



NetBackup in the Public Cloud

Guidelines for AWS Deployments

This technical paper is designed to provide assistance to partners and end users looking to protect workloads with NetBackup deployed in AWS cloud. The guidelines within this technical paper will assist partners and end users as they design and implement data protection solutions based on Veritas products in the public cloud. In addition to these guidelines, partners and end users should also leverage product documentation, Veritas Educational Services and/or Veritas Consultancy Services when necessary.

VERITAS™

The truth in information.

TABLE OF CONTENTS

INTRODUCTION	4
PREREQUISITES.....	4
NETBACKUP OVERVIEW.....	4
KEY CAPABILITIES	4
KEY FEATURES	5
BUSINESS VALUE	5
WHY ARE CUSTOMERS LEVERAGING THE CLOUD?	5
NETBACKUP AND AWS DESIGN OVERVIEW	6
NETBACKUP AND CLOUD CONNECTIVITY	6
NETBACKUP AND CLOUD RESTORE OPTIONS	7
AWS CLOUD VERSUS ON-PREMISES CONSIDERATIONS	7
USE CASES OVERVIEW.....	8
STANDARD BACKUP TO AWS S3 OBJECT STORAGE.....	8
BACKUP TO AWS S3 OBJECT STORAGE WITH CLOUDCATALYST DEDUPLICATION	9
SENDING DATA TO AWS S3 USING A THIRD-PARTY GATEWAY APPLIANCE	10
AMAZON STORAGE GATEWAY VTL	11
BACKUP IN THE CLOUD – AWS-ENABLED NETBACKUP ARCHITECTURES IN EC2	12
DISASTER RECOVERY USING AWS.....	14
AUTO IMAGE REPLICATION (AIR) TO THE CLOUD – HYBRID CONFIGURATION	14
<i>Leveraging NetBackup 8.2 CloudCatalyst Image sharing for Migration and DR</i>	<i>16</i>
CLOUD SIZING AND PERFORMANCE	17
AWS INSTANCE MODEL	18
AWS STORAGE OPTIONS.....	19
ENVIRONMENT DESCRIPTION AND ASSUMPTIONS FOR SIZING	19
NETBACKUP AWS INSTANCE SIZING	20
<i>NetBackup Master Server</i>	<i>20</i>
<i>Master Server recommendations</i>	<i>21</i>
<i>NetBackup MSDP storage</i>	<i>21</i>
<i>NetBackup CloudCatalyst</i>	<i>22</i>
<i>Growing the Media Server</i>	<i>23</i>
<i>Media Server Deduplication Pool recommendations</i>	<i>23</i>
ADDITIONAL ARCHITECTURE REQUIREMENTS	25

SECURITY OF THE INFORMATION	25
<i>Least privileged access</i>	26
<i>Limit access with Resource Tags</i>	27
<i>About AWS Service Quotas</i>	27
<i>AWS Service Quotas and NetBackup</i>	28
<i>Disaster recovery scenarios</i>	28
COST OVERVIEW.....	29
<i>Storage costs</i>	29
<i>Compute costs</i>	29
DEPLOYMENT DETAILS	30
NETBACKUP VPC DEPLOYMENT CONFIGURATIONS	30
SETTING UP NETBACKUP CLOUDCATALYST IN AWS.....	31
<i>Using the AWS marketplace</i>	31
<i>Launching an Elastic Compute Cloud (Amazon EC2) instance</i>	33
<i>Configuring CloudCatalyst on the EC2 instance</i>	35
PROTECTING NETBACKUP ACCESS WITH EC2 SECURITY GROUPS.....	36
TAGGING NETBACKUP RESOURCES	37
ROTATING AWS ACCESS KEYS FOR CLOUDCATALYST	39
PROTECTING NETBACKUP FROM FAULTS, FAILURES, AND DOWNTIME	40
NETBACKUP RISK AND AUDIT MANAGEMENT.....	41
<i>Enabling CloudTrail logging for NetBackup resources</i>	42
<i>AWS Scheduled Service events</i>	43
SUMMARY	43
APPENDIX A – ADDITIONAL INFORMATION	45
APPENDIX B – TERMINOLOGY	46

INTRODUCTION

The purpose of this technical paper is to provide a technical reference on the capabilities of NetBackup and Amazon Web Services. While this guideline is a stand-alone document, additional information can be found using the links in the Additional Resources section. This document is not a replacement for the NetBackup Cloud Admin Guide, links to which can be found at the end of this document.

Veritas has partnered with Amazon Web Services (AWS) to offer a robust backup to and in the cloud experience. Each solution can be tailored to the individual needs of the customer.

NOTE: This document contains recommendations that have been shown to work with customer deployments. Understand that every environment is unique, and changes might be required. In addition to these guidelines, always leverage product documentation and any additional services (Educational or Consultancy) to ensure the best design for their unique environment and workloads.

PREREQUISITES

This document is intended for individuals who have a basic understanding of AWS cloud infrastructure concepts. Users should be familiar with AWS CloudFormation, AWS Identity and Access Management (Roles and Policies), Amazon S3 (Bucket Policies and ACLs), Amazon EC2, Amazon VPC and storage concepts that are related to enterprise backup and recovery solutions.

NETBACKUP OVERVIEW

As an established market leader in data protection, Veritas NetBackup provides unparalleled next-generation data protection by minimizing costs and complexity and ensuring greater business continuity with a solution that unifies data protection across the entire enterprise.

KEY CAPABILITIES

- *Comprehensive* – As a single solution to protect all your data assets, NetBackup provides support for virtually every popular server, storage, hypervisor, database, and application platform used in the enterprise today.
- *Scalable* – High performance, elastic automation, and centralized management based on a flexible, multi-tier architecture enables NetBackup to adapt to the growing needs of a fast-paced, modern enterprise datacenter.

- *Integrated* – From backup appliances to big data platforms, NetBackup integrates at every point in the technology stack to improve reliability and performance. OpenStorage Technology (OST) provides even tighter integration with third-party storage and snapshot solutions.
- *Innovative* – With hundreds of patents awarded in areas including backup, recovery, virtualization, deduplication, and snapshot management, NetBackup continues a long tradition of bringing advanced technologies to market first.
- *Proven* – For over a decade, NetBackup has led the industry as the most popular enterprise data protection software by market share and is used by many of the largest enterprises on the planet. When you need your data back, you can trust NetBackup.

KEY FEATURES

- One platform, one console unifies virtual and physical global data protection
- Unified global management of snapshots, replicated snapshots, backup, and recovery
- Scalable, global deduplication across virtual and physical infrastructures
- Single pass backup, instant image and single file restore for virtual and physical
- Automated virtual data protection and load balanced backup performance

BUSINESS VALUE

With so many Veritas customers considering the cloud either as a supplemental data center, hybrid model, or getting rid of the traditional data center and running everything in AWS, the ability to not only protect the data, regardless of the location, but also to assist in moving the workloads to the cloud with the use of products within the Veritas suite provides extensive value. Whether it is a disaster recovery requirement, or the desire to not have to manage a physical data center, customers are thinking “cloud” more often. And Veritas is there to help them every step of the way.

WHY ARE CUSTOMERS LEVERAGING THE CLOUD?

Customers are using the cloud for several reasons. Smaller customers like the idea of not having to maintain a datacenter and an expensive DR site. Mid-range customers enjoy having an offsite copy of their data that is built on highly scalable offsite hardware. Large customers would have large datacenters already but use the Cloud for offsite workloads in instances where it makes more sense to have it at a cloud provider versus an internal data center. Sometimes a customer will need a temporary space to utilize a workload and instead of ramping up a new rack of disk in a datacenter, they can utilize the same functionality at a cloud provider and once finished, can remove the environment without the additional cost of purchased hardware sitting in the data center. The subscription model works very well for these sorts of projects with a highly scalable, simple to use model.

The current megatrend of moving data to the cloud revolves around driving costs down for business. The cloud model is also very agile when it comes to requirements. Additional disk can be added to a server very quickly and easily versus having to source the hardware and the rack and stack that comes along with it.

There is also the aspect of hardware updates. New firmware for arrays comes out on a regular basis. A customer would need to manage these at a local datacenter whereas in the cloud this requirement is taken care of by the cloud provider.

Each customer will have a different reason for a move to cloud-based computing. And Veritas is here to help with each scenario.

NETBACKUP AND AWS DESIGN OVERVIEW

There are many design cases when it comes to NetBackup and AWS. This section will outline them from a high level. Specific use cases are included in the next section.

NetBackup has created an Amazon Machine Image (AMI) option that will allow for very fast provisioning of NetBackup instances. In an EC2 instance, the AMI can be selected and with several configuration options the NetBackup Master or Media server can be deployed with the necessary CPU, RAM, and disk configuration along with a custom IP address needed for use by the customer. Elastic Block Storage (EBS) is used as the disk storage for these machines which uses block-based storage versus the object-based storage of S3. This provides the necessary storage requirements for NetBackup's storage optimization features like traditional MSDP, Advanced Disk, or CloudCatalyst.

This methodology is simply a "virtual environment" similar to the other virtual environments that can be created using virtualization vendors. Virtual machines are spun up as needed and appear to be physical machines from a management perspective. The AMI can also be used within a Cloud Formation Template (CFT) script that uses a JSON scripting system to automate NetBackup instance creation. Therefore, deployment of NetBackup, either large or small, has been simplified to filling out a short deployment form. The NetBackup CFT is available in the AWS Marketplace for easy deployment.

NETBACKUP AND CLOUD CONNECTIVITY

NetBackup can utilize AWS in several different ways depending on the needs of the customer. As outlined in various places in this document NetBackup can utilize S3 object storage, Glacier, or Glacier Deep Archive to send data to AWS storage similar to a regular disk pool or utilize the NetBackup CloudCatalyst to send deduplicated data to S3 object storage more efficiently. If a customer has resources in the cloud, NetBackup can also be installed in the cloud and used to protect these resources in a similar manner to protecting physical resources in a datacenter.

Cloud instances are managed using a web-based UI whereby systems in the Cloud can be monitored and managed, or other connectivity methods such as RDP can be used to connect.

NETBACKUP AND CLOUD RESTORE OPTIONS

Restores of information in the cloud are as simple as in a local data center. The backup admin has full use of the UI to recover information. Depending on the type of storage (S3, EBS or Glacier) restores will be fast. A section at the end of this document goes over the restore process of basic S3 storage and includes screen shots. It is beyond the scope of this document to go over all restore scenarios.

AWS CLOUD VERSUS ON-PREMISES CONSIDERATIONS

Running traditional IT workloads in the cloud can have significant benefits if designed and architected correctly. However, if architected improperly one could end up paying an unexpected price in terms of cost, workload performance and management headache.

