



RepliStor and VMware — A Good Pair for Robust, Cost-Effective Computing

Analyst: Michael Fisch

Management Summary

When you think of items that pair well, which come to mind? Maybe coffee and cream, oil and vinegar, coat and tie, Antony and Cleopatra, Lennon and McCartney, wine and cheese... Just what makes these pairs *good*?

In each case, the items have different but related characteristics that complement one another. They might be food, clothing, musicians, but the key is that the items work together in a way that the whole is greater than the sum of its parts. Therefore, we say they are well matched.

EMC offers a couple of products that it considers a good solution-pair for *Windows* environments – *RepliStor* and *VMware*. *RepliStor* is server-based software for asynchronous replication. It creates a replica at a target site over virtually any distance, and offers a relatively inexpensive solution for disaster recovery and data migration.

VMware is the well-known purveyor of server virtualization technologies and also is an independently-operated subsidiary of EMC. VMware creates multiple self-contained virtual machines on Intel- and AMD-based servers, each of which can run a discrete operating system and application. This technology is a great facilitator for server consolidation, a more dynamic, efficient infrastructure, and ensuring business continuity.

Together, *RepliStor* and *VMware* address separate and important aspects of enterprise IT, but they also offer some synergistic use cases.

1. A lower-cost infrastructure for remote replication for disaster recovery;
2. Branch office data consolidation that backs up and protects their data centrally;
3. A lower-cost infrastructure for data distribution; and
4. Server migration, especially over remote distances.

Read on for details about how and why *RepliStor* and *VMware* are an effective pair for robust, cost-effective computing.

IN THIS ISSUE

➤ Drive for Cost-Effective and Robust Computing	2
➤ RepliStor	2
➤ VMware	3
➤ Synergistic Use Cases	3
➤ Conclusion	4

Drive for Cost-Effective and Robust Computing

The drive for more cost-effective and robust computing never ceases. Like the pursuit of progress in other fields, there is a general belief that computing can always become faster, cheaper, and more reliable. As predicted by Moore's Law in the 1960s, the price/performance¹ of microprocessors has doubled about every couple of years. A computer that filled a data center 40 years ago can fit in a wristwatch today. Advancements in networking and storage capacities have followed similar trajectories. Since computing continues to live up to its promise of delivering both *more* and *better*, people have come to expect it.

Replication

One such area of expectation is *high availability*, which simply means a computer system is consistently accessible and able to do work. Thus, it is reliable and does not easily break nor require frequent downtime for maintenance. So when it does break, it can recover quickly. High availability is essential, especially because computing is so interwoven into modern communications, transactions, and daily life. The pursuit of high availability is like running an asymptote toward infinity, where one keeps getting closer to the ideal (i.e., 100% availability) without fully arriving. For instance, the first milestone might be an uptime of 95%, then 98%, then 99%, then 99.9%, and so forth. Progress along this line is a matter of adopting increasingly robust technologies and capabilities.

Replication is an important technology for high availability, because it enables long-distance data protection and fast recovery. Replication maintains a full physical copy of data on disk at a remote site in real time. Like the spare tire in the trunk, if the source data becomes unavailable for some reason, a current or near-current copy is available to resume operations in seconds or minutes. Replication protects against computer system failures as well as local and regional disasters, such as fires, floods, or electricity outages. It also is useful for

consolidating data from branch offices for centralized backup, or for sending data in the reverse direction for distribution.

Server Virtualization

Another expectation is regular improvement in *price/performance* – beyond just faster processors. Enterprises spend a significant amount of money on IT. They have to, because so many aspects of modern business rely on computer systems. These systems evolve and grow with time and, therefore, require more resources (i.e., processing, storage, network bandwidth). IT is a moving target. Furthermore, the IT budget consists of multiple cost factors, including hardware and software acquisition, maintenance and support, professional services, management (i.e., skilled administrators), power and cooling, floor space, and disposition. They all add up, of course, but they also present numerous different areas and opportunities to take cost out of the system. For all these reasons, enterprises are looking for new ways to take costs out of computing and do more with fewer resources.

Virtualization of industry-standard servers (based on Intel and AMD architectures) and operating systems (i.e., *Windows* and *Linux*) have emerged as a popular and effective way to lower computing costs and improve flexibility and manageability. Virtualization provides an abstraction layer for a physical resource. It takes something hard and fixed, like granite, and turns it into something flexible and easy to work with, like clay. In the case of servers, it transforms a physical server into dynamic, shared resource that can run multiple applications and adjust to different workloads. Unfortunately, most servers still are dedicated to a single application and have poor utilization. Virtualization opens the door to a greater degree of consolidation in the server infrastructure, as well as the benefits of higher utilization, easier management, and lower overall costs.

RepliStor

EMC RepliStor is a server-based software solution for asynchronous file-level replication in Windows environments. It

¹ Or more precisely, transistor density.

replicates data from a source to a target server over virtually unlimited distance using TCP/IP connections (i.e., LAN or WAN). Salient features of RepliStor include the following.

