Arcserve Solutions for Amazon Web Services (AWS)

Introduction

The public cloud has become a major factor in all IT decision making due to its endless supply of technical resources at an affordable cost. In fact, it is quickly reshaping the way organizations of all sizes are looking at their technology needs as its affordability is now making available new levels of business continuity that were once reserved for only the largest (and best funded) organizations.

Arcserve, too, understands the potential of the public cloud. Working in conjunction with its customers, Arcserve has identified several key areas of data protection that can be transformed by the public cloud; specifically, the Amazon AWS public cloud as it has rapidly become a popular choice for many organizations. Our customers and partners have told us they want solutions that are truly integrated with AWS without requiring deep cloud technical skills, and we've listened. Arcserve solutions for AWS were purposefully designed to meet customer needs with the flexibility to leverage AWS in multiple ways for offsite data protection, application test and development, application failover and full disaster recovery as a service (DRaaS).

- Economical disaster protection for systems, applications and data on physical servers and virtual machines using AWS
- Conversion of Recovery Points to AWS EC2 virtual machine format and automatic snapshot creation for application test and development
- Supports Amazon Elastic Block Storage (EBS) block storage volumes and S3 blob storage volumes for long term retention

Long Term Backup Retention

Long standing best practices dictate storing at least one, if not multiple copies offsite. The common practice has been to copy each backup to disk to tape, and ship the tape offsite in a protected storage location. Third party vendors offer this service to countless organizations as a proven way to protect valuable backup images against fire, floods and other natural disasters; however, two important advances in technology are reshaping this time-proven practice. One is high bandwidth (and affordable) network connections, and the other being low cost cloud blob storage. AWS was the pioneer in this area, offering its S3 blob storage for as little as $0.02 per GB per month. For 1TB of data, this equates to $20 per TB per month. A shockingly low price when compared to traditional on-premise storage.

S3 storage is accessed in the same way as any network file share with no changes required. The limitation of this approach is the complete reliance on the customer to configure the S3 storage account. Customers lacking deep cloud skills found this to be a barrier., while also precluding any assurance that the solution is tested and will perform up to standards.

Arcserve UDP took a different and more thorough approach to using AWS S3 storage, specifically by offering backup plans that include AWS targets. The configuration screen manages key information, such as AWS account credentials, storage selection, compression and encryption. Using this information, an Admin with little cloud experience can configure a backup plan that includes AWS.

Further, Arcserve UDP supports three forms of backup to S3. First is to copy Recovery Points to S3. (Figure One-A.) Recovery Points are backup images created by Arcserve UDP and managed with a retention policy (e.g. daily, weekly, monthly and yearly) along with location. Using it as a location, specific Recovery Points can be copied automatically to S3. For example, a backup plan can be configured to take backup images every four hours, keeping 12 copies, plus six daily copies, four weekly copies, four monthly copies and two yearly copies. Using Arcserve UDP, Recovery Points can be restored back from S3 by downloading the Recovery Point to a local site before performing granular restores from it. One additional use case in this context is to use S3 as offsite tape replacement for storing monthly backups.

In addition to storing backup images in S3, Arcserve UDP supports File Copy and File Archive to S3. Using File Copy, you can copy elected source files to a destination. Using File Archive, you can archive selected source files to a destination whereby the destination can be a cloud account or a network share. File Archive allows you to safely and securely delete the source data after it has been copied to an offsite or secondary storage repository.