When protecting workloads in the cloud consider the following:

- **IOPs available**
 - **On-Premise:** In an on-premises environment if you have specific IOPs requirements you can select the appropriate hardware to meet those needs.
 - **AWS Cloud:** In AWS you can make your selection of EBS volumes based on IOPs requirements. There is a cost associated with guaranteed IOPs.
- **Peer Link Limits:**
 - **On-Premise:** In an on-premises environment you can basically have as many peer-to-peer links as required.
 - **AWS Cloud:** In AWS there are fixed limits on the number of VPC to VPC peer links that are allowed.
- **Storage targets**
 - **On-Premise:** In an on-premises environment you typically write to block storage, deduplication devices or MSDP pools.

- **AWS Cloud:** In AWS your storage targets are typically EBS (Elastic Block Storage) or object storage S3 (Simple Storage Service). EBS is generally more expensive while S3 is more durable.

USE CASES OVERVIEW


There are several different use cases with NetBackup and AWS. These will be outlined in this section. This list is not comprehensive. Use cases will vary with customer requirements, however the use cases presented will outline the more popular options.

STANDARD BACKUP TO AWS S3 OBJECT STORAGE

With NetBackup, the simplest way to move data to object storage is to use the standard cloud connector interface. This interface allows the user to configure a cloud storage target or any number of providers, including Amazon and Amazon GovCloud S3 targets.

BACKUP TO THE CLOUD – STANDARD OBJECT STORAGE

NETBACKUP TO THE CLOUD



- Protect On premise and cloud based workloads with the same methods
- Source and target can be converged Master & storage or separate instances
- Segment data can be encrypted with KMS managed keys before transport and will be stored encrypted

Figure 1 – Using standard object storage

This functionality allows for a straightforward and easy to implement cloud storage target that can be written to from any Master or Media server. Standard charges apply based on data ingress and egress charges as documented on the AWS S3 pricing page:

<https://aws.amazon.com/s3/pricing/>

BACKUP TO AWS S3 OBJECT STORAGE WITH CLOUDCATALYST DEDUPLICATION

The NetBackup CloudCatalyst solution combines the performance and flexibility of NetBackup with powerful data deduplication technology to better leverage the cloud for storing backups for disaster recovery or long-term data retention. By ensuring backup data remains optimized while in transit to the cloud and while at rest in the cloud, the NetBackup CloudCatalyst solution greatly reduces cost and increases performance when using cloud storage.

The NetBackup CloudCatalyst can be delivered on a purpose-built appliance, a virtual appliance, or as a build-your-own (BYO) software solution. It allows customers send backup data to cloud object storage in deduplicated form. CloudCatalyst can receive optimized backup images from existing MSDP volumes or directly from a client and transfer the data to a S3 public or private cloud object storage target. A wide variety of S3 object storage has been certified for use with NetBackup CloudCatalyst.

When using MSDP volumes as the source, the CloudCatalyst server does not rehydrate, or remove optimization from deduplication. This end to end deduplication is a significant difference in how the CloudCatalyst solution operates as compared to other vendors in the market today. The CloudCatalyst Server allows direct recovery of data from the CloudCatalyst Server without first passing through another Media Server. Using CloudCatalyst will provide the highest level of functionality and cost savings when using object storage.

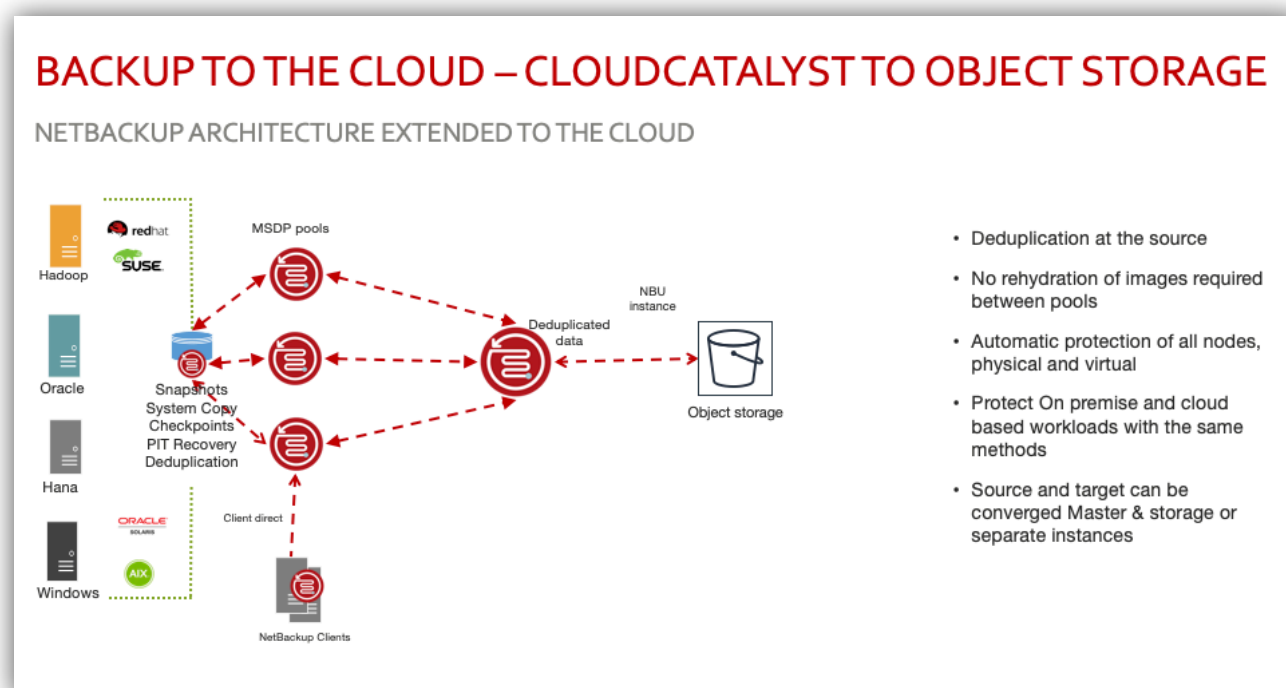


Figure 2 – Sending Data to AWS S3 Using Cloud Catalyst

SENDING DATA TO AWS S3 USING A THIRD-PARTY GATEWAY APPLIANCE

This use case is for the customer that is trying to get away from maintaining a tape infrastructure locally in the datacenter by moving to disk and would like to take advantage of the conveniences of cloud-based storage. This solution uses a 3rd party deduplication appliance on-premises which reduces the amount of data sent to the cloud. From a NetBackup standpoint, the dedupe appliance looks like a disk storage unit. Backups are sent to the appliance like backups sent to any disk pool. The appliance itself will perform the deduplication, therefore only the changed blocks are forwarded on to the cloud via the S3 connector.

This solution will work with any S3 compatible gateway that presents itself as a disk target to NetBackup, allowing data to be deduplicated for the given environment. Unlike CloudCatalyst, third party deduplication appliances are not compatible with MSDP and will require the data be rehydrated prior to sending to the device.

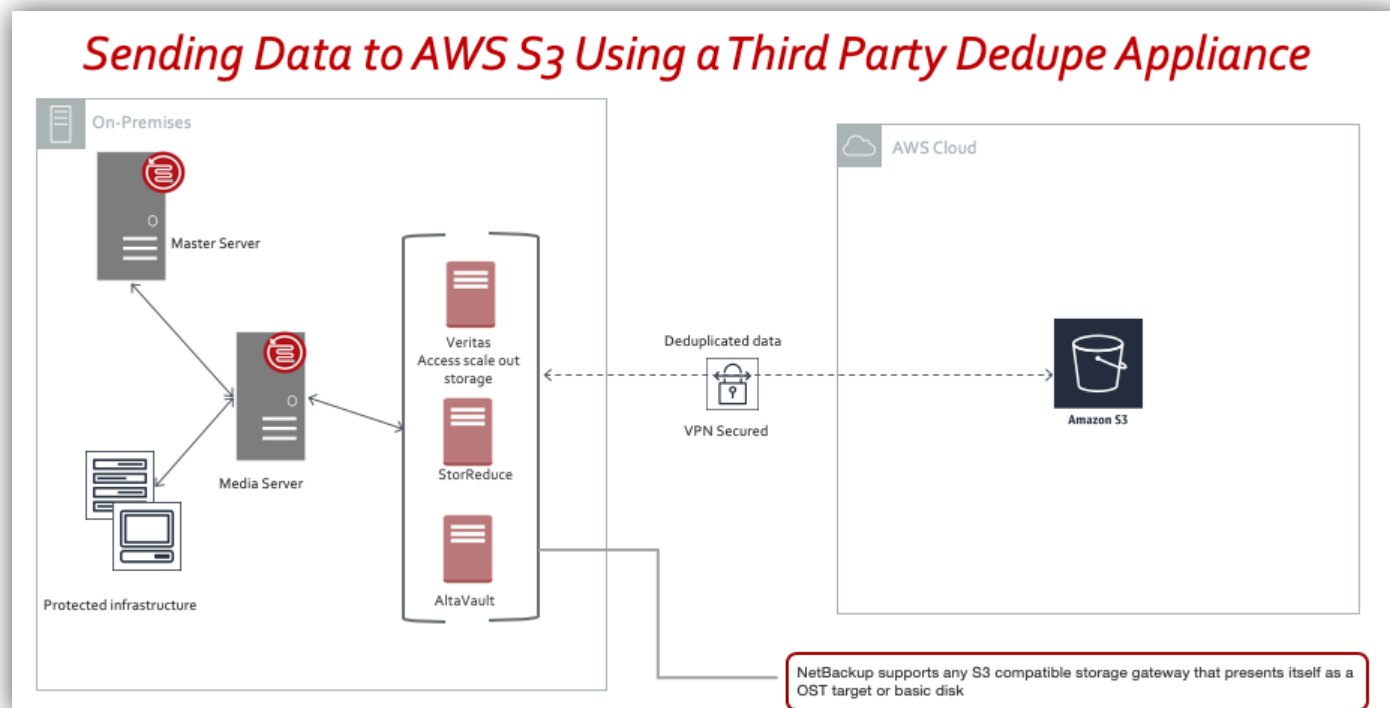


Figure 3 – Sending Data to AWS S3 Using a Third Party Dedupe Appliance

AMAZON STORAGE GATEWAY VTL

This use case emphasizes utilization of the Amazon Storage Gateway. This storage gateway uses iSCSI to provide what looks like a Virtual Tape Library to NetBackup. The ability to provide this access requires the downloading and installation of the Tape Gateway Virtual Machine from AWS. This VM will run in ESX or Hyper-V. It will then present an iSCSI connection to NetBackup. Local storage will be required on the VM to allow for data caching prior to the data being sent to AWS.

