

## ONECLOUD SOFTWARE: DR FOR THE MASSES

*Makes AWS an integral part of the DR puzzle*

**OCTOBER 2014**



With the advent of server virtualization, many adopters erroneously think that disaster recovery (DR) is a problem of the past. They cite the ability of the hypervisors to replace the two most common yet imperfect DR choices: 1) infrastructure replication to a secondary replica site – fast to restore but very expensive, or 2) economical tape backup with off-site long-term storage – economical but slow to recover from.

The reality is that while server virtualization has certainly helped the industry get closer to simpler and less expensive DR products, DR still remains one of the major challenges for IT. This is especially true for applications that fall somewhere between the most mission critical where RTOs and RPOs of a few seconds is needed (and cost is often no object) and those that find RTOs and RPOs of a day or two to be adequate. Today, DR products available for these “intermediate” applications are few and far between, especially when overall cost of DR is considered.

The missing piece so far has been a cost-effective DR solution with excellent RTO and RPO for the *majority* of business applications -- *without* requiring a secondary site. OneCloud steps into the gap by replacing that expensive site with the hyper-scale public cloud. This Profile will discuss how OneCloud works to extend the primary data center onto the cloud, and how this impacts the ease and speed of VM recovery.

### ***How DR has Transformed in the Last Decade***

DR has been the biggest thorn in IT’s side since the beginning of time. For the most part DR has been reserved for the revenue-producing, most mission critical applications where cost was no object and the entire emphasis was on bringing the infrastructure up and running closest to when the problem occurred (RPO) and in the shortest amount of time (RTO). Some IT departments chose a DIY solution, but with priority applications they were more likely to hand over the problem to a DR site provider like SunGard. Whether IT did it themselves or paid for someone else to do it, massive amounts of time and resources were spent ensuring the two sites remained exactly alike, even in terms of software revisions and patches. Even then the only way one found out if the secondary site was recoverable was to declare a disaster and actually test if the failover occurred correctly and if the secondary site was operational.

As a result of this complexity and cost IT used this method to only protect those applications that were absolutely critical to the survival of the business. Other applications either did not have a DR plan or they were simply backed up on tapes and the tapes sent to an independent third party such as Iron Mountain for safekeeping. If a disaster did occur, these tapes were transported back to some secondary site where IT would create a new infrastructure to recover into. Cost was still high but most importantly RTOs and RPOs were measured in days and often weeks. And while the cost was

indeed lower than the first alternative, it was still too high and there was no guarantee that recovery would occur on the first pass.

All this changed for the better when server virtualization entered the picture a decade or so ago. Now the secondary site could be dissimilar to the primary site in all dimensions (server, networking, HBAs, storage, OS and other software revs, etc.). And DR could be applied to a lot more applications because implementation was easier and the cost of the secondary site could be reduced by using fewer and less expensive servers, storage and other gear. However, the cost of the secondary site in relation to the cost of the primary site is such that overall the second site still causes a 1.5x spend on IT gear.

Today there are a myriad of backup products in the market that offer replication functionality. Most of these offer RTOs and RPOs of a few hours to a day or two; some even offer RTOs and RPOs of a few minutes, at least in theory. Regardless, all these require a secondary site to recover into with all its attendant expenses. On the other end of the spectrum there are products designed for the most mission critical applications. They certainly are capable of very low RTOs and RPOs but they are expensive and still require a secondary site. While the secondary site infrastructure can be different from the primary site (and this is a huge improvement due to server virtualization), it is still required.

## ENTER THE CLOUD

The remaining issue is three-fold for applications that are not mission-critical but are business-critical: 1) how to DR-enable them with RTOs and RPOs in the 15-minute to 24 hour range, 2) do so cost-effectively, and 3) not require a separate DR site. Before OneCloud, organizations either had to spend on expensive, redundant infrastructure to hit these more modest RTO/RPO goals or live with the fact that many important applications would be left under-protected

A number of products are available in the market that can send a backup stream into a public cloud. But in most cases, these products use the cloud as a repository for storing data offsite. They present a separate interface for the appliance itself (Taneja Group calls this Cloud Integrated Enterprise Storage) and another one for the cloud of choice.

It is a step towards using the public cloud but it is not fully integrated into the cloud, and does not use the cloud as a DR site. A secondary site may still be needed to perform the recovery, and this requires the VMs to be moved over the WAN to the primary or recovery site. Moving one VM may pose little problem but if the primary site is unavailable for any reason and hundreds of VMs have to be moved into the secondary data center, the issue becomes non-trivial. In fact, it becomes impossible given the speed of affordable WANs. RTOs will be measured in days or weeks. One can consider this two steps forward and one step backwards—still not a complete solution.

It is within this landscape that OneCloud Software enters the market for DR. It does not purport to replace a backup application. Rather it is dedicated DR software that is designed to make the cloud a true partner in crime. It aims squarely at those applications that fit in the middle, those we are calling “intermediate” applications in terms of their RTO/RPO requirements.

