

June 2011

## **Hypervisor-based Replication**

A New Approach to  
Business Continuity/  
Disaster Recovery



### Zerto has introduced a virtual-aware, software-only, tier-one, enterprise-class replication solution purpose-built for virtual environments.

Virtualization of the data center has proven to be a true IT game-changer, providing increased flexibility and control in managing mission-critical data assets, as well as essential cost savings. Organizations large and small continue to move forward with virtualization initiatives, either going completely virtual or operating – in varying ratios – in a hybrid physical/virtual environment.

To more fully realize the benefits of virtualization and get the most out of their investments in the technology, organizations need to optimize all IT processes and activities for their virtual environment. This includes security, compliance, and Business Continuity/Disaster Recovery (BC/DR). Of the three, BC/DR is perhaps the most difficult because to date there have been no virtual-ready remote replication methods that organizations could adopt. That is all about to change.

Until now the most common replication technologies and methods essential to mission-critical BC/DR initiatives have been tied to the physical environment. Although they do work in the virtual environment, they aren't optimized for it, leading to significant operational and organizational challenges. Among them:

- If a replication solution isn't virtual-ready, management overhead could be more than doubled; many of the benefits achieved through virtualization, therefore, could be lost in the disaster recovery sphere.
- Virtualization is scalable, but today's BC/DR methods are not. Customer data is always growing, so a company can find its information inventory expanding exponentially and not have a replication solution to keep pace.
- In an increasingly heterogeneous IT environment, some replication methods remain firmly tied to a single vendor and hardware platform, limiting the organization's ability to get the best solutions – and service – at the best price.

Clearly, there is a critical need for BC/DR to become aligned with the promise and reality of virtualization in the data center. Competitive and compliance-related pressures are at an all-time high, and organizations need every advantage to ensure excellence in their BC/DR capabilities. With the introduction of **hypervisor-based replication**, Zerto has finally brought BC/DR up the stack where it belongs: in the virtualization layer.

## I. Background

Not only are today's array-based, guest/OS-based, and appliance-based replication technologies not optimized for the virtual environment, most were developed specifically for use with physical IT assets. As such, all three have issues that inhibit the efficiency and effectiveness organizations require. A brief review of the structures and limitations of these methods will help to amplify the advantages and benefits of Zerto's virtual-ready, hypervisor-based replication solution.

### Legacy Solutions

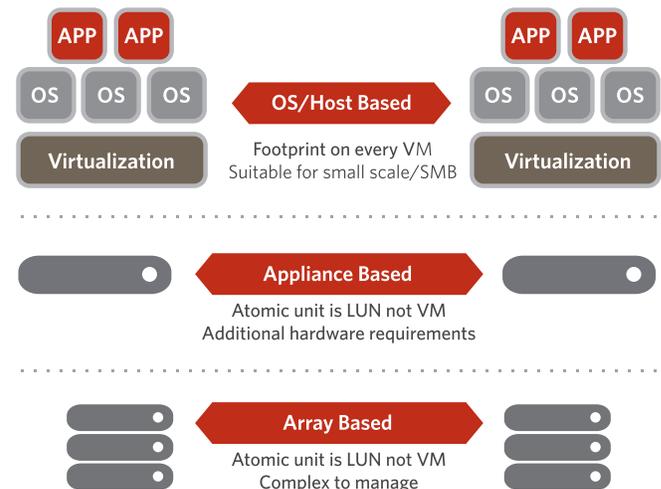


Figure 1 - Legacy Solutions

### **Array-based Replication: Insufficient Granularity**

Array-based replication products are provided by the storage vendors and deployed as modules inside the storage array. Examples include EMC SRDF and NetApp SnapMirror. As such, they are single-vendor solutions, compatible only with the specific storage solution already in use. Currently the most popular replication method in use in organizations, array-based replication does not have the granularity that is needed in a virtual environment.

For example, mapping between virtual disks and array volumes is complex and constantly changing, creating management challenges and additional storage overhead. Often, multiple virtual machines reside on a single array volume, or logical unit. An array-based solution will replicate the entire volume even if only one virtual machine in the volume needs to be replicated. This underutilizes the storage and results in what is known as “storage sprawl.”

