Developing a Complete RTO/RPO Strategy for Your Virtualized Environment

NOVEMBER 2014

WHITE PAPER BY GEORGE CRUMP
Lead Analyst
Storage Switzerland, LLC

Virtualization, whether server, desktop or both, is certainly becoming mainstream in today's data center. The percentage of servers that are virtualized is increasing along with the critical nature of those virtualized systems. As a result, properly protecting and maintaining the right level of uptime for these systems is a key requirement that most virtualization administrators have to address.

What are RPO and RTO?

Answering this key question of how to provide protection and uptime requires an understanding of two concepts, recovery point objective (RPO) and recovery time objective (RTO).

RPO is the desired point in time that a failed server needs to be recovered back to. It is essentially an indicator of the number of transaction that can be lost for a given application.

For example, if a server fails at 4:00pm on Tuesday, and the last backup was at 11:00pm on Monday, the earliest possible recovery point is 11:00pm Monday. All the data that was created or modified from 11:00pm until 4:00pm will have to be recreated. For some virtual servers this amount of lost information may not be as bad as it may sound. If there were a limited number of changes made and if there was a paper trail from which that information could be rescued then bringing the impacted application back to a current state may be feasible.

For other applications it could be a catastrophic situation. This can be especially true if there is alot of data to be re-entered or if there is no audit trail, like paper that would allow this information to be recreated.

Recovery time is the actual time it takes to move this application back to a state that it can be accessed by its users. This essentially involves the transfer time from the backup or secondary storage target, across a network, back to the primary storage system. Obviously, the size of the data set to be transferred and the speed of the network directly impact how fast that data can be moved.

RPO also impacts RTO since, depending on the frequency of data capture, there will likely be some sort of re-keying of information to bring the application back into a production state. Also, the data capture quality impacts RTO. If the data was not in a backup state then some sort of indexing may be required. That indexing typically needs to complete prior to the application being brought back online.

In summary, RPO is the amount of data or transactions that can be lost when restoring from the last backup. RTO is the set amount of time to restore the application to a working state.

Lowering RPO and RTO

The lowering of RPO requires more data protection events occur. Simply put more backup or replication jobs need to occur more frequently. The more frequently that data is being captured, then the less data needs to be re-keyed in the event of a failure.

Lower RTOs can be acheived by improving network transfer speeds or by making a second copy of data rapidly available eliminating the need to copy it across a network. or upgrading the network between the secondary storage and primary storage. In general the lower RPO and RTO need to be, the more expensive and complex the data protection process becomes. One of the benefits of virtualization is that it reduces the cost to deliver improved RPO and RTO to application owners.

All Virtual Servers Are Not Created Equal

Servers and the applications they run have different levels of criticality to the organization. They are not all created equal. The data protection that is applied to these systems should match the specific RPO and RTO demands of the application. This means a balance should be struck between traditional backup and recovery software and some form of high availability (HA) solution. While virtualization has made delivering higher availability more affordable there is still a cost associated with it, typically in the form of additional software and extra disk capacity for live, secondary copies of data. Applying HA across all servers will waste IT budget unnecessarily.

The Role of Backup and Recovery

Traditional backup and recovery's role was to bring an application back to an RPO of 24 hours, assuming that backups were done once per night. The RTO in many cases was determined by the devices used, and usually the application had to be transferred across a network. Most data centers planned for a four hour minimum recovery time, plus the time to re-key information.

Virtualization specific data protection applications like Veeam have offered some improvement in this process. First, thanks to change block tracking (CBT), backups can be taken more frequently since less data has to be transferred over the backup network. But there is a limit to how many CBT backups can be done before they need to be consolidated into the full backup job. CBT backups can't be done continuously. They're usually done two or three times per day for important systems. This reduces the RPO significantly and for some systems and may be all the data protection they need.
Using the above example where the backup was done at 11:00pm and the failure happened at 4:00pm, with CBT another backup could have been executed at 12:00pm. This means that only four hours of information would need to be re-keyed instead of 17 hours.

Also, products like Veeam offer a capability called “Recovery In Place”, where a virtual machine’s volume can be re-instantiated directly from the backup device, assuming it is a disk target. This eliminates some of the transfer of data across the storage network and allows the re-keying process to start immediately. The downside to this capability is that the performance and reliability of the disk backup target do not likely match that of the primary storage system.