Once installed, the NetBackup Media Server will be able to see the VTL and configure it like it would a physical tape robot including drives and tape media with associated bar codes. Volumes (tapes) can be configured in various sizes ranging from 100GB to 2.5TB. The gateway can have up to 1500 virtual tapes with a maximum aggregate capacity of 1PB of storage. Backups are sent to the VTL in the same manner as backups that are sent to a physical tape library. The data is then compressed and sent to AWS across the network connection then stored in the cloud. NetBackup controls this archive process of moving the tape from S3 to Glacier (virtual tape shelf) by ejecting a tape. Therefore, restores from Glacier will require additional time due to the fact that there is a logical performance difference when reading from Glacier.

NetBackup will request the data from the VTL, and the VTL will need to pull the data from Glacier, back to the VTL, then across the network to NetBackup.

This use case is for customers of any size that wish to utilize what looks like a tape-based tracking solution. Many customers currently use VTL technology and with this AWS deployment, there are no changes required to be made to the existing infrastructure. This expands the VTL technology further in that the actual storage of the data is in the Cloud.

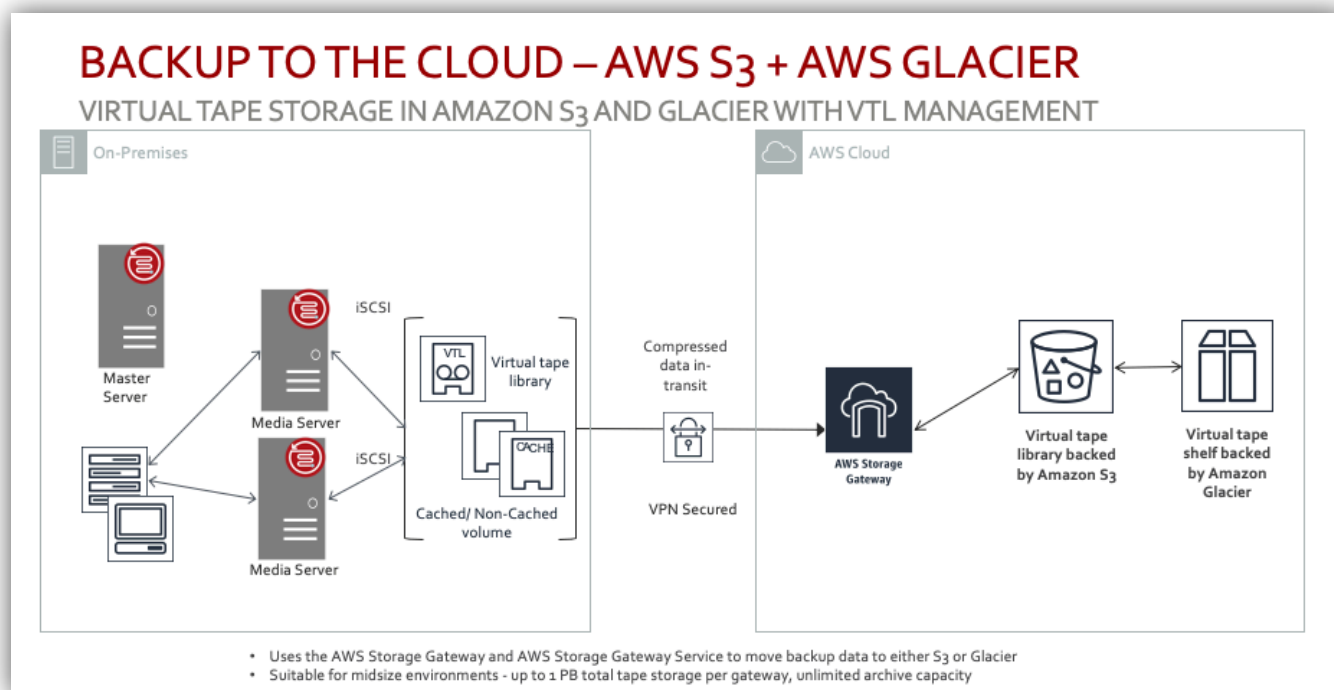


Figure 4 – Amazon Storage Gateway VTL

BACKUP IN THE CLOUD – AWS-ENABLED NETBACKUP ARCHITECTURES IN EC2

In addition to sending data to the cloud, developing a solution that is completely cloud native is also desirable. This is known as Infrastructure as a Service (or IaaS) and many customers are finding that maintaining data centers is not cost effective, so a cloud solution works well for them. This use case involves customer workloads in AWS Elastic Compute Cloud (EC2) which offers the ability to provision VM's similar to any other virtualized environment with the VM and all storage being in the AWS Cloud environment.

Backups of these workloads are typically still required. EC2 functions similar to a datacenter that uses a hypervisor environment for VM's. There are built in safeguards to protect data, however failure can still occur.

NetBackup in EC2 works exactly like NetBackup in a datacenter. A NetBackup master and media server can be provisioned from the marketplace using CloudFormation Templates (CFT) or manually deployed in a BYO fashion.

The image below (Figure 5) shows a media server running a Media Server Deduplication Pool (MSDP) which means that full or incremental backups will be deduped at the storage (EBS) layer. Alternately, S3 object storage, such as Standard, Standard-IA, Glacier or Glacier Deep Archive can be used as a storage target for NetBackup running in the cloud. For optimal storage cost savings, CloudCatalyst can be used to store duplicated data in S3 object storage. Each option has benefits depending on the need of the customer. EC2 instance types and sizing recommendations are covered towards the end of this document.

NetBackup offers an Amazon Machine Image (AMI) and a Cloud Formation Template (CFT) which is a template for the machine and makes provisioning of the Master and/or Media Server effortless.

NetBackup is available in the AWS Marketplace:

<https://aws.amazon.com/marketplace>

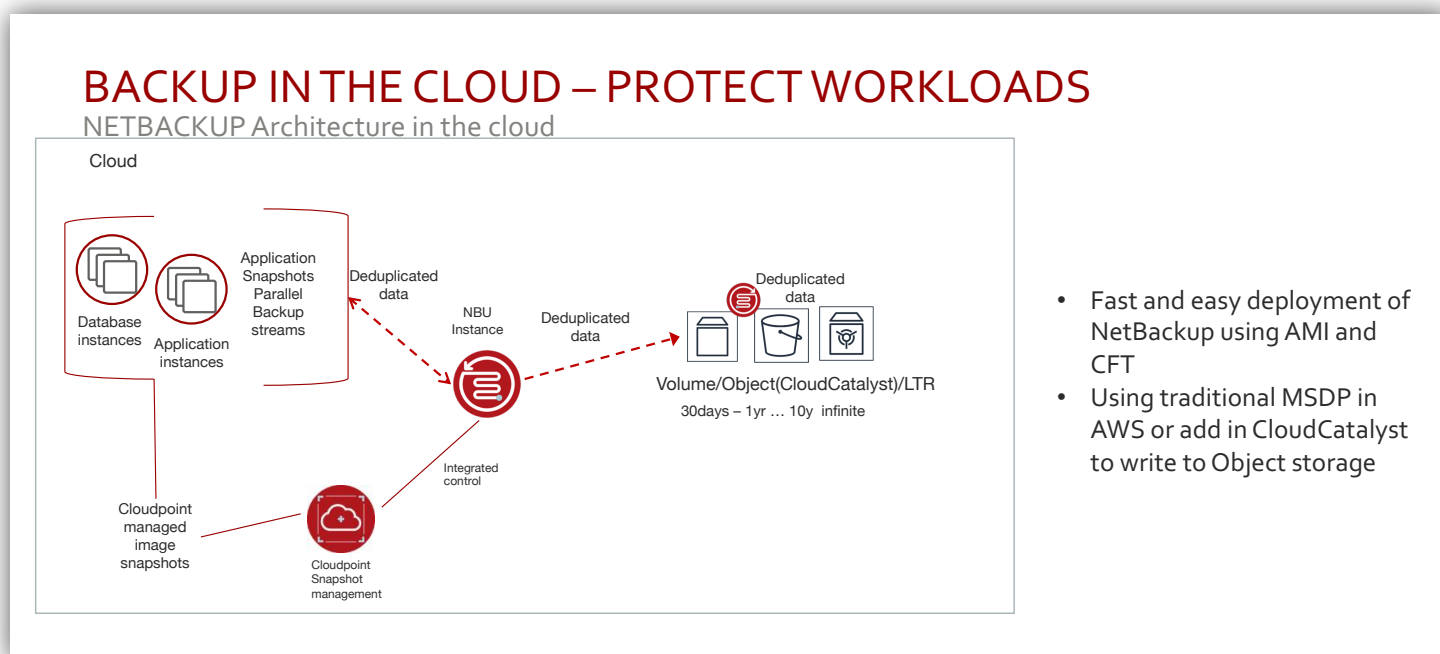


Figure 5 – Workloads in EC2

Similar to Figure 4 discussed earlier; NetBackup can work in EC2 with dedupe appliances that also work in the cloud. NetBackup treats them like “basic disk” exactly the same way it treats them in a data center. The appliance will dedupe the data before sending it on to S3 or to Glacier storage as outlined below.

There are many other configuration options using NetBackup with AWS that can be tailored to the customers’ needs. These use cases outline a handful of them.

DISASTER RECOVERY USING AWS

AUTO IMAGE REPLICATION (AIR) TO THE CLOUD – HYBRID CONFIGURATION

Another option to get data into the Cloud would be to utilize a Hybrid model where part of the environment is in the data center, and a secondary part, for use with Disaster Recovery options, would be to use the functionality of AIR to get the data to the cloud (Figure 5).

