AWS

Zerto In-Cloud for AWS enables highly scalable, orchestrated DR to and from AWS Regions and Availability Zones. Built for scale, Zerto In-Cloud for AWS protects thousands of instances across accounts with speed of recovery, simplicity, and with no agents to manage. API-centric management allows Zerto In-Cloud to easily integrate with the automation tool of your choice (i.e., Ansible, Jenkins, and Terraform) to make DR part of your automated management strategy.

- **Scalability.** Zerto In-Cloud's agentless, native integration approach enables easy scaling to protect more than 1,000 workloads, promoting resilience for enterprises of all sizes in Amazon EC2. Whether you have one AWS account or hundreds, you can easily orchestrate protection across your organization.

- **Orchestration.** Zerto In-Cloud automates protection of EC2 instances, eliminating the need for excessive, manual, and error-prone steps to execute recovery plans. Failover and failover testing features are then automated for scalability and simplicity.

- **Simplicity.** Using native, agentless AWS integration rather than creating another layer of software across AWS, Zerto In-Cloud takes advantage of Amazon EBS snapshots and native AWS replication. Analytics provide you with insights and reports on RTOs and RPOs.

- **Nondisruptive Failover Testing.** Zerto In-Cloud automation and Zerto orchestration makes it possible to nondisruptively failover test all instances in an AWS Region or Availability Zone to validate recovery plans and readiness when an outage or disaster strikes.

- **Flexible Management.** A complete REST API is at the forefront of Zerto In-Cloud management so it can be easily integrated with other management systems and combined with other automation and integration solutions.

- **Application Groups.** Zerto In-Cloud orchestration enables application-centric VPGs to recover applications as a single entity to another region or availability zone, enabling you to protect and recover large complex applications like SAP.

- **Stateless Management.** The Zerto In-Cloud Manager appliance runs from any Amazon EC2 region and can be quickly redeployed to another region in the event of an outage to recover and manage workloads.

Zerto In-Cloud Architecture

The Zerto In-Cloud Manager is a Linux-based virtual appliance that works with the Amazon DynamoDB and Amazon EBS to provide automation of snapshots and replication for EC2 instances across all associated accounts. The Zerto In-Cloud Manager is stateless, resides in any EC2 region, and quickly redeploys into another zone or region if needed. Only a single Zerto In-Cloud Manager appliance is needed for management, regardless of how many EC2 instances are protected.

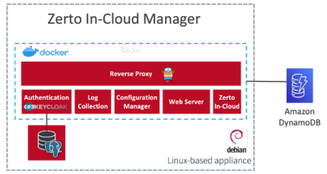

Figure 8 – Using EBS snapshots and native AWS replication, Zerto In-Cloud protects Amazon EC2 instances from local and regional outages and other potential disasters

Zerto In-Cloud coordinates Amazon EBS snapshots and replication to protect instances across regions. Should an outage, disaster, or ransomware attack occur, your data and applications are protected to a separate region or zone to be brought online again quickly.

Applications and the instances they run on in Amazon EC2 are protected in VPGs that coordinate replication and recovery across the entire application group. This application-centric approach to protection means entire applications are restored with consistency and write-order fidelity. Analytics, included in the offering, provide insights and reports on RTOs and RPOs for your VPGs.

SECTION 4

Summary

Planning for disaster recovery is a complex process that involves many considerations. Choosing the right DR solution to meet your needs can be difficult with so many emerging and legacy solutions to choose from, and choosing the wrong solutions can end up costing millions in losses due to downtime and data loss.

By understanding the key considerations in this guide and choosing a DR solution like Zerto, you can not only meet all your SLAs for RTOs and RPOs easily, but also gain many more benefits over the legacy solutions that have not kept up with today's needs. Look to Zerto to provide you with these benefits and more, now and in the future.

To learn more about Zerto, take our interactive product tour or sign up for our free trial.

About Zerto

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 eliminates the risk 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 solve for ransomware resilience, disaster recovery, and multi-cloud mobility. Zerto is trusted by over 9,500 customers globally and is powering offerings for Amazon, Google, IBM, Microsoft, and Oracle and more than 350 managed service providers.
www.zerto.com

RITM0066172