## *OneCloud Software*

Headquartered in Boston, OneCloud Software was founded in 2012 and is managed by a team of experts from VMware, EqualLogic, Altiris, and ExaGrid. OneCloud Software recognizes that it is difficult for IT staff managing complex infrastructures to easily and cost-effectively consume hyperscale public cloud resources to solve real problems such as DR. OneCloud is solving this

problem by using the hyperscale public cloud as a seamless extension of their customers' virtual infrastructure. In place of a primary data center stack replicating to a secondary data center, OCS creates a single flexible data center consisting of on-premise virtualized environments and a cloned environment in the cloud. The user can now have limited to no expertise in how to directly manage the cloud platform due to OneCloud's extensive cloud automation.

OCS does this with their Automated Cloud Engine (ACE) with OneCloud Recovery. ACE replaces the need for redundant infrastructure for protecting VMs and is the enabling technology for OneCloud's product portfolio. ACE capitalizes on Amazon Web Services (AWS) to create a single flexible data center comprising on-premise virtualization with the hyper-scaled public cloud. The result is a seamless, centrally managed, highly scalable and economical data center. The OneCloud Recovery product runs atop ACE to provide economical data protection to VMs, replacing the expensive and redundant data center.

## INSTALLATION

- ACE works with the customer's existing VMware environment. No agents need to be installed on VMs themselves. IT installs ACE as a virtual appliance onto an ESX host. ACE utilizes the VMware APIs for all on-premise operations and creates an instance of itself in AWS. ACE retains AWS' customer benefits such as specifying geography for regulated data.
- ACE launches Auto Discovery to discover and replicate the on-premise virtualized data center in the cloud. The resulting blueprint includes VMs, storage, and network architecture. ACE deduplicates and compresses before replicating this copy into the cloud.
- ACE's Policy Engine leverages stated RTO and RPO goals to automatically provision resource allocation, prioritization and configuration settings in AWS.

## ONGOING

- Subsequent failover can occur on-demand. Upon failback, ACE only brings changed data back to the primary site.
- ACE maintains two application level consistency groups: 1) individual VMs and 2) bring-up order to restore an application consisting of multiple VMs. This enables customers to quickly recover key applications and data. ACE also validates successful replication to the cloud, and data is secure with in-flight and at-rest encryption. Only the user has access to the cloud account and keys. Auditing, reporting, alerts and monitoring complete the security picture.
- ACE uses VMware's change block technology along with deduplication and compression to keep data movement and storage capacity needs down.

## COSTS

- Customers do not need to be AWS experts to tap value from its service. ACE automatically moves the data between EBS, and S3 storage (Glacier in the future), based on the RTO/RPO policy. This process keeps customer costs to a minimum, without requiring the customer to know the inner workings of AWS. All that complexity is handled by ACE.
- The cloud replaces the secondary site. There are no capital requirements to buy and equip a secondary data center, nor to pay high operating expenses to co-locate one. There are very limited compute, memory, flash or database costs except in a failover condition and even

then only for as long as the failover condition lasts. The only recurring costs are for storage and even those are kept down by using the right tier and deduplication/compression technologies. Data movement to/from the primary data center to/from AWS use WAN optimization technologies to keep the bandwidth requirements down.

- IT does not have to be familiar with AWS' complicated billing practices. OCS estimates AWS costs and optimizes customer configurations accordingly. Initial settings will be at the lowest cost for the SLA configurations, and as applications and data grow ACE will continue to seek the lowest possible cost.

### ***Taneja Group Opinion***

DR has come a long way over the past decade. And server virtualization has a lot to do with it. Backup vendors added virtualization support and replication software to relieve the pressure of building an identical secondary site. Now IT could DR-protect a VM at the primary site by taking a snapshot and replicating it to a secondary site very cost effectively. These technologies have been a very strong step forward and made DR feasible for most applications. Server virtualization has been at the heart of this transformation.

However, while virtualization solves a lot of problems nothing is perfect. Virtualization vendors often state explicitly or implicitly that once IT virtualizes their environment, their DR problem is solved. Cloning, snapshots and replication are simple to do; problem solved.

Except it's not. Clones and snapshots must go somewhere *that is not the primary site*. Traditionally this means either a secondary site with potentially high expenses and ongoing management, or the cloud with slow and expensive data movement and still the requirement for a secondary site to recover into.

OneCloud's ACE with OneCloud Recovery gives customers another choice. They replace the secondary site's capital costs and lower ongoing costs. And they do it with RTOs and RPOs designed for many tier 1 as well as all tier 2 and 3 virtual applications.

There are a myriad of DR solutions in the market today from backup vendors to specialty DR vendors to VMware itself. However, there is a big difference between these vendors and OneCloud: OneCloud provides better than acceptable RPOs/RTOs for business-critical applications, *and* automatically optimizes data movement for the lowest possible pricing within AWS. It is high time someone made the public cloud an integral part of the DR picture. OneCloud Software just did. We expect this product to appeal to companies from mid-size to large enterprises who are looking for cost optimization and simplicity, and the opportunity to protect many more applications at reasonable costs than ever before.

NOTICE: The information and product recommendations made by Taneja Group are based upon public information and sources and may also include personal opinions both of Taneja Group and others, all of which we believe to be accurate and reliable. However, as market conditions change and not within our control, the information and recommendations are made without warranty of any kind. All product names used and mentioned herein are the trademarks of their respective owners. Taneja Group, Inc. assumes no responsibility or liability for any damages whatsoever (including incidental, consequential or otherwise), caused by your use of, or reliance upon, the information and recommendations presented herein, nor for any inadvertent errors that may appear in this document.