Meriplex

Disaster Recovery Guide

DR in Virtualized Environments

Powered by Zerto

Disaster Recovery Guide
DR in Virtualized Environments

CONTENTS

PREFACE
Disaster Recovery in a Virtualized World

In today's always-on, information-driven organizations, business continuity depends completely on IT infrastructures that are up and running 24/7. The costs of downtime are huge and data loss can put a company out of business. Data loss is not only caused by natural disasters, power outages, hardware failure and user errors, but more and more by software problems and cyber security related disasters. Therefore thorough security and business continuity strategies are crucial for modern businesses, minimizing data loss and downtime. Especially now, as data centers become more and more software-defined, these private, hybrid and public clouds become more vulnerable to these kinds of threats.

In a software-defined, virtualized environment, applications run on virtual machines (VMs), independent from the hardware. Though this brings a lot of efficiency benefits to the business, these benefits are not extended into the realm of disaster recovery (DR) and business continuity (BC). Most DR/BC solutions are still based on physical entities, arrays and appliances, and lack the ability to scale with the amount of data modern organizations produce. Many of the benefits achieved through virtualization, therefore, can be lost because of the management overhead and the complexity of aligning a virtualization strategy with disaster recovery tools designed for physical environments. Virtualization-aware BC/DR solutions are needed to overcome this.

In this booklet we provide insights in the challenges, needs, strategies, and available solutions for disaster recovery and business continuity, especially in modern, virtualized environments and the public cloud. We will also explain which benefits and efficiencies Zerto Virtual Replication delivers and how it compares to other BC/DR technologies. With this overview we want to provide businesses with the right information to choose the best possible BC/DR solution for their situation. If reading this booklet leads to any questions, please contact us at **connect@meriplex.com**.
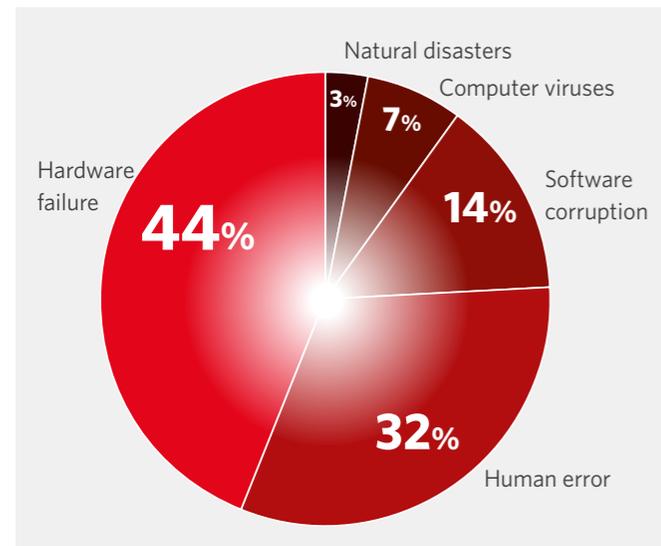
TRY IT YOURSELF
Zerto Virtual Replication can be installed and configured in under 1 hour. With simple VM-based replication enabling RPOs of seconds and RTOs of minutes. Go to www.zerto.com/trial and download a free trial today!

Disaster Recovery: Needs & Technologies

The causes & costs of data loss

Modern businesses can't afford to lose data. Whatever the cause – natural disaster, human error, or cyber attack – data loss is costly and extremely risky. Research from various institutes shows that the volume and costs of data loss are increasing year over year. The need for a business continuity strategy to ensure uptime, diminish data loss, and maximize productivity in the midst of any compromising situation is a necessary digital assurance policy for any company. Because the question is no longer *if* a disaster will strike, but *when*.
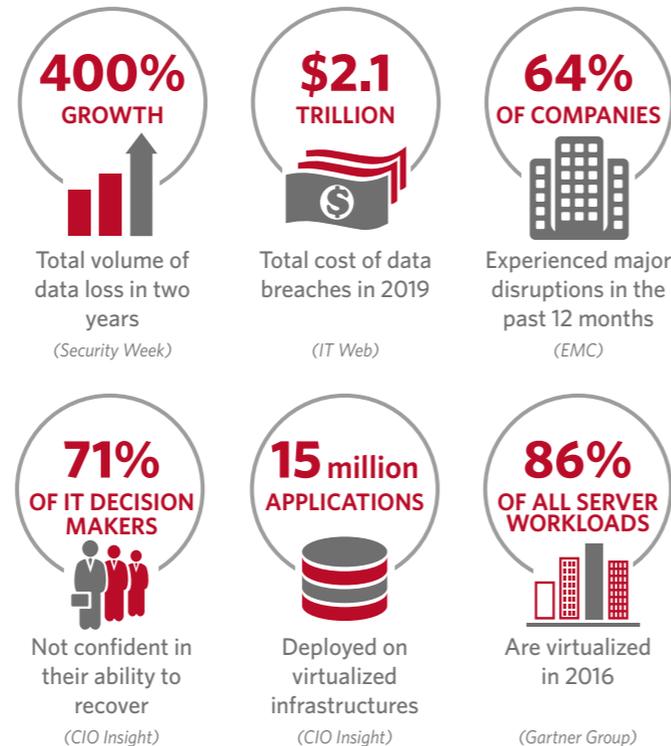

TOP CAUSES OF DATA LOSS AND DOWNTIME

(source: World Backup Day 2015)

Natural disasters 3%
Computer viruses 7%
Software corruption 14%
Human error 32%
Hardware failure 44%

...in a virtualized world

Many organizations have virtualized their production environment and gained real efficiencies and realized measurable savings. However, many of the management efficiencies are lost in the BC/DR sphere, as BC/DR technologies are typically based on dated technologies – array-based replication or agent-based replication – that are not virtualization-aware. As more and more technologies are leveraged within the disaster recovery planning, it is very difficult to create a consistent and repeatable disaster recovery plan.

400% GROWTH
Total volume of data loss in two years
(Security Week)

$2.1 TRILLION
Total cost of data breaches in 2019
(IT Web)

64% OF COMPANIES
Experienced major disruptions in the past 12 months
(EMC)

71% OF IT DECISION MAKERS
Not confident in their ability to recover
(CIO Insight)

15 million APPLICATIONS
Deployed on virtualized infrastructures
(CIO Insight)

86% OF ALL SERVER WORKLOADS
Are virtualized in 2016
(Gartner Group)

The concept of Disaster Recovery

What is Disaster Recovery (DR)? Literally, it means to recover from a disaster, the time and labor it takes to be up and running again after data loss or downtime. This time depends on the solution that is chosen to protect the business against data loss. DR it is not just about the time systems and employees cannot work, it is also about the amount of data lost when having to fall back on a previous version of the data. Businesses should ask themselves, how much an hour of downtime will cost. And foremost: is it possible to remember and reproduce the work employees (or systems) did in the last few hours? 95% of all companies are not able to answer this question...