Virtualization specific backup solutions have enabled standard backup and recovery to play a larger role in meeting RTOs and RPOs. Backup should be considered a foundational part of any strategy designed to meet these objectives. But for most organizations, there will be certain applications that need a higher level of availability.

Reducing the Cost of Low RPO and RTO

Reducing RPO and RTO to a point below what basic data protection can provide used to mean a large capital outlay. To meet a very low RPO and RTO requires that data be replicated, synchronously or asynchronously, to a secondary server. In the past this required an expensive storage network, two enterprise class storage systems with this type of replication and a standby server waiting for the primary server to fail. In the same way that virtualization led to backup applications that expanded backup’s role in reducing RPO and RTO, virtualization provides a way to reduce the cost of providing very low RPO and RTO.

Virtualization eliminates the need for a dedicated standby server waiting for a failure to occur. All the hypervisor needs is access to the original virtual machine’s data. Of course, that data has to be constantly updated so that the RPO is minimized. Then, when there is a failure, the virtual instance can be started on another physical system accessing this second copy.

It is the role of software based storage solutions like those from StarWind Software to provide this constantly updated second copy of data. These software applications move the storage intelligence off of the expensive enterprise array and onto the same physical hosts that drive the rest of the virtual machines. This software can then replicate certain virtual machine volumes to other hosts within the environment.

If this replication is done synchronously, meaning that the write has to be confirmed in both locations prior to being acknowledged, then the second copy is in a perfect state and the RPO should be basically zero. The possible exception would be data that may be in the application’s cache. If capturing cache state is critical, this can be addressed by using the faul tolerance technquies within the hypervisor. A synchronous type of replication can be a mirror to a secondary system or, to provide even greater protection from failure, distribute data between multiple datacenters. In either case, if there is a host

failure critical virtual machines can be re-started on one of the alternate hosts rather quickly, since data is already in place and in a usable form.

The downside to either of these techniques is potential performance loss since the application has to wait for both writes to be acknowledged. This performance loss can be designed-around with proper network configuration and investment. Also some solutions like StarWind Software's address this with write back caching techniques.

The other downside is the consumption of additional storage capacity, up to 3X, so that these extra copies can be made available. Capacity consumption concerns can be offset by using high capacity SATA drives, as well as in-line deduplication to eliminate data redundancy, which is typically very high in virtual environments.

Software like StarWind greatly reduces the cost of providing higher levels of availability. This allows data centers to provide HA to a broader range of applications. But there is still some cost in additional resource consumption, so it remains important to identify only the VMs that truly need low RPO and RTOs.

Conclusion

Data protection is not a one-size fits-all situation, especially in virtualized environments. Virtualized backup applications can form the foundation of the data protection strategy, including disaster recovery. They can even create tape-based backups, providing comfort that data is protected on an entirely different type of media all together. But there are some applications that need to have as little downtime as possible. They need very low RPOs and RTOs. For these environments virtualization provides a way to greatly reduce the cost of delivering these aggressive objectives. Software based storage solutions like StarWind can help customers meet the requirments of very strict RPOs and RTOs while at the same time reducing the excessive cost by leveraging server side storage and eliminating the need for dedicated storage hardware.

George Crump

George Crump, Lead Analyst, Storage Switzerland, LLC
George Crump is President and Founder of Storage Switzerland. With 25 years of experience designing storage solutions for data centers across the US, he has seen the birth of such technologies as RAID, NAS and SAN. Prior to founding Storage Switzerland he was CTO at one the nations largest storage integrators where he was in charge of technology testing, integration and product selection.

Storage Switzerland

Web-site: http://www.storageswiss.com
E-mail: gcrump@storage-switzerland.com
Twitter: http://twitter.com/storageswiss
YouTube: http://www.youtube.com/user/storageswiss

Veeam Software

Global Headquarters	Americas Headquarters	EMEA Headquarters
+41-41-766-71-31	+1-678-353-2140 +1-614-675-9494	+33 1 75 61 27 40

Web-site: http://www.veeam.com/

StarWind Software

US Headquarters	EMEA and APAC
+1-617-449-7717 +1-617-507-5845	+44 20 3769 1857 (UK); +49 302 1788 849 (Germany) +33 097 7197 857 (France); +7 495 975 94 39 (Russian Federation and CIS) 1-866-790-2646

Support Portal: https://www.starwind.com/support
Support Forum: https://www.starwind.com/forums
Sales: sales@starwind.com
General Information: info@starwind.com

In 2016, Gartner named StarWind "Cool Vendor for Compute Platforms".

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact.

Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.