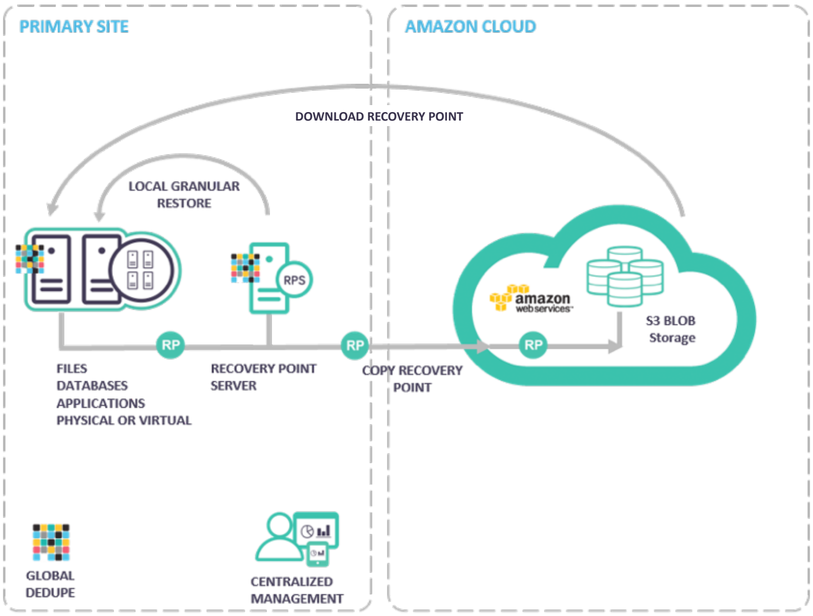

Figure 1A. Arcserve UDP Copy Recovery Point to AWS

Arcserve UDP Tape Module VTL Gateway

The tape capability in Arcserve UDP also allows one to leverage the Amazon VTL Gateway. In this scenario, Arcserve UDP 6.5 (or Arcserve Backup v17.5) will treat it like any other physical tape library. Backups from all agents can be written to this Virtual Tape Library, with Amazon Gateway managing the backups on the Amazon side.

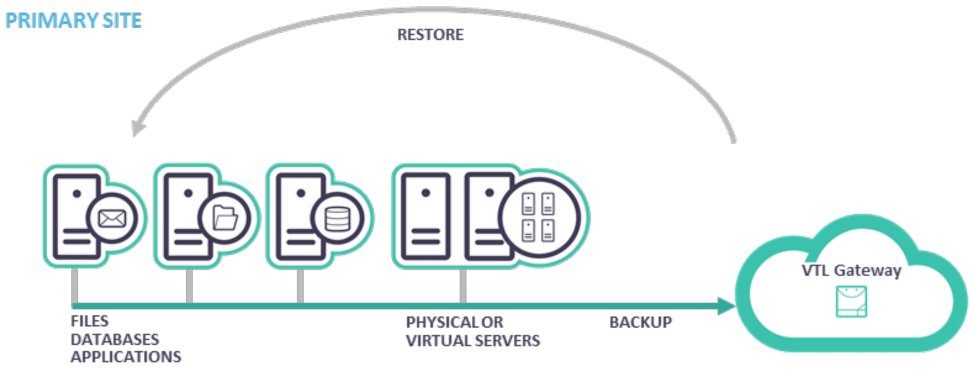

Figure 1B. Amazon VTL Gateway

Boot EC2 VM from Recovery Point

For the majority of organizations, the first foray into public cloud storage is long term retention of backup images on S3 as a replacement for tape. The next use case takes matters one step further. Consider that in addition to low cost S3 blob storage, AWS also provides a vast array of compute options at a very affordable cost. Compute charges are based on actual usage and are billed monthly, whereas S3 costs are based on capacity. Modern data protection solutions offer the ability to convert an image-based backup to a virtual machine (ESX and Hyper-V), connect to a hypervisor and boot up. Arcserve UDP offers the capability to perform P2V and V2V failover from any physical, Hyper-V and ESX to Amazon EC2 format.

For AWS, Arcserve looked closely at its existing virtual machine recovery capabilities and integrated them closely with AWS - supporting two methods of virtual machine recovery:

- Arcserve UDP Virtual Standby supports AWS EC2 and EBS volumes to protect Windows systems. To create a virtual standby machine using AWS EC2, a backup task performs a backup of the source nodes and a Virtual Standby task (using AWS EBS as a destination) transfers the Recovery Point "on-the-fly" to AWS and converts to Amazon EC2 virtual machine format. The Recover Point data stored in EBS incurs no EC2 charges. A Windows Backup Agent, installed on an AWS EC2 instance, automatically converts the Recovery Point to EC2 format. The Proxy running in EC2 will incur some charges for compute.