Backup is not a true disaster recovery solution
Disaster recovery involves many different concepts, which might get confusing: disaster recovery, business continuity, backup, RTO and RPO. Best known is **data backup**, which means replicating data or VMs to another device or location, done consistently at a regular interval (say, every 24 hours), for recovery, or as a long-term retention solution for compliance. However, in case of a disaster, backup is an empty concept without a **disaster recovery (DR)** solution, the ability to recover the files, software and functionality. Usually this consists of more than just copying the data back to its original system. If a server is down, it needs to be re-installed, reconfigured and maybe even replaced. Backups are not a true disaster recovery solution. With a backup, VMs must be completely reconstructed as there is no automation built into a backup process.

RESPONDING TO RANSOMWARE

The last few years a rising trend is visible, of hackers attempting to extort money from both private users and businesses, via various Ransomware Trojans such as CryptoLocker. These malicious pieces of software encrypt data, files, or even complete server systems, by generating a private-public pair of keys. The data is impossible to decrypt without the private key, which is stored on the attacker's server until the ransom is paid.

The best way to protect businesses against this threat is to use up-to-date Anti-Virus and security software, a solid recovery strategy, as well as educating users on the threat of phishing mails, where the viruses usually originate from. In the unfortunate case a business is the victim of a ransomware attack, Zerto can help mitigate data loss:
- Rewind the systems to the last point-in-time before the infection struck, to within a matter of seconds.
- Recover all the critical systems within the space of a few minutes, with only a few clicks of a button.
- Not only restore entire applications and databases with consistency, but restore individual files as well.
- Perform non-disruptive failover tests at any time, to be sure the business can be brought back online straight away when needed.
- Create off-site data copies for longer-term data retention in addition to giving the business a Continuous Data Protection for up to 14 days.

Business Continuity

Many companies have a – preferably remote – **disaster recovery site** where data is replicated to on a continuous basis, ready to be leveraged in the event of an outage. If this disaster recovery site is on a remote location it can also provide **business continuity (BC)**: the ability for a business to continue to operate after a major disaster, like a fire, power outage or a natural disaster. In case the original site is down, the services on the production site can be run on the DR site. This switching process is called **failover**. Once the normal production site is back up and running, the work that has been done at the disaster recovery site must be replicated back to ensure that all that work is not lost. The ability to **failback** applications and data from the DR site to the production site is a critical attribute of a solid disaster recovery solution.

DR sites used to be a copy of the production site in another office location, but nowadays they are often located at the datacenter of a cloud service provider, or in the public cloud.

RTO AND RPO EXPLAINED
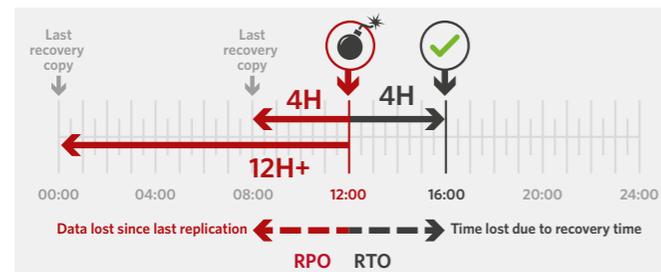
Any enterprise with a **stock market quotation** has to comply with rules regarding data security; loss of data will result in loss of revenue, reputation and shareholder value.

An **online business** that loses 4 hours of business data might end up with angry customers wondering when their bought and paid for goods are coming.

If a **transport company**'s systems are down for a few hours, it is impossible to plan deliveries and pick-ups efficiently, which has an enormous effect on revenues that are already under pressure

Complex **robotized production processes** that are down after a hardware or software failure, cause enormous loss of productivity and revenue.

RTO and RPO

When it comes to business needs, translated into Service Level Agreements, recovery is usually expressed in two types of objectives: RTO and RPO. **The Recovery Time Objective (RTO)** is the amount of time the business can be without the service that needs to be recovered, without significant losses or risks. The **Recovery Point Objective (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 up to as long as 24 hours. In modern, ubiquitously digital enterprise environments both RTO and RPO need to be as low as possible, no longer expressed in hours but in minutes or even seconds. Though many organizations focus on RTO to get the business up and running as soon as possible, it is the inability to reproduce the loss of data (RPO) that will haunt an organization for a long time after any disaster.

High Availability

A concept that is often confused with disaster recovery and business continuity is **high availability (HA)**. This is functionality that helps avoid downtime by hardware issues, and involves technologies like RAID and redundant parts like power supplies and cabling, but can be applied within virtualized environments as well. HA technologies are necessary to keep systems running, but will not help recover after a disaster. High Availability is mostly expressed in a percentage, somewhere in the 99%. But don't forget that 99.9% uptime still means that a system has 8 hours of unplanned downtime in a year.

HIGH AVAILABILITY IN % AND IN TIME *(source: Wikipedia)*

Availability %	Downtime per year		Per week	
90% ("one nine")	36.5	days	16.8	hours
99% ("two nines")	3.65	days	1.68	hours
99.9% ("three nines")	8.76	hours	10.1	minutes
99.99% ("four nines")	52.56	minutes	1.01	minutes
99.999% ("five nines")	5.26	minutes	6.05	seconds

Disaster Recovery as business strategy

Since data loss and downtime have direct impact on the business, disaster recovery is an issue that should be decided upon based on strategic business criteria and goals. Questions like how much downtime a business can survive and how much data loss would be acceptable – setting RTO and RPO – are impossible to answer from only a technical level. The answers depend on revenue streams that come from IT systems, the value that is associated with corporate data, logistics and other business processes that heavily depend on IT. The bottom line: though a lot of technology is involved, disaster recovery is a key element to effectively support business goals.

Decide what is truly critical

When it comes to developing a disaster recovery strategy, it is important to realize that not all systems, applications and data are equally business-critical. For the core applications, a working DR strategy involving a remote DR site, low RTO/RPO (low data loss and short recovery time), and a tested recovery plan, is essential. For other applications and types of data, less expensive solutions and higher RPO/RTO might be more acceptable.