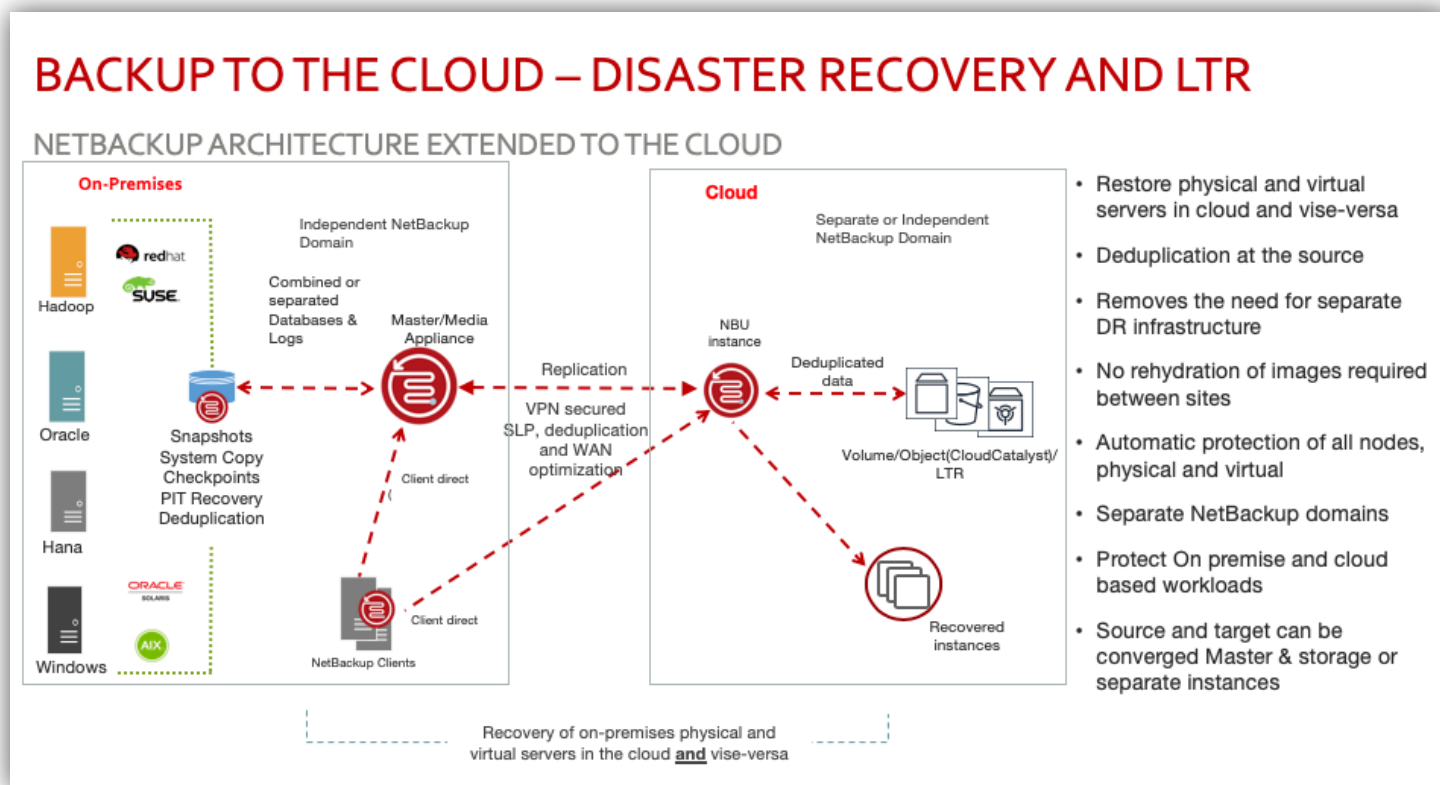


Figure 6 – NetBackup Auto-Image Replication for Cloud Recovery

This concept is very simple and ties into a number of these use cases. A NetBackup Master and Media Server with MSDP is configured in the data center, and a Master and Media Server with MSDP is configured in AWS. From there, an AIR process can be used to automatically send data from MSDP in data center A to MSDP in AWS. The data can then be imported into the AWS Master for use in a DR scenario. This is the same as using AIR to move data from a datacenter in San Francisco to a datacenter in London. The fact it is in the cloud is not noted by NetBackup. It is just an AIR target.

This option is ideal for a customer that would like to have an offsite DR copy of the data at a datacenter. It is also a very good way to migrate to the cloud from a NetBackup perspective.

Veritas offers additional products, such as Veritas Resiliency Platform (VRP), where workloads can be migrated into the cloud with the help of NetBackup. This method is a perfect blend of creating dual instances of a workload for test/dev/QA while maintaining the original data in the data center. In an existing NetBackup environment, VRP can offer orchestration workload recovery. Figure 7 provides an example of how NetBackup and VRP can be used to recover workloads into Amazon.

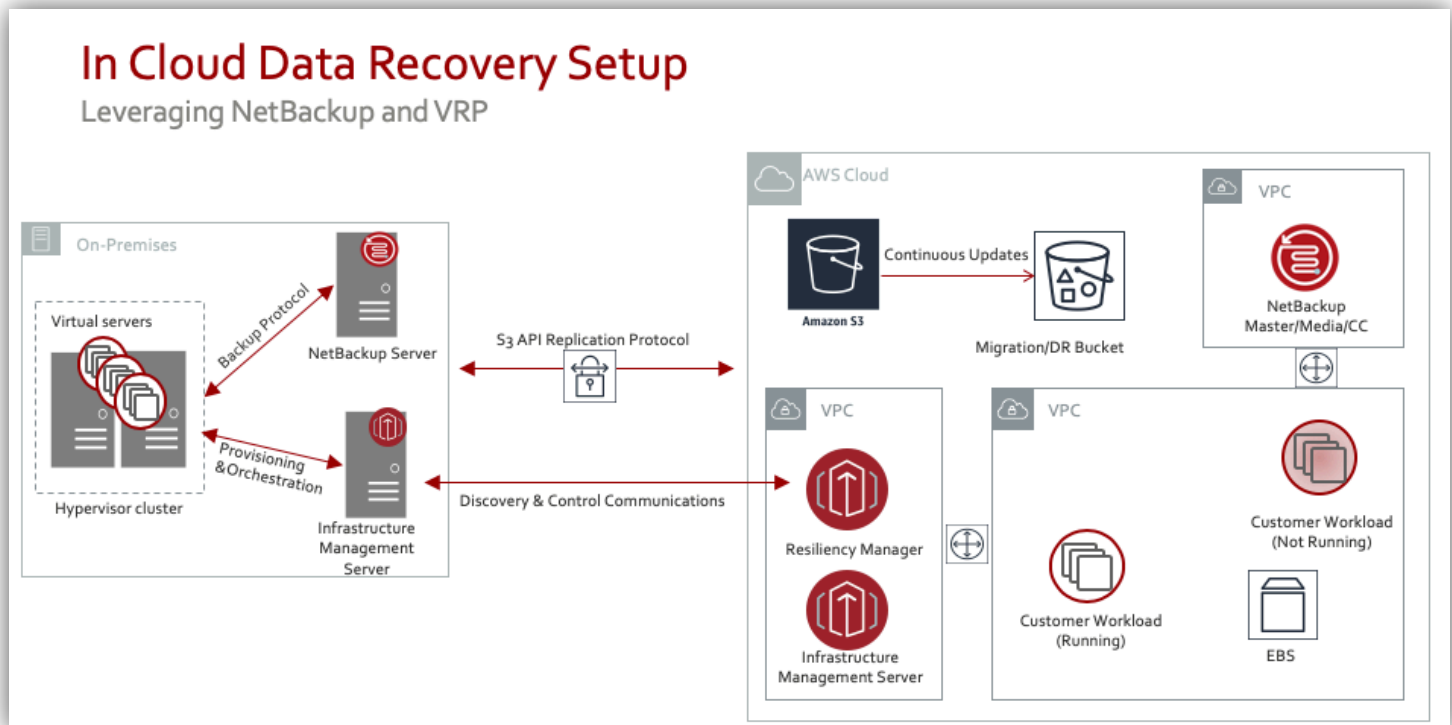


Figure 7 – Cloud DR with NetBackup and VRP

Leveraging NetBackup 8.2 CloudCatalyst Image sharing for Migration and DR

Starting with 8.2, NetBackup has added a new image sharing capability when using CloudCatalyst to write to AWS s3 object storage. This new functionality essentially makes the storage bucket self-descriptive for reuse by an instance other than the one that initially used it. The only data needed to access the data is the bucket name and the necessary authorization.

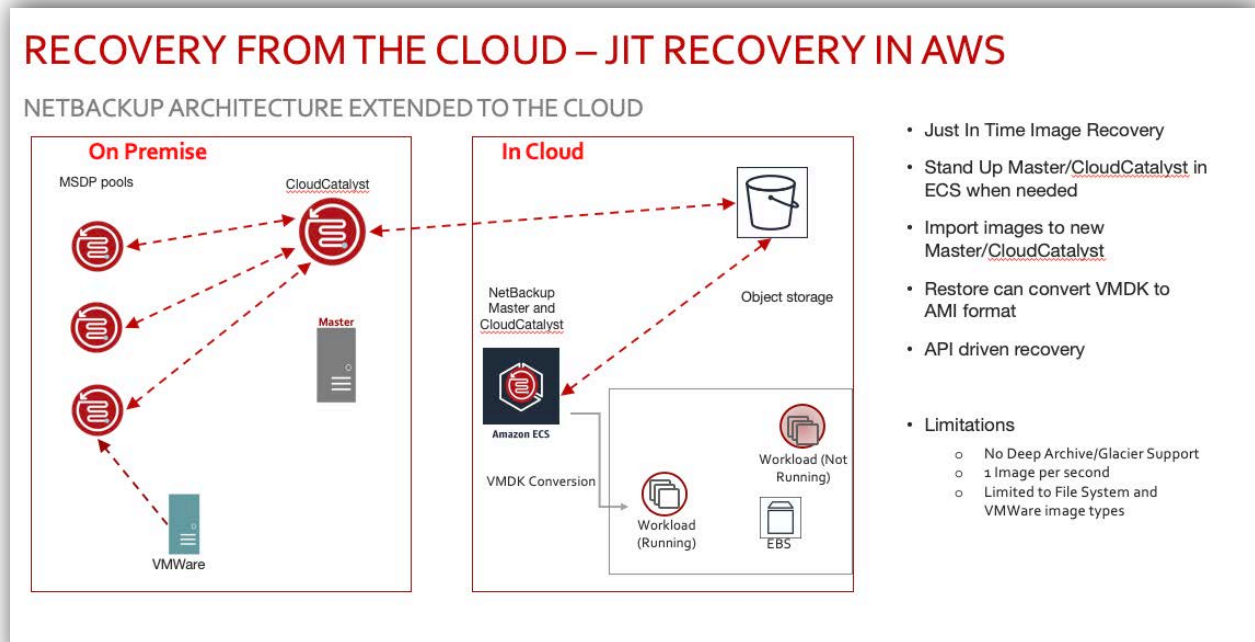


Figure 8 – 'Just in Time' Recovery in AWS

Using Image sharing, an on-demand NetBackup instance can be launched from the AWS marketplace CFT and attached to the existing CloudCatalyst bucket. The new instance will be able to read the bucket data from within the cloud infrastructure and leverage image data to restore workloads in the cloud, or even convert VMware images into an AMI format for host recovery and migration.

For more advanced migrations of complex environments and their infrastructures, the Veritas Resiliency Platform (VRP) integrates with NetBackup to orchestrate recovery and migration operations with pushbutton simplicity. This includes automatic deployment of NetBackup in AWS on demand instances to leverage CloudCatalyst data stores in object storage.