- Arcserve UDP Instant VM supports AWS EC2 and EBS volumes to protect Linux systems. To create an Instant VM using AWS EC2, a backup task performs a backup of the source nodes and uses AWS EBS as the destination to store the Recovery Point. A Linux Backup Agent is installed on an AWS EC2 instance and converts the Recovery Point to Amazon EC2 format in-place when manually started. The Recover Point data stored in EBS incurs no EC2 charges.

Arcserve UDP VSB and IVM both perform virtual machine failover using AWS EC2 format, but there is one important difference: VSB requires additional storage for the new virtual machine, whereas IVM does not. For the fastest recovery time, VSB has the advantage however it requires additional storage. Figure two illustrates the basic methodology for both VSB and IVM.

Using a proxy server running on EC2, a Recovery Point is converted to an Amazon EC2 virtual machine. EC2 compute charges are incurred for the proxy server and the new virtual machine for the duration that they are running.

Using either VSB or IVM, Arcserve UDP enables virtual machine failover in AWS. AWS charges are incurred to store Recovery Point image files and EC2 compute charges when the proxy server and new virtual machine are running.

EC2 can only boot from EBS volumes and not from S3. As for VSB, it should be noted that the Recovery Points are stored in EBS volume and not in S3.

A popular application is test and development. On-premise applications under the protection of Arcserve UDP can be easily converted to running VMs in AWS for little cost. When asked by engineering for new machines to test development, IT can spin up machines in AWS and avoid the cost of new on-premise hardware.

Another great benefit is, of course, disaster recovery. In this case, Recovery Point Objectives (RPOs) and Recovery Time Objectives (RTOs) can be significantly improved over other methods, including transferring backup data “back home”. In some cases, there is no more “home” if premises have become unavailable for whatever reason.

For the protection of virtual machines running in EC2 or to run a mirrored site, a Recovery Point Server is installed in AWS.

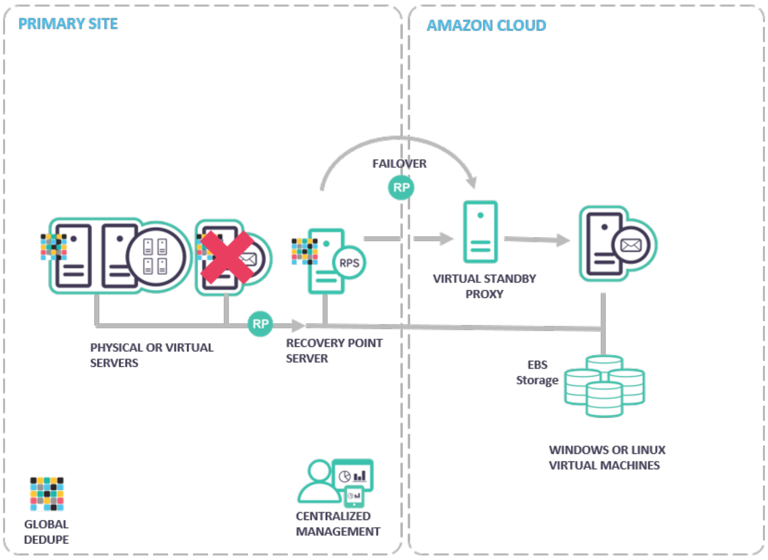

Figure 2. Arcserve UDP VSB and IVM in AWS

Arcserve UDP with RPS/Console Running in EC2

Since the inception of Arcserve UDP, it has supported the ability to run the Recovery Point Server and Console in EC2. In this case, and unlike the use of the Arcserve UDP proxy, compute is needed as data protection workflows are processed between the primary site(s) and EC2. By using a Recovery Point Server (RPS) running in EC2, customers can replicate from the RPS on the primary site with built-in WAN optimization and Global Deduplication Optimization that speeds data transfer and requires less cloud storage.