Prioritization is a key element for disaster recovery planning. Review what downtime can be tolerated for each application with line of business owners. It will become clear which ones need to be available fast with minimal data loss. There may even be some applications that a business can do without for hours. The important step here is to set the service level across the board, so there is agreement and there are no surprises when a disaster occurs.

BC/DR is a financial decision

To design a DR solution various structures and many solutions are available. They all come with a cost. The cheapest solution is probably an old-fashioned backup, but this is not enough in modern environments. When implementing a remote DR site, there is a choice between a standby copy of the production site at another location, or a cloud-based service (DRaaS). Choosing between these two is also a choice between Capex and Opex, between the expenses and costs of owning a solution or the operational costs of an online service.

Another consideration is the number of tools that will be included in the BC/DR plan. 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 when they are the most costly.

DR is governance

When it comes to DR larger enterprises encounter compliance and governance issues as well. Data regulations are becoming stricter and meeting them must be part of a DR strategy as well. To comply, procedures need to be documented and solutions need to be tested and reliable. The question whether and where data is stored in the cloud and who is in control of the cloud data is a question of complying as well.

DR technologies and virtualized environments

Virtualization of the data center has proven to be a true IT game-changer, providing increased flexibility and control in managing production workloads, as well as significantly streamlining the implementation and operational support. To fully realize the benefits of this software-defined environment – this private or hybrid cloud – organizations need to optimize all IT processes and activities for the virtual environment: security, compliance, and Business Continuity/Disaster Recovery (BC/DR). When it comes to BC/DR, many organizations still see this as a costly insurance policy, especially since the available solutions are often very expensive and inadequate in a virtualized environment.

Hardware and software
Many DR solutions focus on minimizing downtime based on hardware failures, power outages or natural disasters. Most data center issues aren't whole data center outages. IT is typically tasked with recovering an accidentally deleted file or even a single VM that was patched and is now acting out or other smaller events. This means that disasters are not always caused by hardware failures nor is it strictly solved by hardware solutions.

What makes BC/DR in a virtualized environment different?
- **Software-defined** – In a virtual environment, replicating at the hardware layer is not adequate. Replication must take place in the hypervisor, so that the business needs will be available in the event of a disaster. The end-users aren't looking for a logical storage unit, they are looking for the application - Oracle, Microsoft Exchange, or another application.
- **Virtualization-aware** – With a virtualization strategy in place, the disaster recovery solution needs to be virtualization-aware as well. This ensures that any changes in the production environment will be reflected in the disaster recovery strategy, ensuring protection is not compromised while still delivering the flexibility and agility that virtualization offers.
- **Application consistency** – Many critical applications use more than one VM and these are interdependent. That means they need to be replicated together to remain consistent.
- **Scalability** - It is critical to look for a solution purpose-built to protect many VMs. Data and applications are growing at exponential rates. The disaster recovery solution needs to scale with the growth without adding significant complexity and overhead.
- **Change** – Due to their dynamic nature, virtual environments tend to sprawl which makes BC/DR more complex.
- **Granularity** – To be able to respond to some of the most important causes of data breaches, like data corruption and accidental user errors, granularity is needed to be able to restore single VMs or files.
- **Frequency** – Now IT environments – not only the virtualized ones – are crucial for businesses to survive, replication of VMs, data and files needs to be done in a high frequency. A daily backup won't do.

Available DR Solutions
A number of disk-to-disk based DR solutions haven been introduced over the years, none of them fully virtualization-aware. In a short overview we will summarize their structures and their shortcomings in a virtualized environment.

Array-based Replication
Array-based replication products are provided by the storage vendors and 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, whether it is 40% or 90% utilized.

- **Hardware-defined** – Array-based replication is designed to replicate physical entities. It doesn't "see" the virtual machines and is oblivious to configuration changes.
- **Not independent** – Though optimized to work with the existing storage array, it locks in the organization 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** – Replicating the entire LUN, array-based replication lacks the granularity needed in a virtual environment.
- **Costs** – The entire LUN is replicated, whether it is 40% or 90% utilized, increasing power, cooling, networking and storage costs.
- **Single point for recovery** – Many array-based solutions do not have the ability to store a history of the performance of the LUN. With that limitation, if the last data point was corrupted, that is what the business has to use for recovery, rendering the DR solution useless.
- **Time** – Recovery is very time-consuming and complicated as there is no automation so the VMs and applications have to be built from scratch.
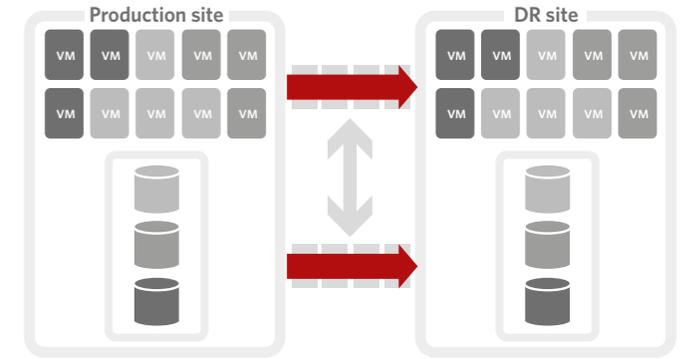

Figure 1. Array-based and appliance-based replication requires coordinating two replication products, both for the physical and the virtualized environment. This increases management complexity and undermines the investment made in virtualization.

Appliance-based Replication
Appliance-based replication solutions are also hardware-based and specific to a single platform. The main difference is that replication runs on an external, physical appliance instead of inside the storage array itself. This makes it more flexible and less-consuming in array resources. But the disadvantages are more or less the same as for array-based replication.

- **Hardware-defined** – It is also designed to replicate physical entities rather than virtual entities.
- **Not independent** – Though it is more flexible than array-based replication, it 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** – It 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 virtual machine. This lack of granularity conflicts with the requirements and promise of virtualization.
- **Costs** – Since the entire LUN is replicated, costs for power, cooling, storage and networking increase.
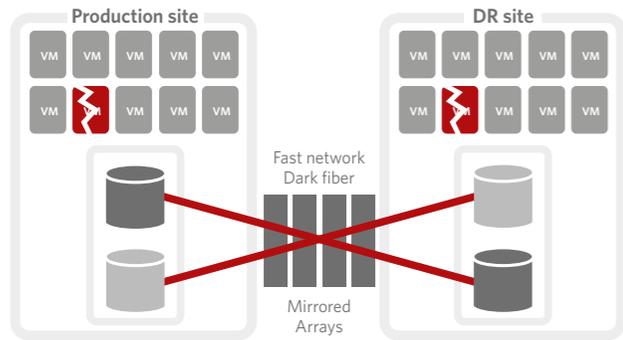
Production site   **DR site**

Fast network
Dark fiber

Mirrored
Arrays

Figure 2. By mirroring systems over a fast network a very high availability is possible, but corrupted software components are replicated as well.