Because array-based replication lacks the visibility and granularity to identify specific virtual machines in different locations, organizations tend to put all disks from an enterprise application into a single storage logical unit, when in fact there are operational advantages to splitting them up over a number of logical units.

Array-based replication has other important disadvantages as well, such as:

- It is designed to replicate physical entities rather than virtual entities. As a result, it doesn’t “see” the virtual machines and is oblivious to configuration changes – and due to their dynamic nature, virtual environments have a high rate of change.
- It requires multiple points of control: in addition to the physical storage array’s management console, IT also is managing virtual assets from a virtualization management console, such as VMware’s VSphere Client.
- Though optimized to work with an organization’s existing storage array, it locks in the organization to a single vendor.
- Array-based replication is inflexible, negating the advantages of virtualization.

### **Guest/OS-based Replication: Impossible to Scale**

Guest/OS-based replication solutions comprise software components that have to be installed on each individual physical and virtual server. Such solutions include Double-Take from Vision Solutions and Veritas Volume Replicator from Symantec. Although much more portable than array-based solutions and typically simpler to manage (because of their small scale), guest/OS-based replication solutions are not fit for enterprises for the following reasons:

- 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.
- Since it is designed to replicate only single servers, a guest/OS-based solution is incapable of protecting full multi-server enterprise applications. For example, an organization that wants to protect multiple virtual machines can install a module on each virtual machine; however, each virtual machine is still protected independently. Users cannot get a consistent point-in-time view across multiple virtual machines, so an application that spans multiple virtual machines will not be protected correctly or consistently.

### **Appliance-based Replication: More Moving Parts**

Appliance-based replication solutions – the primary current exponent of which is EMC RecoverPoint – are similar to array-based solutions in that they are hardware-based and specific to a single platform. The main difference is that the replication code runs on an external, physical appliance instead of inside the storage arrays themselves. This gives it an advantage over array-based solutions in that it is more flexible and does not consume array resources.

When it comes to protecting virtual environments, however, appliance-based solutions suffer from the same disadvantages as the array-based products. Specifically:

- They replicate physical entities rather than virtual entities. Their focus, therefore, is the logical unit rather than the virtual machine. This lack of granularity conflicts with the requirements and promise of virtualization.
- They require dual points of management: the physical management console and the virtualization management console.

Ultimately, the key disadvantage of appliance-based replication is that in introducing more moving parts, it also introduces more management complexity. And with an additional hardware footprint it inevitably leads to additional hardware bottlenecks. As with array-based and guest/OS-based replication, appliance-based replication is not virtual-ready and so is not ideal for virtual or hybrid environments.

## II. Introducing Hypervisor-based Replication

We have seen how three different categories of replication technologies designed for physical IT environments have critical limitations in a virtual context. This does not imply that there is anything inherently wrong with any of these technologies. But to fully benefit from virtualization without compromising on BC/DR, a new approach is required.

Before virtualization, replication was managed at the storage layer, which made perfect sense because that's where the information was. If there is a physical box you want to keep track of, you could track it with physical sensors. But in a virtual environment, the boxes aren't (or aren't all) physical, so putting a physical sensor on a virtual box isn't going to help you protect its contents.

The problem is actually common in historical terms: one technology often advances at a faster rate than others, creating a capability gap. Virtualization offers extraordinary capabilities and benefits, but they cannot be fully enjoyed unless and until other technologies within the data center evolve to enable them. Managing a virtual or hybrid environment from the physical storage layer makes it harder to fully realize the benefits of virtualization and inhibits the move to cloud computing.

That's why Zerto realized the need to move replication up the stack – above the resources abstraction layer – into the virtualization/hypervisor layer. And that's how hypervisor-based replication was born.