For the cost of running an RPS on EC2, many backup and recovery options are available:

- Protect EC2 Virtual Machines with an agent inside the EC2 VMs
- Back up an EC2 instance to local /NFS/CIFS/RPS
- Restore an EC2 instance (through IVM for Linux and VSB for Windows with auto recovery option)
- Set a longer retention policy to optimize the primary (on-premise) site
- Protect Exchange Online running in Office 365
- Replicate data from EC2 back “home” to keep a local copy of cloud workloads

As shown in figure three, an RPS can be installed in AWS EC2 where it serves as a backup destination for the on-premise RPS. This configuration supports WAN optimized replication and deduplication, as well as full incremental backup integration; in addition to VSB and Instant VM for Windows and Linux systems. An RPS running in AWS EC2 incurs added compute and storage fees; without an RPS in AWS EC2, no EC2 fees are incurred until the installation of the Backup Agent on EC2 instance.

To protect systems, applications and data running in AWS EC2, Arcserve UDP supports local backup of EC2 instances using the Linux Backup Agent. This feature offers improved RPOs and RTOs in AWS if you desire to enhance standard AWS recovery capabilities. EC2 Windows instances can also be protected with Windows Backup Agent. Full VM and file-level recovery is supported for Linux systems. You can restore an EC2 instance through Instant IVM for Linux and VSB for Windows to Amazon EC2 with auto recovery option. File level recovery is also supported for Windows EC2. Also with BMR to alternate hardware, backup of Windows EC2 instances could be restored to any on-premise physical/virtual machine.

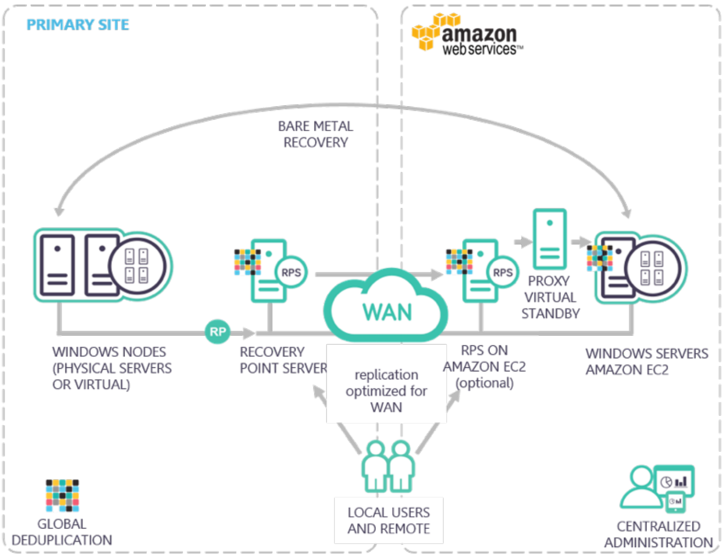
PRIMARY SITE
amazon web services™
BARE METAL RECOVERY
RPS
RP
WAN
RPS
PROXY VIRTUAL STANDBY
WINDOWS NODES (PHYSICAL SERVERS OR VIRTUAL)
RECOVERY POINT SERVER
replication optimized for WAN
RPS ON AMAZON EC2 (optional)
WINDOWS SERVERS AMAZON EC2
GLOBAL DEDUPLICATION
LOCAL USERS AND REMOTE
CENTRALIZED ADMINISTRATION

Figure 3. RPS Server running on EC2

High Availability and Continuous Data Replication

Applications that require 24x7 continuous operation with no loss of data are candidates for Arcserve High Availability (HA). Arcserve High Availability provides full system HA of physical or virtual machines to the AWS EC2 cloud. This allows users to not only protect their local site with continuous replication into EC2, but to failover with multiple bootable copies that can be powered on instantly in case of disaster at the source site. Arcserve High Availability can be used for failback to on-premise using bare metal restore. In this case, compute is needed to run the HA engine in EC2. It is a proven solution with support for many-to-one replication, by running only one active instance in EC2.

Figure four shows a typical Arcserve High Availability deployment in Amazon. With Agents deployed on each application server, application data is continuously replicated between on-premise and EC2. A server heartbeat monitors the connections and is ready to failover in an instant, should the connection go down. For EC2 failover, the standby replica initiates the failover procedure if the master server becomes unresponsive. During failover, a new EC2 recovery instance of the same major OS version as the master is started using one of the predefined and supported AMIs. Replicated EBS volumes are detached from the Arcserve High Availability Appliance and are attached to the Full System EC2 recovery instance. The Full System EC2 recovery instance is started.