Synchronous Replication

Another option is to have a complete copy of an infrastructure on another location, and copying or striping every write to that location as well. In case of a disaster 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, but it is completely based on hardware and is more a high availability (HA) solution than a DR solution. Failover will work in case of a hardware failure, power outage or a natural disaster, but if the problem is a corrupted database, a virus or any other software-based issue, these issues will have been 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.

- **Lock-in** – An exact copy of the hardware is needed, from the same vendor, on another location.
- **Expensive** – Of course 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.
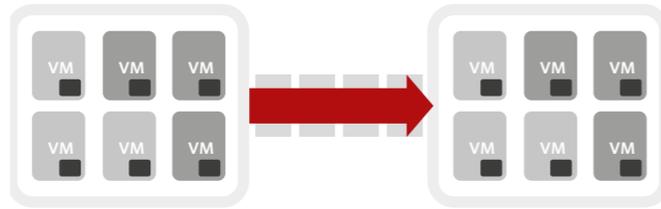

Figure 3. Host-based replication requires an agent on each VM, greatly increasing complexity.

Guest/OS-based Replication

In a 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.

- **Growth and change** – The requirement to install a module on every single 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 the implementation, putting increased burden on the IT team with increased management complexity.
- **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 for an environment with very few VMs, as the environment grows to over 20 VMs management and maintenance, including guest reboots, become a much greater negative business impact. Maintenance and updates for the DR strategy are now a weekend task, often requiring downtime.

Snapshots

Many solutions use snapshots as a method to enable a quick restore. A snapshot is a way to "freeze" a live storage system or VM at a moment in time. Changes continue to be made to the files beyond the snapshot capture. If changes are made beyond the snapshot capture and the VM or storage system encounters an issue, there is a choice to reject those changes by reverting the VM or storage system back to the time of the snapshot creation. A snapshot is especially useful when making changes to a single VM that where a rollback may be necessary.

There are two types of snapshots: **storage snapshots**, based on the hardware, and **hypervisor snapshots**.

Storage snapshots: expensive

Storage snapshots are taken on the entire storage volume as a whole and can expand exponentially in size, using a lot of first tier storage space. 30% of the net disk space is not unusual. This is particularly true when there are a lot of changes to the data on the storage after the time the snapshot is first taken. Most storage snapshot technologies also rely on the original disk.

Virtual Machine snapshots: incomplete

VM snapshots apply only to the specific individual VM and do not create copies of VMs. It's just a file that enables a virtual machine that already exists to be returned to a previous state (likewise for most storage-based snapshot technologies as well but with regards to the entire storage volume rather than only a single virtual machine). They are not protected in the case of hardware failure. If the files containing a virtual machine are lost, the associated snapshot files are rendered useless.

Backup   Snapshot   Snapshot   Snapshot

4H

16H+

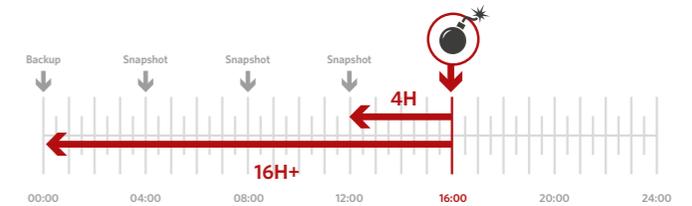00:00   04:00   08:00   12:00   16:00   20:00   24:00

Figure 4. Though snapshots shorten RPO, they usually are taken every 4 or 8 hours (because otherwise they will have too much impact on the performance and us to much disk space). This results in a better RPO than a traditional back-up, but still up to 4 or 8 hours of work can be lost when an incident occurs.

The question is, are Snapshots adequate as a DR solution?

- **No real DR** – Snapshots are used to save a point in time temporarily, not for a long-term solution. To create a copy of a VM to store, a backup or a DR site is needed, not a snapshot.
- **Performance** – Virtual Machine snapshots have an enormous impact on the performance of a virtual machine, and can also impact the entire environment with additional hypervisor and storage overhead.
- **Management** – Large numbers of snapshots are difficult to manage.
- **Frequency** – Because snapshots are typically taken every 4 hours (more would have too much impact on performance and storage), still 4 hours of data is lost after roll back (see figure 4). The claim that an RPO of 15 minutes is feasible, doesn't scale beyond very small environments. In modern environments a more continuous replication solution is needed, without performance impact on the production environment.
- **Snapshots & the cloud** – Though some DR solutions use snapshots and store them in a cloud environment (DRaaS), the snapshot still has to be created in the production environment before being replicated to the cloud. In this way the storage and performance impact remains the same. What is also important to know is which type of snapshot is being used by these solutions: is it a storage snapshot or a VM snapshot. Storage-based snapshot technologies require identical hardware, limiting cloud providers and hardware life cycles between the two organizations involved.

Hypervisor-based replication

All these categories of replication technologies have critical limitations in a virtual context. In this way they undermine the promise of virtualization and limit its functionality. To fully benefit from the investment made in virtualization without compromising on BC/DR, a new approach is required: hypervisor-based replication. Zerto moved replication up the stack, from the storage layer, above the resources abstraction layer into the virtualization/hypervisor layer. **Section 2** describes how Zerto's innovative hypervisor-based replication solution delivers enterprise-class, virtual replication and BC/DR capabilities for the data center and the cloud.

DR offered by hypervisors

Hypervisor vendors, like VMware, also 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. Even combined with VMware Site Recovery manager (SRM), the recovery time and scalability may not be enough to satisfy the business needs. While SRM adds capabilities around the planning, testing, and execution of a disaster recovery, it can't overcome the replication limitations of vSphere Replication, as VMware vSphere Replication utilizes virtual machine snapshot technology.