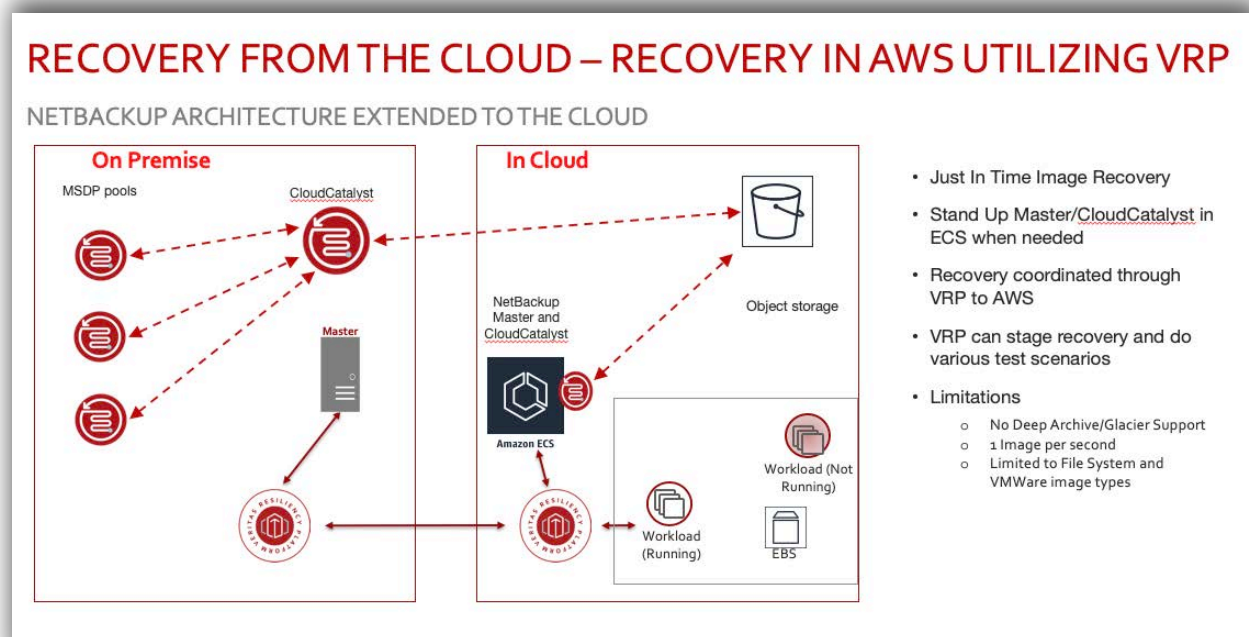


Figure 9 – Recovery in AWS utilizing VRP

CLOUD SIZING AND PERFORMANCE

Sizing and performance of data in the Cloud is based on customer need thus will vary from customer to customer. This makes providing exact sizing and performance information difficult.

To get data to the cloud, customers can utilize a simple Internet connection if the data to be transmitted is less than the amount of available bandwidth.

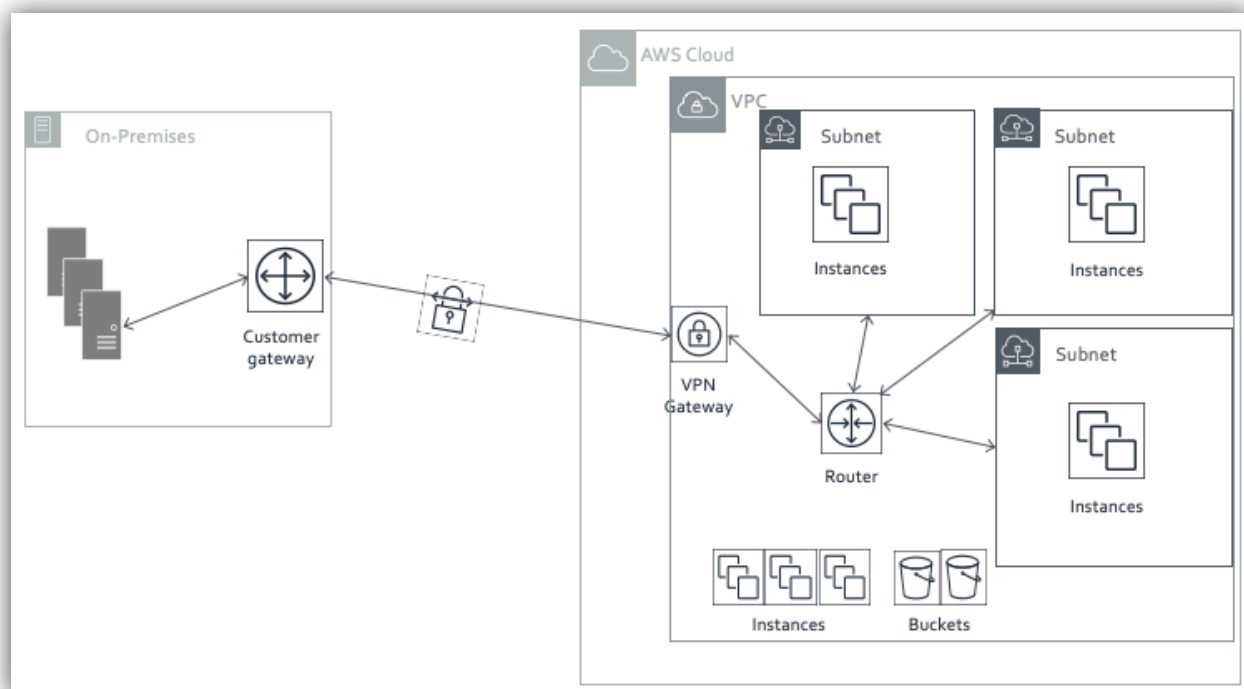


Figure 10 – A general view of cloud infrastructure connected with an on-premises datacenter

With AWS Direct Connect, customers can get a dedicated link to AWS with performance similar to LAN bandwidth inside a data center. Data can be compressed at the datacenter prior to being sent across the network to AWS or use CloudCatalyst to optimize the data being sent before it is stored in the cloud. Bandwidth can also be “throttled” if desired to prevent over-saturation of the network pipe.

AWS INSTANCE MODEL

Amazon Web Services has a Regional model. Various regions across the globe have been configured. Each Region has Availability Zones which are more or less datacenters within the Region that communicate with each other over high bandwidth connection similar to a customer having multiple physical datacenters in a geographical region that are close enough for low-latency connectivity, yet far enough to not be impacted by the same natural or man-made disaster.

Data within the Region will typically stay within the region but the option to select a geographically dispersed region is available for regional disaster recovery. Data can be replicated between Availability Zones to provide high availability within the cloud for the customer data. The loss of a single Availability Zone would not impact the operations of the others. Customers would choose to operate within the Region that is closest to

them (typically) to provide optimized bandwidth for moving data in and out of the AWS cloud, as well as have the option to select a geographically dispersed region to provide regional DR.

AWS STORAGE OPTIONS

One of the many benefits of the AWS storage model is the ability to quickly add storage to environments. Customers don't pay for the storage until it is provisioned. This model is much different from a traditional data center where racks of disk may sit idle until needed thus increasing TCO. If the disk is spinning and generating heat, additional cooling and power could be needed to keep the disk spinning even if it is not currently in use. Next-gen SSD arrays require less cooling and power, however idle disks still increase TCO.

Once data is in the cloud, AWS utilizes various types of storage including object (S3), object with infrequent access (S3-IA) and block (EBS) depending on the type of use case. Other options include EFS for scale out storage targeted at big data solutions and Glacier for long term storage of archive data that is rarely accessed. Sizing of the environment is based on the needs of the customer, and the workloads places in the cloud. Pricing is based on the type of storage chosen and is typically priced per GB. For example, standard S3 storage typically runs approximately \$0.023/GB per month whereas Glacier storage currently runs about US \$0.004/GB per month and even less for Glacier Deep Archive. Cost does depend on the region where the data is stored (see AWS [pricelist](#) for current costs). Glacier is typically used as long-term archive storage whereby data is moved there automatically using a variety of methods and a restore from the storage could take hours to access (vs. seconds for S3).

ENVIRONMENT DESCRIPTION AND ASSUMPTIONS FOR SIZING

The below sizing guidelines are based on the assumptions listed and were created using the standard NetBackup Appliance Calculator to determine the storage required for each type of workload. This is purely for AWS and backup in the cloud workloads only.

The following assumptions were used to size this environment:

- Data Assumptions
 - Data split – 80% FS / 20% DB [no hypervisor level in cloud]
 - Daily retention 2 week / Weekly – 4 weeks / Monthly 3 months
 - Daily change rate 2%, and YoY growth 10% [sizing done for 1 year only]
- Instance Type workload descriptions:
 - Small - FETB <=100TB <= 100 concurrent jobs

- Medium - FETB <=500TB <= 500 concurrent jobs
- Large - FETB <=1000 TB <= 1000 concurrent jobs
- Extra-large - FETB > 1PB >1000 concurrent jobs

NETBACKUP AWS INSTANCE SIZING

The architecture is based on a single NetBackup domain consisting of a NetBackup Master Server, several MSDP Media Servers and a single NetBackup CloudCatalyst Server in AWS EC2 Cloud.

Typically, backups are written directly to MSDP Storage for an immediate copy, then opt-duped to CloudCatalyst to send dedupe data to S3 object store. However, there is no requirement that backups must go to standard MSDP before CloudCatalyst. Requirements consist of:

- NetBackup Master Server
 - Single NetBackup Master Server can be on any supported Operating System.
- NetBackup MSDP Media Servers
 - MSDP Media Servers receive the initial backups from clients and performs deduplication.
- NetBackup CloudCatalyst Server
 - The CloudCatalyst Server is dedicated to performing NetBackup dedupe writes to S3 Object store. It is a dedicated high-end RedHat server that meets the minimum requirements for CloudCatalyst. It takes the deduplicated backups images from the MSDP media servers and stores them in S3.
- Backup Workloads (Clients/Agents)
 - These are the systems or applications that are being protected.

NetBackup Master Server

The NetBackup Master Server should be sized according to the standard Veritas guidelines depending on the load placed upon the NetBackup domain as a whole. Plan accordingly for the initial needs of the environment. AWS does have added benefits of being able to scale up the systems as workloads grow. The solution can scale out by adding additional media server nodes.

Master Server memory and CPU requirement

The table below details the minimum processor and memory requirements for the various environment sizes.

<i>Number of processors</i>	<i>Minimum RAM</i>	<i>Maximum number of jobs per day</i>	<i>Maximum number of media servers per master server</i>
4	16GB	10000	20
8	32GB	20000	50
16	64GB	30000	100

These estimates are based on the number of media servers and the number of jobs the master server must support. The amount of RAM and number of processors may need to be increased based on other site-specific factors.