Arcserve High Availability is truly a full feature high availability solution to provide near instant failover for your most mission critical applications with no loss of data.

The following illustration represents a Full System EC2 scenario protecting an on-premise server with two EBS volumes before failover occurs:

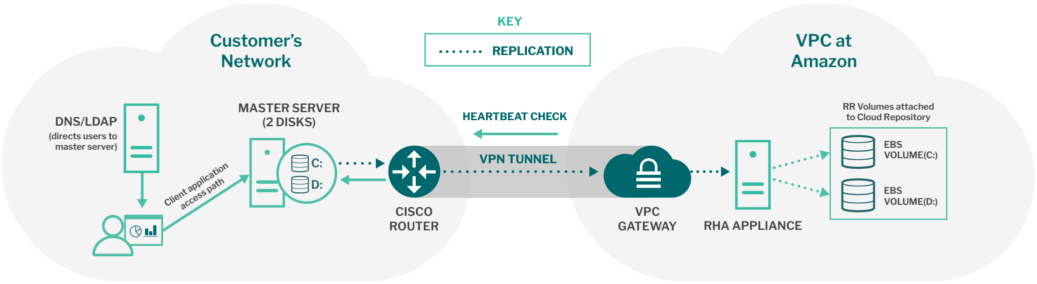

The following illustration shows what happens once failover occurs:

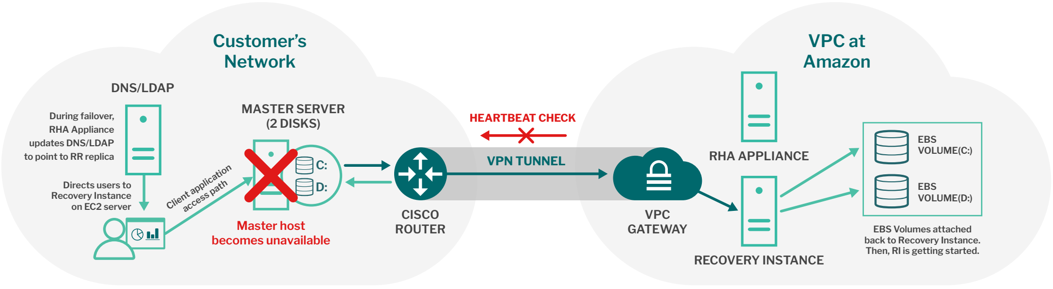
Customer's Network
DNS/LDAP
During failover, RHA Appliance updates DNS/LDAP to point to RR replica
Directs users to Recovery Instance on EC2 server
Client application access path
MASTER SERVER (2 DISKS)
C:
D:
Master host becomes unavailable
HEARTBEAT CHECK
VPN TUNNEL
CISCO ROUTER
VPC at Amazon
VPC GATEWAY
RHA APPLIANCE
RECOVERY INSTANCE
EBS VOLUME(C:)
EBS VOLUME(D:)
EBS Volumes attached back to Recovery Instance. Then, RI is getting started.

Figure 4. Arcserve HA installed in Amazon EC2

Summary

The high level of flexibility in leveraging AWS afforded by Arcserve UDP is simply unmatched in the mid-market. Simply said, the unified architecture of our solution and its robust enterprise class feature set were designed for cloud. As organizations evolve their infrastructure to become inherently hybrid (on-premise/cloud), a solution that effectively and efficiently guarantees data protection SLAs is a must have. In designing the most recent version of Arcserve UDP, we placed a high premium on developing robust integration with Amazon AWS, each well suited for a different level of RPO/RTO and cost profile depending on what services are consumed.

For more information on Arcserve, **please visit** arcserve.com

Copyright © 2017 Arcserve (USA), LLC and its affiliates and subsidiaries. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective owners. This document is for your informational purposes only. Arcserve assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, Arcserve provides this document "as is" without warranty of any kind, including, without limitation, any implied warranties of merchantability, fitness for a particular purpose, or non-infringement. In no event will Arcserve be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, business interruption, goodwill or lost data, even if Arcserve is expressly advised in advance of the possibility of such damage.