Disaster Recovery & the Cloud

With the cloud becoming more of an option, enterprises of all sizes are looking for the cloud, be it public, hybrid or private, to become part of their BC/DR solution. Virtualization has created the opportunity, but depending on the solution, there still can be a significant technology gap. Mission-critical applications can be effectively virtualized and managed; however, they cannot be effectively protected in a cloud environment when the wrong tools are chosen.

Disaster Recovery as a Service (DRaaS)

The cloud as a DR solution is a smart choice, since the cloud is more flexible and usually less expensive than implementing a self-owned DR site. When it comes to choosing a cloud service provider and a DRaaS service it is important to realize that DRaaS is not a technology. It is a service based on one of the technologies mentioned above. The only big difference lies in the place where DR files are stored. This means that, if a solution is based on snapshots, it will come with all the shortcomings associated with snapshots including performance and storage impact on the production site. And foremost: a 15 minutes RPO based on snapshots is still unrealistic.

Looking for a DRaaS service, a closer look has to be taken at the technology the service is based upon, with assurance that the RTO and RPO it offers is realistic and proven and without extra investments upfront. To help with this choice we set up a DR and DRaaS requirements checklist that can be 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 is the impact on the production site, due to the technology being used (e.g. snapshots)?
2. Which RTO and RPO does the solution offer? Is it 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 application-aware; does it offer application-consistent groupings of VMs?
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, when VMs are moved to other storage locations or when you want to do a migration?
10. **DRaaS** – Does it support multiple sites and is it multi-tenant? Does it offer securely isolated data streams for business critical applications and compliance?

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

Compliance
15. 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?
16. **DRaaS** – Are there any license issues or other investments upfront?
17. **DRaaS** – Where is the data being kept? Does the service provider comply with EU regulations?

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

The Zerto Revolution
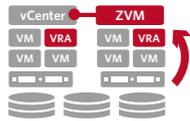
DR on a new level

Zerto has moved replication up the stack, from the storage layer into the virtualization/hypervisor layer. This results in an innovative hypervisor-based, virtualization-aware replication solution providing enterprise-class replication and BC/DR capabilities for the data center and the cloud.

Always-on replication

Zerto Virtual Replication continuously replicates I/O as it is created, delivering RPOs of seconds.

Scalable & Granular

With a software-based solution, scaling the infrastructure to support disaster recovery processes is simple and easy. As a new virtual host is added, simply install a new virtual appliance. Although Zerto Virtual Replication scales to support very large environments, it provides the same granularity that is needed in environments of all sizes, with the ability to recover files, VMs, applications, and entire sites.

Easy to manage

Zerto Virtual Replication installs seamlessly into the existing infrastructure with no configuration changes required in the hypervisor, application, or storage. The console is accessible anywhere and provides a complete view of the environment, making it easy to determine issues, leading to reduce troubleshooting and quicker resolutions. ZVR also provides a powerful dashboard with a consistent look and feel across all platforms.

Application consistent

Since applications often consist of more than one VM, Zerto developed Virtual Protection Groups (VPGs), a unique feature that enables replication of multiple VMs together as a group. Replicating VPGs gives the ability to restore the entire group from a single point in time, consistently, and with write-order fidelity.

Technology Agnostic

Zerto is hypervisor and hardware agnostic, removing barriers to innovation while greatly increasing efficiencies. There is no vendor lock-in so older or less expensive arrays and hypervisors can be used to further reduce costs. Additionally, it is now much easier to leverage or test new and innovative technologies, such as flash arrays.

Simple recovery

In case of a disruption, a simple failover and recovery process can be initiated from the same console. It is easy to test, configure, and execute disaster recovery automation. Testing can be done without any impact to the production site or replication, with test result reports for compliance.

Complete

Zerto's functionality extends to a complete solution for virtual environments, supporting private, hybrid, and public clouds with disaster recovery, long-term retention, testing, and migration capabilities.

SECTION 3

Zerto Virtual Replication

When production is virtualized, there is an obvious gap in the data protection strategy as it is usually based on older technologies that have reliance on physical asset limitations. Zerto Virtual Replication aligns production and disaster recovery strategies with a hypervisor-based replication solution.

In this section we will explain how the Zerto technology works and the advantages it offers.

Architecture

The heart of Zerto's replication technology is formed by two components:
- **Zerto Virtual Manager (ZVM)** – Zerto Virtual Manager manages disaster recovery, business continuity and offsite backup functionality at the site level; plugs into VMware vCenter and/or Microsoft System Center Virtual Machine Manager, and also includes a browser-based option.
- **Virtual Replication Appliance (VRA)** – Replicates the VMs and associated virtual disks; one VRA is installed per ESXi/Hyper-V host.

How replication works?
The Zerto Virtual Replication Appliances (VRA) copies I/O as it is created before it leaves the hypervisor. This continuous block-level replication delivers RPOs of seconds, minimizing data loss in the event of an outage.
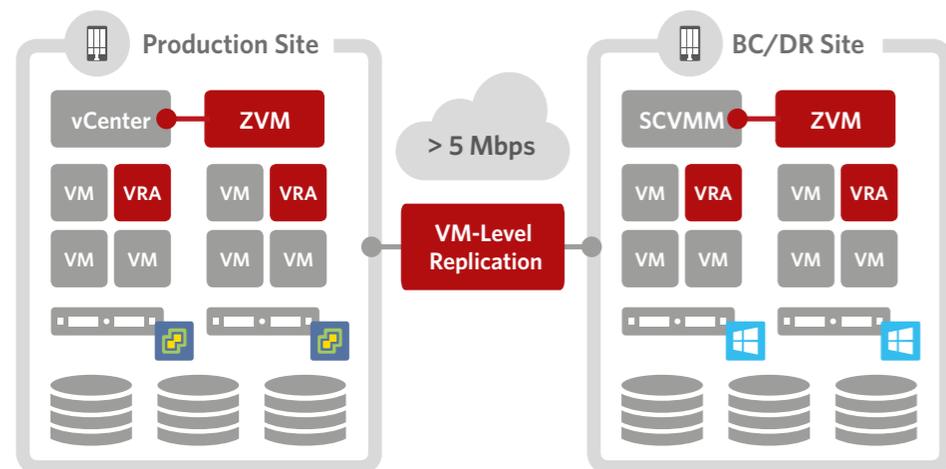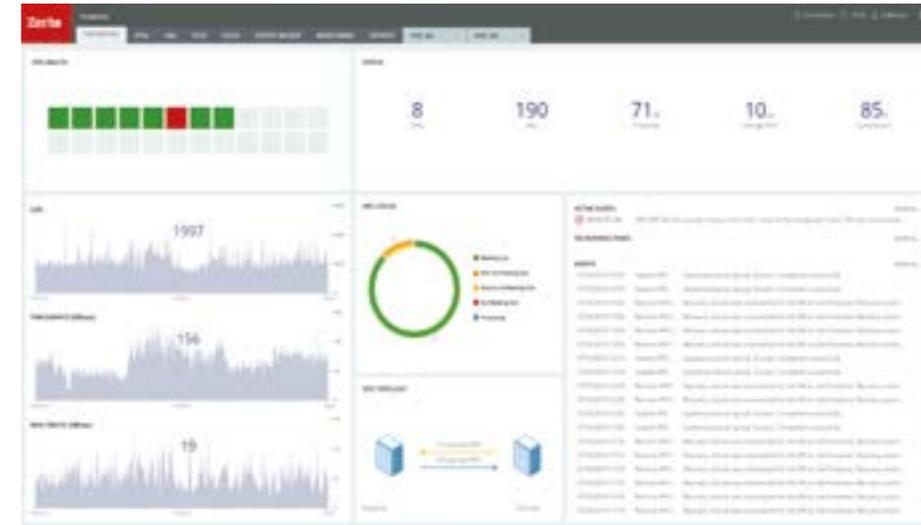
Figure 5. *Zerto's Architecture*