Master Server recommendations

Environment Size / Components	Small	Medium	Large	Extra Large
Master Server	32 GiB / 8 vCPUs Storage: EBS – 500GB EBS – Catalog 5GB Sample instance: m4.2xlarge	64 GiB / 8 vCPUs Storage: EBS – 500GB EBS – Catalog 5GB Sample instance: r3.2x.large	64 GiB / 16 vCPUs Storage: EBS – 500GB EBS – Catalog 10GB Sample instance: m4.4xlarge	122 GiB / 16 vCPUs storage: EBS – 500GB EBS – Catalog 10GB Sample instance: r4.4xlarge

NetBackup MSDP storage

NetBackup MSDP Storage can reside on either a NetBackup Appliance, a Virtual Appliance, or a BYO virtual or physical host, including a cloud based virtual instance. This section will outline MSDP in AWS built on an EC2 instance with EBS storage.

Specifications for MSDP Media Server in EC2

The host computer's CPU and memory constrain how many jobs can run concurrently. The storage server requires sufficient capability for deduplication and for storage management. Processors for deduplication should have a high clock rate and high floating-point performance. Furthermore, high throughput per core is desirable. Each backup stream uses a separate core.

<i>Hardware Component</i>	<i>MSDP Media Server</i>
CPU	<ul style="list-style-type: none"> Veritas recommends at least a 2.2-GHz clock rate. A 64-bit processor is required. At least four cores are required. Veritas recommends eight cores. For 64 TBs of storage, Intel x86-64 architecture requires eight cores.
RAM	<ul style="list-style-type: none"> From 8 TBs to 32 TBs of storage, Veritas recommends 1GB of dedicated RAM for 1TB of storage. However, beyond 32 TBs storage, Veritas recommends more than 32GBs of RAM for better and enhanced performance.
Operating System	<ul style="list-style-type: none"> The operating system must be a supported 64-bit operating system. See the operating system compatibility list at http://www.netbackup.com/compatibility

NetBackup CloudCatalyst

The CloudCatalyst storage server is a dedicated Linux media server for MSDP deduplicated cloud storage. If a BYO server is used in AWS, specifications are shown below.

Specifications for a Linux media server

The dedicated media server that will be configured as a CloudCatalyst storage server should meet or exceed the minimum system specifications of a small NetBackup CloudCatalyst. The requirements for the CloudCatalyst media server are smaller than for a regular MSDP media server and are noted below and can be found at:

https://www.veritas.com/content/support/en_US/doc/NB_CC_MIN_SYS_REQ

- Red Hat Enterprise Linux 7.3 or later
- NetBackup 8.1 or later
- 4 Cores Minimal (8 preferred)
- 16GB RAM Minimal (32GB preferred)
- 1TB of gp2 EBS cache volume

Growing the Media Server

As the amount of data protected by a server increases, the load requirements on that host will increase. In that scenario, there is a simple solution. Any EC2 instance or EBS volume can easily be expanded to meet higher requirements that may happen over time.

Refer to the following for more information:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html>

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modify-volume.html>

Media Server Deduplication Pool recommendations

Running traditional MSDP in a cloud environment requires specific resources to be available such as 10G network, EBS volumes with Provisioned IOPS, etc. The recommendations below have been formulated using AWS kits that addresses MSDP pools of different sizes. These are just recommendations and specific customer environments may have different needs. Depending on the AWS footprint, any of the environments below would work based on the sizes.

MSDP considerations

- Example MSDP storage pool size is up to 96TB on Linux
 - Can be a direct backup target, use Fingerprinting Media Servers or a Client-Side Deduplication target
 - MSDP will be storing all data in EBS volumes
 - The pool will be able to replicate to any Veritas Deduplication compatible target, including CloudCatalyst

Below are recommended NetBackup Media server sizing guidelines based on the size of the intended deduplication pool:

Dedupe Pool	Instance Type	Storage	Cores	RAM	Networking	IOPS
10 TB (Small)	R3.2xlarge	1x160 SSD 1x16TB EBS SSD	8	61		
1-20 TB (Small)	C4.8xlarge	1x80 EBS-SSD 1x16 TB EBS-SS	36	60	10 GB	EBS Provisioned IOPS (SSD)
32 TB (Medium)	C4.4xlarge	1x80 EBS-SSD 2x16 TB EBS-SSD	16	30	10 GB	
	R3.2xlarge	1x160 SSD 2x16TB EBS-SSD IOPS - 12,000	8	61		12,000
32-64 TB (Large)	m4.10xlarge	1x80 EBS-SSD 2-4x16 TB EBS-SSD	40	160	10 Gb	
	c4.8xlarge	1x80 EBS-SSD 2-4x 16 TB EBS-SS	36	60	10 Gb	
	r3.4xlarge	1x160 SSD 2x16TB EBS-SSD IOPS - 12,000	8	61	10 Gb	12,000
32-96TB (xLarge)	m4.10xlarge	1x80 EBS - SSD 2-4x16TB EBS-SSD	40	160	10 Gb	
	r3.8xlarge	2x320 SSD 2-6x16 TB EBS-SSD IOPS - 12,000	32	144	10 Gb	12,000

* Instance info at a glance: https://www.ec2instances.info/?min_memory=64&min_vcpus=16

Product	Role	Instance Type	EBS/EFS Storage	CPUs	RAM (GB)
NBU	CloudCatalyst Min	RHEL M5.xlarge	250 GB SSD (gp2) 1TB Cache	4	16 GB
	CloudCatalyst Large	RHEL M5.2xlarge	500 GB SSD (gp2) 1TB Cache	8	32 GB

For NetBackup CloudCatalyst servers, the above table lists a minimum configuration and a large configuration along with the EC2 instance type and the storage configuration. Customers should expect to start with the larger instance recommendation, unless using CloudCatalyst for basic functionality testing. The SSD disk listed below will contain the operating system and NetBackup installation files. The 1TB volume represents the local cache volume required for CloudCatalyst cloud deployments (Figure 11).

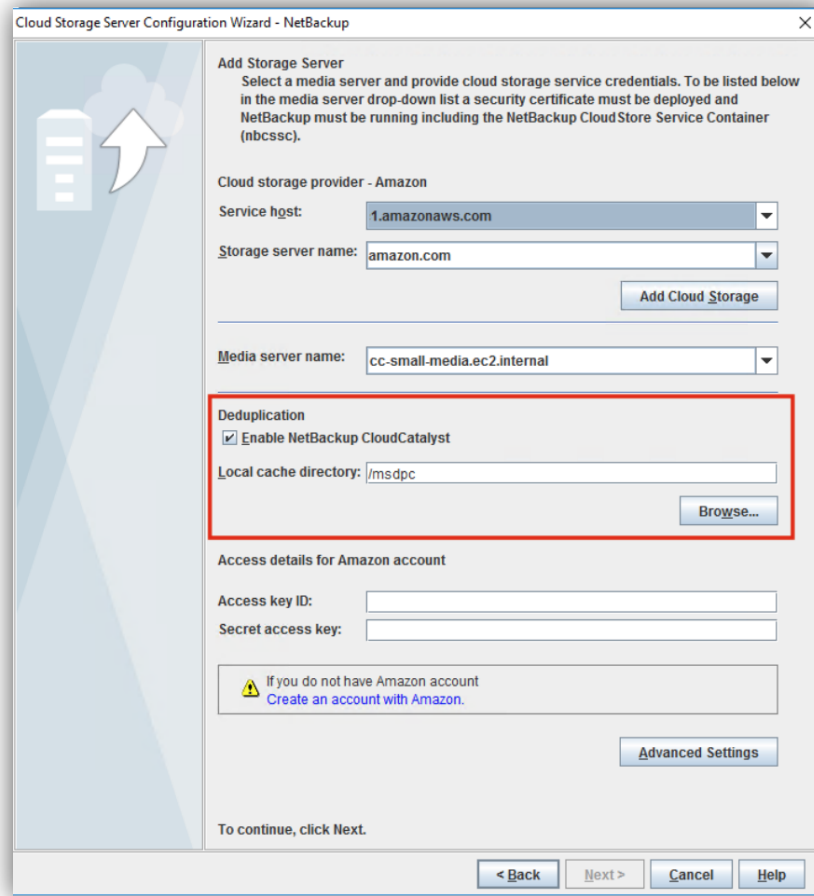


Figure 11 – Configuring CloudCatalyst

ADDITIONAL ARCHITECTURE REQUIREMENTS

In addition to the use case architectures noted in this document, there are several other topics that customers will need to consider when looking at a move or partial move to the cloud.

SECURITY OF THE INFORMATION

In-flight

Starting with NetBackup 8.1, data security has been heightened as more data is now going to the cloud and out of the ownership of the data center. With NetBackup, the use of SSL and certificates guarantees that the servers and clients that are being protected and the data being received are from authenticated endpoints.

At-rest

NetBackup MSDP can deliver source side encryption from the client end that encrypts data in transit and at rest. In addition, any data from that client that is sent to another pool will maintain that encryption, even when going to the cloud via CloudCatalyst.

When using standard cloud storage servers, data coming from NetBackup moving into the cloud can utilize encryption before the data is sent to the AWS environment from the media server. This encryption can be done at the Client or Media Server level and uses Key Management from the NetBackup master to handle the keys (KMS). The data in the cloud at rest will be encrypted. The only drawback to this option is that during a restore the KMS server must be available in order to have the keys available to decrypt the data. In most cases this would not be an issue unless the original Master is not available. Starting with 8.2 CloudCatalyst can use additional KMS based key protection as well.

Least privileged access

For security concerns, it's always important to practice a "least privileged" approach. This means that users are limited to only be able to access what they absolutely need. When running NetBackup on AWS, or in the cloud, the minimum permissions needed for NBU to write and retrieve data from S3 are:

- s3:CreateBucket
- s3:ListAllMyBuckets
- s3:ListBucket
- s3:GetBucketLocation
- s3:GetObject
- s3:PutObject
- s3>DeleteObject

For Amazon Glacier, you need following additional permissions:

- s3:PutLifecycleConfiguration
- s3:GetLifecycleConfiguration
- s3:PutObjectTagging

For more details, refer to this following Knowledge Base article:

[How to configure Amazon Web Services when using I AM Role with NetBackup](#)

These permissions should not be simply assigned to a user directly. They should be applied in a limited context using Resource Tags and prefixes of S3 bucket names.

Limit access with Resource Tags

Resource tags are a recommended way to manage access to resources. For example, when working with S3 buckets, NetBackup can be limited to the above permissions for specific buckets using a defined resource tag that only grants the above permissions when working with the specified buckets.

For example, in the following policy the permissions indicated above are only granted to the user when accessing buckets that have names beginning with the string "nbu".