### Architecture

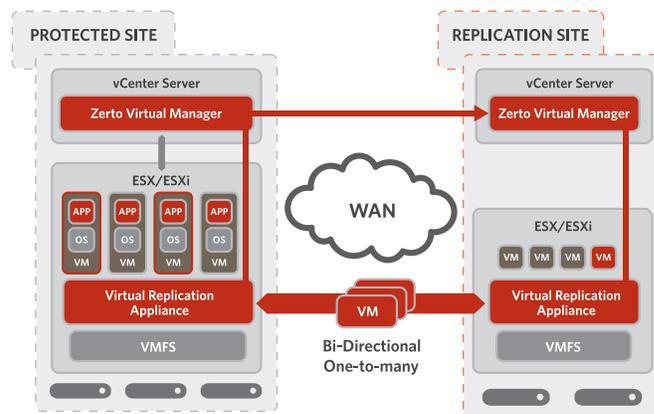


Figure 2 - Architecture

### Zerto Architecture: Simple, Effective, and Virtual-ready

Zerto has introduced a virtual-aware, software-only, tier one, enterprise-class replication solution purpose-built for virtual environments. The company's innovative hypervisor-based replication solution is currently the first and only solution that delivers enterprise-class, virtual replication and BC/DR capabilities for the data center and the cloud.

At the heart of this patent-pending replication technology are two key components:

- **Zerto Virtual Manager (ZVM)** – The ZVM plugs directly into the virtual management console (such as VMware's vCenter), enabling visibility into the entire infrastructure. The ZVM is the nerve center of the solution, managing replication for the entire VSphere domain, keeping track of applications and information in motion in real time.
- **Virtual Replication Appliance (VRA)** – The VRA is a software module that is automatically deployed on the physical hosts. The VRA continuously replicates data from user-selected virtual machines, compressing and sending that data to the remote site over WAN links.

Because it is installed directly inside the virtual infrastructure (as opposed to on individual machines), the VRA is able to tap into a virtual machine's IO stream. Therefore, each time the virtual machine writes to its virtual disks, the write command is captured, cloned, and sent to the recovery site. This is more efficient, accurate, and responsive than prior methods.

Unlike some replication technologies that primarily offer data protection along with cumbersome snapshots and backup paradigms, Zerto's solution provides continuous replication *with zero impact on application performance*.

Hypervisor-based replication is fully agnostic to storage source and destination, natively supporting all storage platforms and the full breadth of capabilities made possible by virtualization, including high availability, clustering, and the ability to locate and replicate volumes in motion.

### Application Protection: Virtual Protection Group

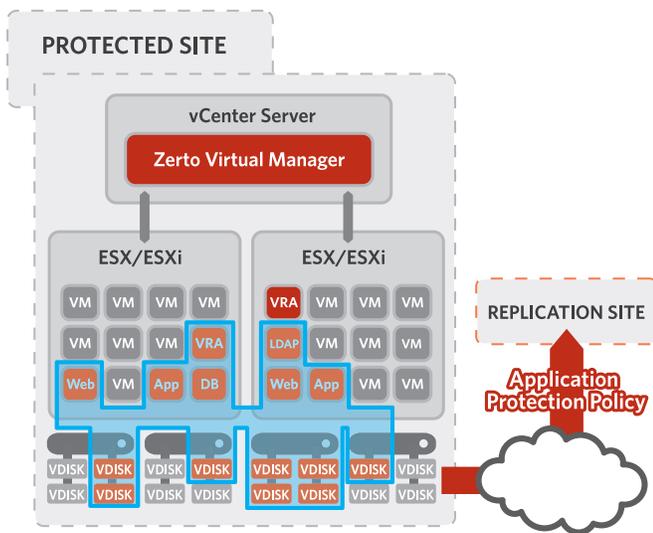


Figure 3 - Application Protection: Virtual Protection Group

### Application-centric Protection: An(other) Important Differentiator

Enterprise applications often consist of more than one server. These servers are interdependent so when they are in need of recovery, they must be recovered from a single consistent point-in-time image. Before Zerto, this has been impossible to achieve.

Leveraging the concept of storage-based consistency groups that operate strictly at the storage array's logical unit level, and the unique capabilities of the virtual platform, Zerto has developed an important innovation: Virtual Protection Groups (VPGs), a user-assigned group of virtual machines and their related virtual disks that have dependencies and must be recovered from a consistent image.

Zerto VPGs ensure that enterprise applications are replicated and recovered with consistency, regardless of the underlying infrastructure. This enables organizations to deploy the application across different physical devices to maximize performance, capacity, and/or to reduce the complexity of the infrastructure. The ZVM recognizes VPGs and the component relationships within each one, assuring complete, accurate, and timely replication. While the core idea is readily achievable at the physical storage level, only Zerto has figured out how to do this at the virtual level for both server and storage locations.