Figure 6. *Zerto Virtual Manager interface*

Features and Benefits
- **Journaling capabilities** – Provides continuous block-level replication with zero impact on application performance delivering point in time recovery with between 1 hour and up to 14 days worth of recovery points.
- **Hardware and hypervisor agnostic** - Remove barriers to innovation with a replication solution that has no dependencies on hardware or hypervisors.
- **Simple and seamless installation** – Installs seamlessly into the existing infrastructure with no downtime or configuration changes required.
- **Protect Production Workloads** - Ensure application consistency with groups of VMs which are protected, managed, replicated and recovered as one entity.
- **Scalable** – As a software-based solution it grows with the infrastructure, no matter how fast the business expands.
- **Simple, centralized management** – Centralized management for two sites with the Zerto Virtual Manager and for multiple sites with the Zerto Cloud Manager.
- **Aggressive Service Levels** – Achieves a Recovery Point Objective (RPO) of seconds and a Recovery Time Objective (RTO) of minutes.
- **Complete Orchestration** - Automate failover, failback, reverse protection is executed in just a few clicks.
- **Non-disruptive DR Testing** - Test the full recovery process without impacting production environments or ongoing replication, giving the team confidence they are covered in the event of a disaster.
- **Enterprise-class support** – Zerto delivers enterprise-class support services that are built into all of its products. These services include real time alerts when RPO/ RTO targets are not being met, network degradation alarms and reminders to check configurations and Virtual Protection Groups. Zerto solutions are backed by global support service centers that provide on-demand access to an expert team of support engineers.
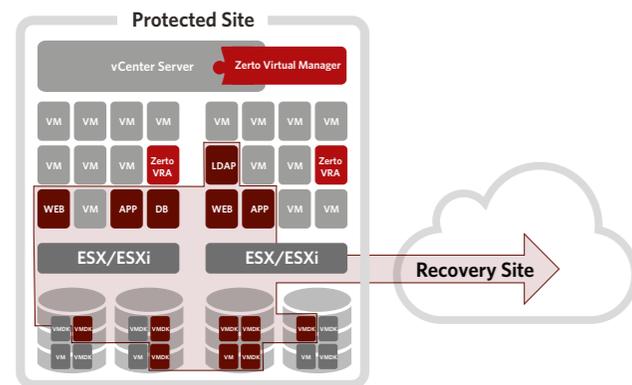
Management
The Zerto Virtual Manager (ZVM) plugs in at the virtual management console and gives a graphical overview of the site's VMs and their performance. If any problem occurs, it is represented visually, and alerts are sent as well. In the tabs at the top, all other functionality is available for orchestration and automation of failback and recovery processes, like boot order, re-IP, scripts, test and validation options.

Application-centric protection: Virtual Protection Groups

Many enterprise applications consist of more than one virtual server – a web server, application server, database server, etc – which are interdependent. When recovery is needed, all servers must be recovered from a single consistent point-in-time. To be able to do that, Zerto developed Virtual Protection Groups (VPG), which ensure consistency across a group of VMs. In this way the Zerto solution ensures that enterprise applications are replicated and recovered with consistency, regardless of the underlying infrastructure. Zerto Virtual Replication recognizes and preserves these relationships while enabling critical VMware features such as DRS, vMotion and Storage vMotion.

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

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

Fully automated and orchestrated

Replicating the data to the recovery site is only half the issue. The information that is there to protect a business in the event of a disaster needs to be easy to use. Zerto recognized this issue and built in automated and orchestrated processes that can be executed in just a few clicks when IT is in the middle of a high-pressure situation.

Fully configured failover process

Part of the VPG configuration is to set up the failover process. As part of this configuration boot order, re-ip on failover, length of journal, and other parameters are configured. With all this up front work done, this greatly simplifies the recovery process, reducing it to just a few clicks.

Failover as a business decision

Since every disaster is different, Zerto believes that failover needs to be a business decision and not an automated process. Because it is possible to pick a moment in time, this decision phase is essential for a correct failover. After clicking the failover button, an automated and orchestrated process will be started to bring services back online. In this way a failover can be done with the ability to choose a point in time, for example the point in time just before a database corruption occurred.
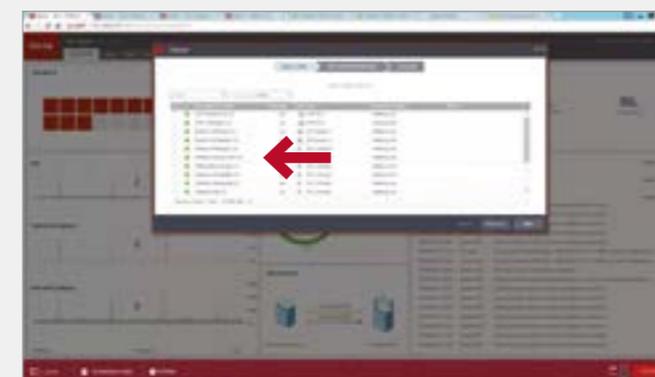
QUICK 4 STEPS FAILOVER PROCESS