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:PutObject",
        "s3:GetObject",
        "s3:CreateBucket",
        "s3:ListBucket",
        "s3:DeleteObject",
        "s3:GetBucketLocation",
        "s3:ListAllMyBuckets"
      ],
      "Resource": [
        "arn:aws:s3::nbu*"
      ]
    }
  ]
}
```

For more info, see details from Amazon on s3 [access tags](#), [user specific access](#), [folder access](#) and [object tagging](#).

About AWS Service Quotas

Every AWS account has default Service Quotas. These used to be referred to as service limits. These limits implement constraints on the number of RDS Database clusters, a limit on the number of a specific instance type can be running at once, or how many elastic IPs the user is allowed.

For more information on Service Quotas, review Amazon documentation here: [AWS Service Quotas](#).

AWS Service Quotas and NetBackup

Since NetBackup is running on dedicated EC2 instances and VPC that have predefined architectures, it is not likely to be impacted directly by service quotas. A specific limit that NetBackup could see an issue with is the max number of EBS snapshots (default 100,000). If this limit was hit some snapshot-based backup operations will fail with a snap create failure. To avoid this problem, the user will need to request a quota increase in the AWS [Service Quotas console](#).

Disaster recovery scenarios

Disasters can strike anywhere and at any time, even in cloud infrastructures. That's why it's always important to have a plan for how to recover from failures. Part of an effective plan is to have proper recovery testing. With NetBackup, there are two levels of disaster recovery. One is to recover from an EC2 based snapshot of the NetBackup server, another is to perform DR with the NetBackup's built in Disaster Recovery capabilities. Using NetBackup's built in recovery capabilities is preferred.

To test recovery from a snapshot, do the following:

- Stop NetBackup Services to ensure consistency or the image will only be crash consistent
- Perform EC2 snapshots using the EC2 console, AWS cli or APIs.
- Recover to a new instance in an isolated, dedicated VPC from the snapshot
- The VPC should prevent connectivity to production resources during testing
- Test functionality to validate the recovery was successful.
- Discard the test instance(s).

The Preferred method of recovery NetBackup is to use the built-in recovery capabilities. This involves a Catalog Backup Policy, an isolated DR backup location, such as S3 or an EFS share, and a recovery wizard that can be used to recover in the following scenarios:

- Corruption of the database
- Loss of important configuration information
- Rollback after an upgrade failure or other significant change
- Recovery from complete loss of the system

The recovery wizard will walk the user through the recovery process. Depending on the location of the catalog backup, there might be some configuration required to access the DR data, such as mounting the EFS share, or Complete details on Disaster recovery and NetBackup in available in the NetBackup Troubleshooting Guide's [Disaster Recovery](#) section. If using MSDP or CloudCatalyst, be sure to enable the MSDP DR policy, details are [here](#).

COST OVERVIEW

The cost of the cloud will vary depending on what is needed. Simple backups to S3 storage can be a very cost-effective solution for a customer that wants to send important data to an offsite location. This solution is probably not ideal for a large customer with a large amount of data sending to S3 due to the bandwidth constraints of the network. In addition, S3 storage is limited in options based on object vs. block-based storage. There is a cost to send the data, store the data and retrieve the data.

The cost of Gets and Puts

When writing data to Amazon S3 object storage, there is a cost associated for each time you or an application uploads/updates a file or object (PUT) and when you retrieve an object (GET) from Amazon S3. To optimize the data transfer to S3, NetBackup breaks data down to 64MB of deduplicated data before sending it to the configured S3 bucket. Each 64MB chunk write or read will incur a GET or PUT request.

<https://aws.amazon.com/s3/pricing/>

Storage costs

Just as in the above example of costs associated with putting objects into Amazon S3, there is also a cost of using S3 storage. The costs and available S3 storage types will vary based on region, so be sure to check prices in your intended region when calculating costs: <https://aws.amazon.com/s3/pricing/>

Compute costs

EC2 environments whereby machines are configured in the cloud using EBS disk will have additional cost based on the number of processors needed, RAM usage, amount of disk provisioned, etc. Backups in this environment will also incur costs based on moving the data from the client to the NetBackup environment. This cost is dependent on the location of the NetBackup media server in relation to the source client or data. Most options in cloud-based computing come “a la-carte” whereby the customer pays for what they use. In some cases, these costs can be less than maintaining a data center, in some cases the cost is more. That said, the peace of mind knowing the data is highly available and at a relatively secure offsite datacenter might be worth the extra expense of the cloud environment.

To better understand the cost structure, AWS has created a number of calculators that can be used to determine approximate cost options. For example, the TCO calculator can be found at this link:

<https://aws.amazon.com/tco-calculator/>

This allows the customer to put in information about a planned deployment and receive cost estimates. Another comprehensive calculator to determine basic cost models can be found here:

<https://calculator.s3.amazonaws.com/index.html>

When you deploy an Amazon EC2 Instance the cost of the instance is determined by the hardware type (CPU, memory) EBS volume usage, and utilization. For example, to run a small NetBackup instance to protect less than 100TB for front end data, it would cost about \$301 (*compute of \$250 and EBS storage of \$50.50*) a month with NetBackup running 16 hours a day. If NetBackup does not need to run for 16 hours, then the total cost decreases.

As noted, cost should not always be the determining factor when it comes to a cloud-based solution. There is more to a datacenter than hard costs. The AWS infrastructure can provide additional options that might be more expensive up front, but the flexibility provided with on-demand storage, uptime guarantees, and staffing may make a move to AWS more sense. And NetBackup will be there to protect the data that is in the cloud.

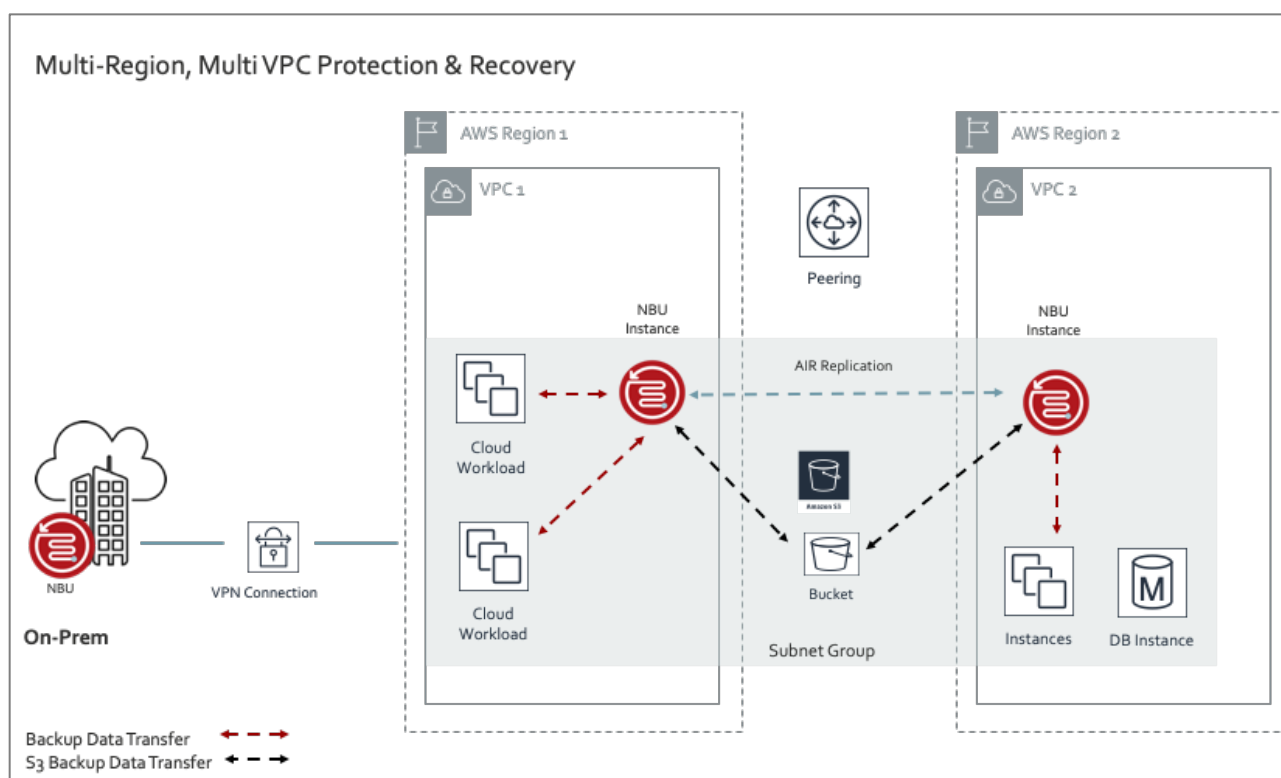
DEPLOYMENT DETAILS

NETBACKUP VPC DEPLOYMENT CONFIGURATIONS

Veritas NetBackup utilizes a 3-tier architecture that consists of a single Master Server, multiple Media Servers, and clients which allows flexibility when deploying NetBackup in AWS. By default, NetBackup supports Single-AZ, Multi-AZ as well as Multi-Region environments.

NetBackup deployment options

To optimize data movement as well as costs, it is recommended that you place the NetBackup Master server in the region where management of all EC2 instances takes place, and in a subnet where it can communicate with the deployed media servers. Media servers should be placed as close to the target EC2 instances as possible, but as long as network connectivity exists, a media server in one AZ can protect clients in different AZs.