### III. Hypervisor-based Replication Benefits

The benefits of virtualization are well known and include increased efficiency, flexibility, and savings in space, equipment, and energy costs. Hypervisor-based replication enables many of those benefits to be fully realized. As a replication technology in and of itself, hypervisor-based replication offers many important benefits to organizations as well.

#### Granularity

The ability to replicate at the correct level of any virtual entity, be it a single virtual machine, a group of virtual machines, or a virtual application (such as VMware VApp), is critical. Zerto's solution can replicate all virtual machines consistently and all of the meta data as well - including VApp meta data. This ensures that the entire application infrastructure can be recovered in the event of a disaster.

A typical enterprise application includes a web server, application server, database server, etc., and all have their respective disks. Today, administrators tend to put all those disks in a single logical unit in storage so they can replicate the entire application at once without having to search for its individual components. The problem is that this means the entire logical unit must be replicated, even though it may contain other applications that are not in need of replication. That lack of granularity - where administrators cannot identify specific applications and application components to replicate - is inefficient.

For example, a given CRM application may span across eight virtual machines deployed on four physical servers, using five different data stores located on three different logical units. With hypervisor-based replication - and



## Hypervisor-based Replication

A New Approach to Business Continuity/Disaster Recovery

only with hypervisor-based replication – centralized management through the hypervisor layer enables the solution to find what it's looking for no matter where it is.

That is simply impossible with prior, non-virtual- ready replication technologies. The goal is full consistency between and among all application components. With hypervisor-based replication, that goal is achieved.

### Scalability

There are two aspects of scalability: deployment and management. As a virtual infrastructure grows, an organization's DR capabilities must grow with it seamlessly, without having to purchase, install, and configure additional proprietary hardware. Zerto's hypervisor-based replication solution is software-based so it can be deployed and managed easily, no matter how fast the infrastructure expands. The solution also enables administrators to perform operations and configure policies at the level of the virtual machines or applications.

### Ease of Management

With no guest-host requirements or additional hardware footprint, Zerto's solution is easy to manage. It simply resides in the hypervisor, enabling centralized management from the virtual management console (such as VMware vCenter). Organizations can now manage everything from the same console. Because it is software-based, it is user-installable (the VRA install process itself is automated), user-configurable, and scalable.

### Server and Storage Motion

One of the great advantages of the virtual environment is the ability to quickly move virtual machines around from one physical server or array logical unit (data store) to another. This might be done for load balancing or other strategic data management reasons. With VMware, this is accomplished manually through VMotion or automatically using Distributed Resource Scheduler (DRS). Only hypervisor-based replication supports this capability, continuing to locate and replicate data no matter where it resides or is moved to.

### Hardware-agnostic

Hypervisor-based replication is hardware-agnostic, supporting all storage arrays, so organizations can replicate from anything to anything. In today's increasingly heterogeneous IT environment, this allows users to mix storage technologies such as Storage Area Network (SAN) and Network-Attached Storage (NAS), and virtual disk types such as Raw Device Mapping (RDM) and VMware File System (VMFS).

### Remarkably Effective

In Business Continuity/Disaster Recovery, the two key metrics are Recovery Point Objective (RPO) and Recovery Time Objective (RTO). The former is an indication of the amount of data at risk of being lost between data protection events and how long until all the data at risk is recovered. The latter refers to the amount of time it takes to recover from a data loss event and return to normal operation and availability. Zerto's hypervisor-based replication solution achieves RPO in seconds and RTO in minutes.

## IV. Conclusion

If you have a virtual or hybrid environment, you cannot realize the full benefits and promise of virtualization unless your replication solution is virtual-aware and virtual-ready. Zerto's hypervisor-based replication technology is the first and only solution that delivers tier one, enterprise-class, virtual replication and BC/DR capabilities for the data center and the cloud.

Your information will only grow over time; replication demands will also grow, as will the scope of your virtualization initiatives. Purpose-built for the virtual environment, hypervisor-based replication will position you for growth and optimize your business continuity and disaster recovery activities. ■



Contact us today to learn more or request a free trial at [www.zerto.com](http://www.zerto.com) or [info@zerto.com](mailto:info@zerto.com).

930 Winter Street, Suite 2500  
Waltham, MA 02451  
Phone: 781.478.6806  
Fax: 781.478.6601