- **Real-time replication** – Maintains a virtually current copy of data at a local or remote site with minimal impact on the application server. After the initial synchronization, it sends byte-level or whole-file changes, on either a continuous or a scheduled basis.
- **Consistent snapshots** – In addition to a replica, it can create application-consistent snapshots at the target for fast recovery to a previous point in time (such as in the case of data corruption). It leverages Microsoft's *VSS* capability, currently supporting *Exchange 2003*, *SQL Server 2005*, and Windows file system.
- **Automated or manual failover** – The system automatically detects a failure, fails over to the target, and then can be manually failed back after the primary server has recovered.
- **Write-order consistency** – Ensures data at the target is consistent with the source.
- **Standard hardware** – Runs on Industry-standard servers, as opposed to proprietary hardware.
- **One-to-many configurations** – Replicates from one server to many, or vice-versa, for purposes of data distribution and consolidation.

This combination of features makes RepliStor is a good solution for disaster recovery, backup consolidation, data movement, and data repurposing in Windows environments.

VMware

VMware's premier virtualization suite is *VMware Infrastructure 3*, which includes *ESX Server*². *ESX Server* installs directly on Intel/AMD-based server hardware³ and

² *VMware Server* (formerly *GSX Server*) is a free, entry version.

³ Also called "bare metal", it can be most any server with Intel or AMD processors.

creates partitions or virtual machines. Each virtual machine runs an operating system and creates an environment in which to run associated applications. The virtual machine is a self-contained, portable, and complete system consisting of processing, memory, network, and storage resources. In short, *ESX Server* turns a physical server into multiple virtual ones.

ESX Server supports up to 128 virtual machines per physical server. The servers themselves can have up to 32 logical CPUs and 64 GB RAM partitioned among the virtual machines. Windows, Linux, *Solaris*, or *Novell NetWare* operating systems are supported without modification. To handle changing application workloads, *ESX Server* can enforce policies that allocate CPU, memory, disk, and network resources among virtual machines.

VMware Infrastructure 3 includes a number of other products for management and optimization of the virtualized environment.

- *VirtualCenter* is a tool for centralized management, automation, and optimization, which eases the work for IT administrators.
- *VMotion* moves virtual machines running live (active) applications between physical servers without disruption. This is a great tool for minimizing downtime and optimizing the infrastructure. There are also solutions for centralized backup, application failover for high availability, dynamic application placement among servers, dynamic resource allocation among virtual machines, and tools for migrating physical to virtual environments.

Synergistic Use Cases

RepliStor can run in a VMware virtualized environment and replicate from physical to virtual, virtual to virtual, and virtual to physical. In addition to *ESX Server*, RepliStor also supports *VMware Server*. The combination of RepliStor and VMware opens up several new infrastructure possibilities for enterprises that deploy both.

Remote Replication for Disaster Recovery

VMware can bring flexibility and lower costs to disaster recovery. Because virtual machines truly are hardware independent, using VMware eases the movement between servers, regardless of type and vendor. In the past, restarting a replica traditionally required the same server hardware at the remote site, due to dependencies in the registry. However, using RepliStor with VMware allows a replica to reside and restart on any server hardware, such as a lower-cost model. Multiple replicas could also reside on a single physical server, which provides a form of consolidation for disaster recovery. This additional flexibility lowers the cost threshold for asynchronous replication and the robustness it brings to enterprise IT.

Branch Office Data Consolidation

The RepliStor/VMware combination is also a solution for the perennial problem of branch office backups. Branch offices traditionally lack the resources, focus, and skilled personnel to make sure their data is properly protected. However, data consolidation offers a way to alleviate this branch office burden as well as make sure their backups are reliable. Use RepliStor to replicate data from the branches to a server in the central data center – either continuously or on a scheduled basis – and create snapshots. With VMware virtual machines, a single server can be the target for multiple branches. Then, use the backup system in the main data center to back up and restore the branch data. The result is a more efficient overall backup process and better data protection for branch offices.

Data Distribution

Some enterprises need to distribute data to remote sites for convenient, local access. RepliStor can broadcast in a one-to-many configuration, sending data to sites in real time or periodically according to a schedule. With VMware, the target server at each site can be shared, as opposed to dedicated and probably underutilized.

Server Migration

RepliStor and VMware can facilitate

server migration and consolidation, especially over distance. First, RepliStor copies data to the target site. The *VMware Converter*⁴ transforms the physical system image to a virtual machine, so it can run in a virtualized environment. The original server is now free and available, either to be disposed of or repurposed.

Conclusion

RepliStor and VMware may not be as well-known a pair as coffee and cream, but they blend well (synergistically) in an enterprise Windows environment. RepliStor brings asynchronous replication and snapshots for disaster recovery and moving data over distance. VMware's virtualization layer makes the server infrastructure much more flexible and efficient. Together, they bring both cost-effectiveness and robustness to Windows computing. RepliStor and VMware are good solutions individually, but don't forget to consider them – as a pair.



⁴ *VMware P2V Assistant* is being renamed *VMware Converter*.

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over ten years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, and "clipper.com" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "Navigating Information Technology Horizons", and "teraproductivity" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.