The failover process consists of 4 simple steps. After an incident is visible in the management console:
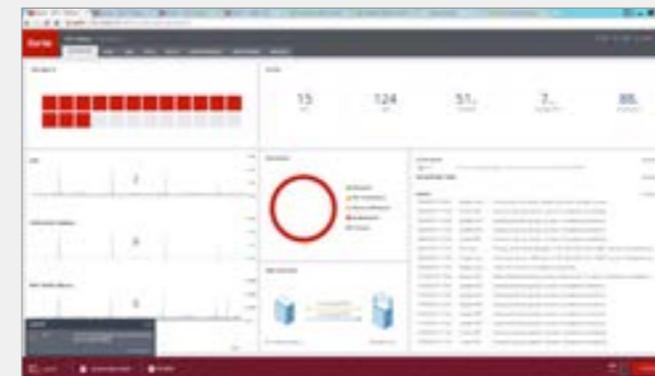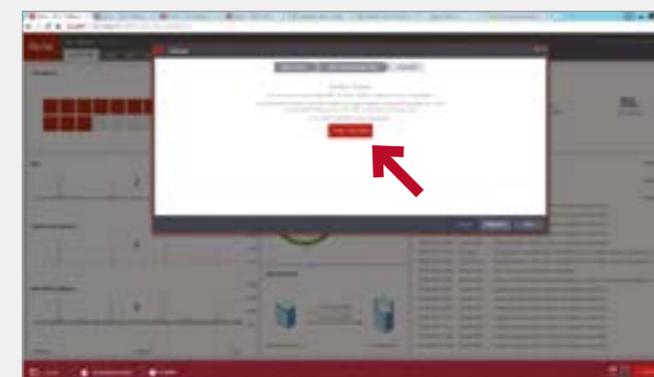
1. **Click failover.**
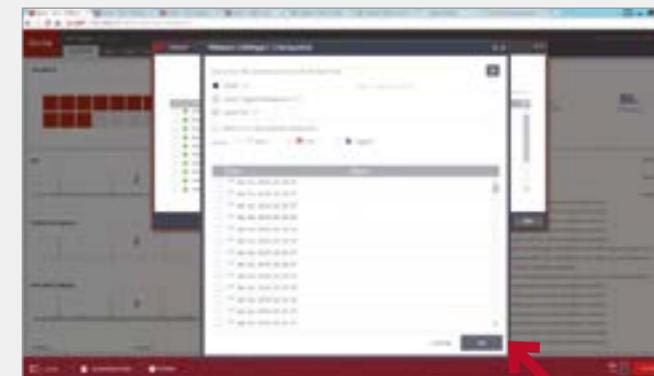
3. **Verify the point in time** to which the apps need to be restored. To avoid corrupted applications to be restored, it is necessary to go back to the point when they were not corrupted.

2. **Select the applications** (Virtual Protection Groups) that need recovery from the list.

4. **Start failover process.** The failover process begins and virtual machines are booted and reconfigured as needed.
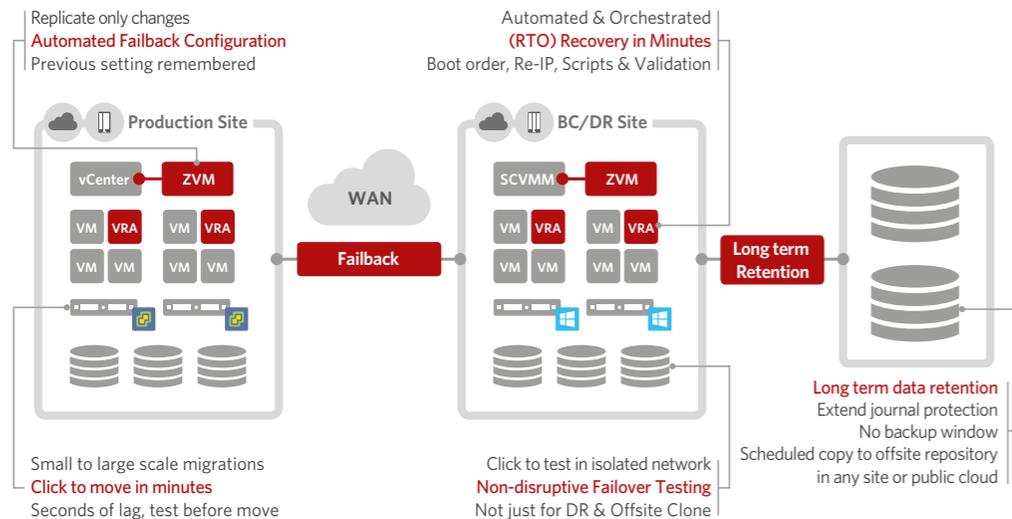
Figure 8. *When it comes to failback, the Zerto Architecture delivers full automation and orchestration, and non-disruptive failover testing. The same functionality can be used for sandbox testing and data migration as well. For complete protection options, the data at the DR site can also be used to make a long-term retention copy of the data, without impacting production.*

Automated Failover and Failback

As stated, 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 is executed in just a few clicks. Even when the disaster recovery process is initiated, there is the opportunity to rollback the failover should there be issues at the recovery site unrelated to Zerto, like a network is down. Upon a successful failover, reverse protection makes the failback process even easier. Reverse protection begins syncing the additional work that was done at the recovery site to the production site, when the production site is ready for use. After the applications have been updated to the original production site, failback is again just a few clicks. Many organizations will not failover because failback is so cumbersome; with Zerto Virtual Replication, everything is easy.

Non-disruptive disaster recovery testing

Organizations need to be able to demonstrate that disaster recovery processes will work in the event of the disaster to support internal and external compliance requirements. Zerto Virtual Replication enables non-disruptive testing in a sandbox environment, fully demonstrating the success of a failover. During the test the environment is still protected and replication is still in process. This means that DR testing and personnel resource scheduling no longer requires weekend test windows, as none the production environment needs to be taken down to fully exercise the test.

Figure 9. *Non-disruptive disaster recovery testing results in audit reports that can be used for compliance*

Sandbox testing

With the failover testing functionality Zerto can also create a test- and development environment.

Data migration

Data center migrations and consolidations are massive time and resource consuming projects that must be carefully scheduled and planned to try to minimize downtime and loss of productivity. With Zerto's hypervisor-based replication technology however, migrations can be a near painless activity. Using the core attributes of ZVR, virtualized applications can be tested ahead of time, and migrated in just a few minutes with minimal downtime.

- **Simplicity** – Migrating VMs is as simple as pointing the replication to the target datastore of choice and allowing the data to be replicated to the new site in the background from other business activities.
- **Granularity** – Migrations can be very granular with the ability to migrate at the VM Disk (VMDK) level, which can be pointed to different tiers of storage.
- **Flexible** – Support for a heterogeneous environment allows for migrations between different types of hardware and different VMware and Hyper-V versions, from a vCenter environment to a vCloud environment, and between different versions of ZVR
- **Fully automated moves** – Leveraging the VPG configuration, moving VMs to a new location is accomplished simply and in just a few clicks. This dramatically reduces the application downtime to just a few minutes, ensuring revenue generating activities are not impacted.

Long-term retention data copies

Since the data is replicated to the DR site, it is easy create an offsite copy of the data for long-term retention or for compliance. This process and its infrastructure are not a part of the production site, removing the overhead and management burden, and can be managed from the same ZVR user interface.

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. ZVR has solved this most frequent disaster problem by providing the ability to recover a single file or folder from up to 14 days in the past, using the Journal. The continuous block-level replication 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 extremely minimized.

- **Risk** – Minimizes data loss across files, folders, VMs, applications, and sites with the ability to recover at any level, at any point in time.
- **Simplicity** – Reduces mean time to recovery with the ability to leverage a automated workflow to recover files, applications and data.
- **Protect productivity** - When a file or folder is accidentally deleted, end-users no longer need to recreate hours or a day of lost work, preserving productivity and employee morale.

Zerto Virtual Replication Use Cases

The advantages of Zerto Virtual Replication are available for many different use cases. The basic technology is storage agnostic and supports mixed hypervisors, which means that any site can be replicated to any other site, whether it is a private cloud, a public cloud, a service provider or a branch office.

Private Cloud to Private Cloud
The most traditional example is setting up a DR site as a remote version of the internal datacenter. The less traditional aspect is, that any type of storage can be used, any storage vendor and a mix of hypervisors. Zerto supports it all with no limits on distance.

Satellite Offices Protection and Migration
Another application can be that Zerto is used to protect or migrate applications between branch offices, again using any type of storage, any storage vendor and a mix of hypervisors.

Hybrid cloud, Disaster Recovery as a Service (DRaaS)
Various Cloud Service Providers give the option to use the cloud as a DR site. If this Disaster Recovery as a Service is based on Zerto, it will offer all Zerto's advantages. Within these services it is possible to control DR from a self-service portal, or have it remotely managed by the service provider (Managed Disaster Recovery as a Service), or even use a private cloud as DR site for a cloud-based production site.
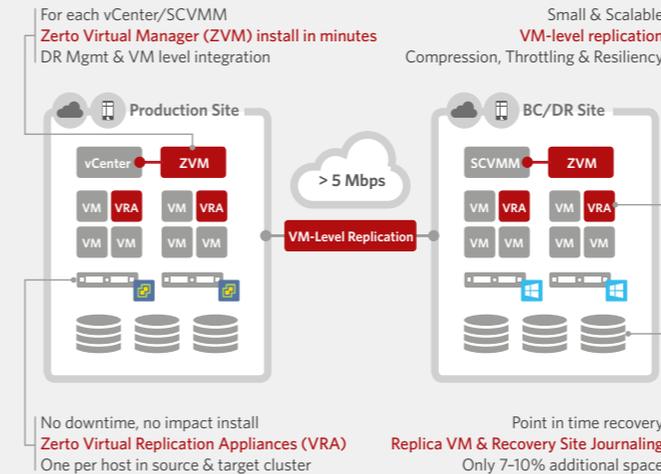

Self-Service Disaster Recovery as a Service
It is also possible to use a public cloud service, like Amazon Web Services, as a DR site. Here businesses have to configure the service themselves, as if they were setting up a remote DR site.

Zerto Summary

ZERTO ARCHITECTURE

For each vCenter/SCVMM
Zerto Virtual Manager (ZVM) install in minutes
DR Mgmt & VM level integration

Small & Scalable
VM-level replication
Compression, Throttling & Resiliency

Production Site

> 5 Mbps

BC/DR Site

VM-Level Replication

No downtime, no impact install
Zerto Virtual Replication Appliances (VRA)
One per host in source & target cluster

Point in time recovery
Replica VM & Recovery Site Journaling
Only 7–10% additional space

ZERTO FEATURES

Hypervisor-based replication – Continuous block level replication across VMware and Microsoft Hypervisors across multiple sites.

Granular, VM-level protection – Protect exactly what is needed - VMs and VMDKs/VHDs.

Full orchestration – Automate failover, failback, reverse protection and disaster recovery testing.

Continuous data protection – Recover from any disruption with 14 days worth of checkpoints available every couple of seconds.

Complete data protection – One solution for BC/DR across storage & hypervisors with long-term retention.

Simple management – Centralized management for two sites with the Zerto Virtual Manager or multiple sites with the Zerto Cloud Manager.

Aggressive service levels – Realize recovery point objectives of seconds and recovery time objectives of minutes.

Simple installation – Complete in one hour without changing application or storage configurations – architect for performance not DR.

Protect production workloads – Ensure application consistency with groups of VMs which are protected, managed, replicated and recovered as one entity.

Offsite retention – Increase the usefulness of replicated data with long-term retention capabilities at the replication site.

Storage agnostic replication – Remove barriers to entry with storage independent replication.

Support for public cloud – Replicate, protect and migrate workloads to public cloud services like Amazon Web Services.

WANT TO TRY IT OUT?
Zerto Virtual Replication can be installed, configured and replicating VMs in under 1 hour. With simple VM-based replication enabling RPOs of seconds and RTOs of minutes. Why not go to www.zerto.com/trial and click to download a free trial today?

www.zerto.com/trial

About Meriplex

A managed solutions provider specializing in intelligent network solutions, cloud enablement and managed services for the mid-enterprise market, Meriplex develops lasting customer relationships by demonstrating positive attitudes in providing creative, reliable, and personalized solutions.

Meriplex

WANT TO TRY IT OUT?
Zerto Virtual Replication can be installed, configured and replicating VMs in under 1 hour. With simple VM-based replication enabling RPOs of seconds and RTOs of minutes. Why not go to www.zerto.com/trial and click to download a free trial today?

www.zerto.com/trial

meriplex.com