Where multi-region protection is required there are two deployment options; a NetBackup domain in each region or single NetBackup domain with media servers in each region. By utilizing a NetBackup domain for each region, traffic will be isolated in each region and NetBackup AIR can be used to offer disaster recovery in case of a region failure. The downside to this approach is that each NetBackup domain will need to be managed independently.

Utilizing a single NetBackup domain and deploying Media servers into each region eliminates the need to manage multiple NetBackup domains independently but if there is an outage or failure, all backup operations are impacted for all regions.

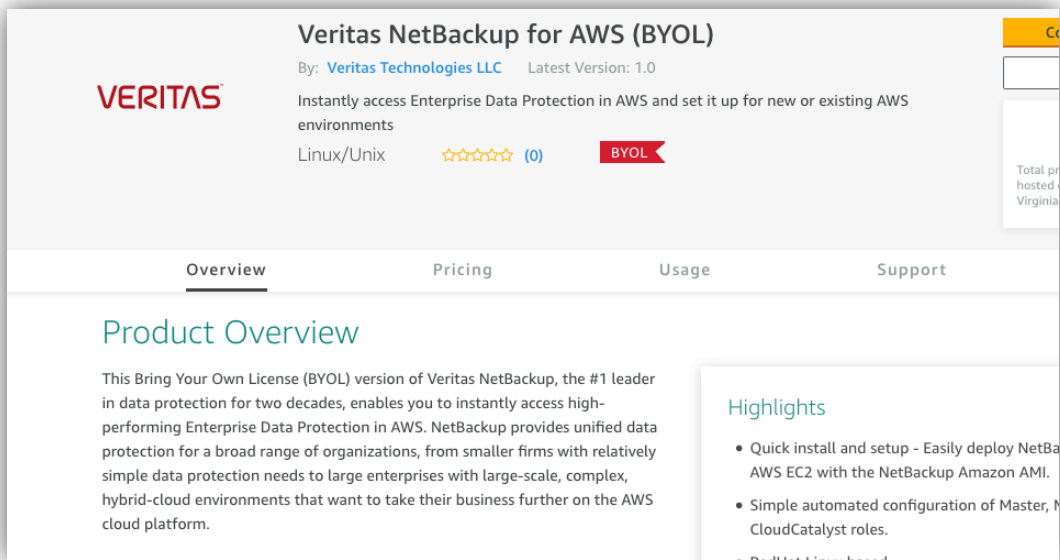
SETTING UP NETBACKUP CLOUDCATALYST IN AWS

Note: This document assumes that you already have set up a NetBackup 8.1 or later master server instance in EC2. Your environment may vary.

Using the AWS marketplace

NetBackup is available on the AWS Marketplace (<https://aws.amazon.com/marketplace>).

The NetBackup page is available [here](#).



From the AWS Marketplace, NetBackup can be deployed as a master or media server using a straightforward CloudFormation Template (CFT) form.

Application Parameters

NetBackupRole
Install NetBackup as Master or Media server, or for Cloud Recovery

Master

NBUMasterServerName
Use only lowercase(a-z),digits(0-9),minus sign(-) and period(.) for NetBackup master server name

netbackupmaster

NBUMediaServerName
Use only lowercase(a-z),digits(0-9),minus sign(-) and period(.) for NetBackup media server name (not required for Cloud Recovery)

netbackupmedia

LicenseKey
NetBackup License Key

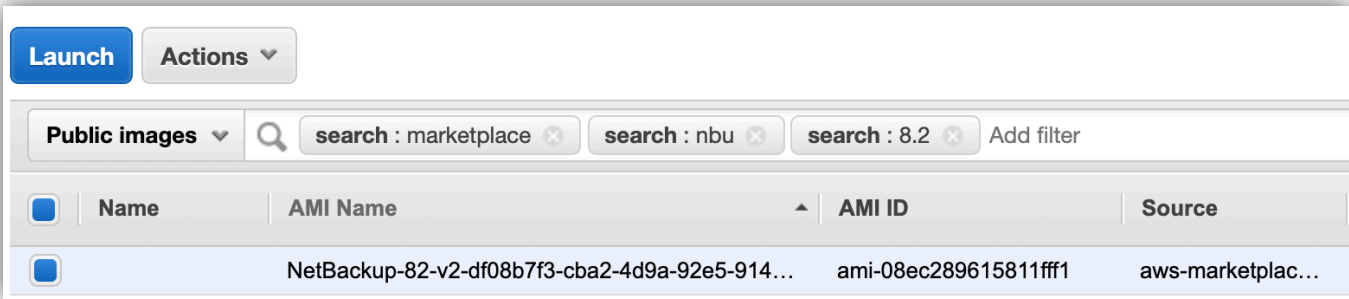
CCBucketName
Name of the bucket where CloudCatalyst images are stored (use only when installing Cloud Recovery)

S3 Bucket used by CloudCatalyst

The CFT will help deploy NetBackup with an existing or a new VPC, including setting up the necessary Security Group requirements for secure communication.

Launching an Elastic Compute Cloud (Amazon EC2) instance

Alternately, launch an EC2 instance manually from an AMI from the **Ec2 -> Images -> AMIs** dialog.



- 1. On the **Choose an Instance Type** page, select an appropriate size for your performance requirements.

<input type="checkbox"/>	General purpose	m4.xlarge	4	16
<input checked="" type="checkbox"/>	General purpose	m4.2xlarge	8	32
<input type="checkbox"/>	General purpose	m4.4xlarge	16	64

For more information about specifications, review the following document:
https://www.veritas.com/content/support/en_US/doc/NB_CC_MIN_SYS_REQ

Note that this is a minimum system requirement. Based on your objective, a larger system may be required. Remember the instance can easily be resized after deployment. In AWS, a m4.2xlarge is a good system to start with.

- 2. Click **Next: Configure Instance Details**.
- 3. On the **Configure Instance Details** page:
 - a. Under **Network**, select the appropriate VPC.
 - b. Enable **Auto-assign Public IP** to access the instance from your desktop.
- 4. Click **Next: Add Storage**.
- 5. On the **Add Storage** page, select the storage size and volume type. Select an appropriate EBS volume size for the root disk. It's recommended a second disk be added for the CloudCatalyst Cache.

The cache size should be a minimum of 1TB for production usage.

Volume Type ⓘ	Device ⓘ	Snapshot ⓘ	Size (GiB) ⓘ	Volume Type ⓘ	IOPS ⓘ	Throughput (MB/s) ⓘ
Root	/dev/sda1	snap-0fcae4949a75c5312	100	General Purpose SSD (gp2) ⌵	300 / 3000	N/A
EBS ⌵	/dev/sdb ⌵	Search (case-insensit	1024	General Purpose SSD (gp2) ⌵	3072	N/A

- 6. Click **Next: Add Tags**.
- 7. On the **Add Tags** page, add tags if desired. However, tags are not required.

1. Choose AMI2. Choose Instance Type3. Configure Instance4. Add Storage5. Add Tags6. Configure Security Group7. Review

Step 5: Add Tags

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. A copy of a tag can be applied to volumes, instances or both. Tags will be applied to all instances and volumes. [Learn more](#) about tagging your Amazon EC2 resources.

Key (127 characters maximum)

Value (255 characters maximum)

Instances ⓘ

Volumes ⓘ

This resource currently has no tags

Choose the Add tag button or [click to add a Name tag](#).
Make sure your [IAM policy](#) includes permissions to create tags.

Add Tag (Up to 50 tags maximum)

Cancel

Previous

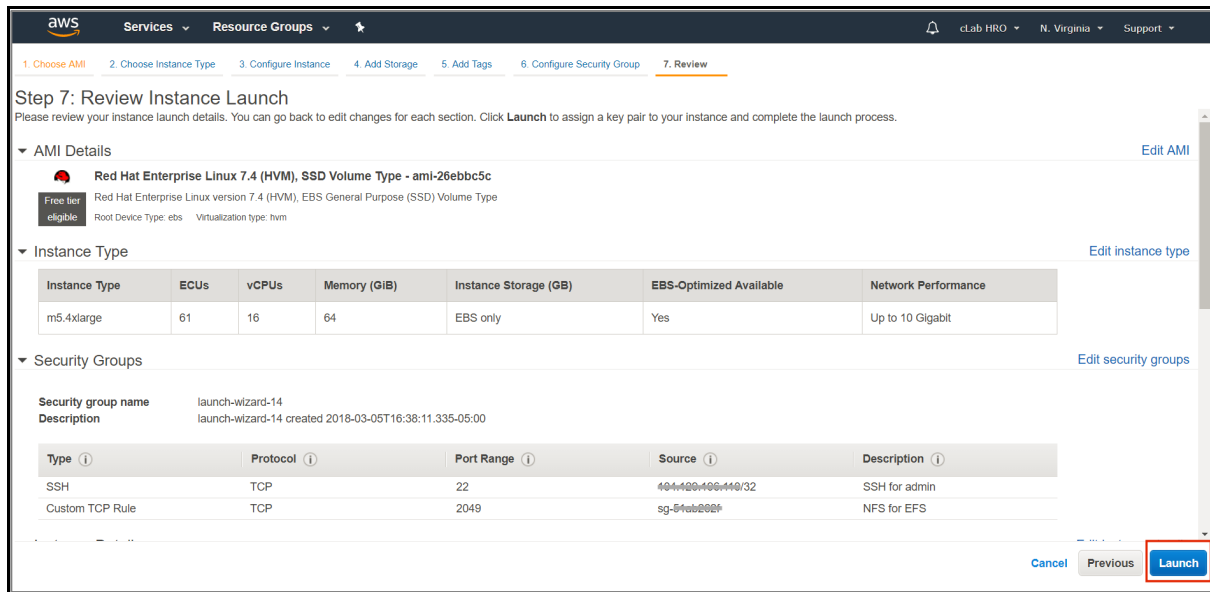
Review and Launch

Next: Configure Security Group

- 8. Click **Next: Configure Security Group**.
- 9. On the **Configure Security Group**.
In this example, we've allowed access from sources outside the VPC to NetBackup services.

Type ⓘ	Protocol ⓘ	Port Range ⓘ	
Custom TCP Rule	TCP	10102	6
Custom TCP Rule	TCP	1556	6
Custom TCP Rule	TCP	1556	0
Custom TCP Rule	TCP	1556	::
SSH	TCP	22	0
SSH	TCP	22	::
Custom TCP Rule	TCP	10086	6
Custom TCP Rule	TCP	5637	6
RDP	TCP	3389	0
RDP	TCP	3389	::
HTTPS	TCP	443	0
HTTPS	TCP	443	::
Custom TCP Rule	TCP	10082	6

- 10. Click **Review and Launch**.
- 11. On the **Review Instance Launch** page, review the EC2 instance, and then click **Launch**.
The instance is available in approximately 15 minutes.



12. Once available, change the host name of the machine to a fully qualified domain name (FQDN) using the following procedure:

<https://aws.amazon.com/premiumsupport/knowledge-center/linux-static-hostname-rhel7-centos7/>

13. When using the second volume for the CloudCatalyst cache, be sure to configure it for use prior to configuring CloudCatalyst:

[Making an Amazon EBS Volume Available for Use on Linux](#)

Configuring CloudCatalyst on the EC2 instance

If using the CFT, NetBackup should be up and running, ready to be configured. If manually deploying from the AMI, NetBackup can be deployed for use by connecting to the host by ssh as directed in the "Connect to your instance" dialog in the EC2 interface. After connecting, run the NetBackup installer in: **/root/NBUSetup/**.

During the CloudCatalyst configuration, as part of the Cloud Storage Server Configuration Wizard, enter the EBS mount point as the CloudCatalyst Local cache directory. In the following example, the EBS mount point is: **/msdpc**.

Follow the configuration wizard prompts, selecting the desired bucket type such as S3 Standard or Glacier. Only supported bucket types will be selectable. The wizard GUI is already in the doc earlier (Figure 11).

PROTECTING NETBACKUP ACCESS WITH EC2 SECURITY GROUPS

Security and data protection are still a requirement as data is migrated to the cloud. When deploying NetBackup into AWS there are two ways to limit network connectivity and access to the NetBackup instance—Security Groups or VPC access control lists.

An EC2 A security group acts as a virtual firewall for your instance to control inbound and outbound traffic. These security group firewall settings are attached to the EC2 network interface to only allow connections to NBU and out.

To limit the inbound connectivity to NetBackup, create a security group to control access to the EC2 instance or a Network ACL to control access to the subnet with the following custom TCP settings:

VPC Dashboard

Filter by VPC:

Select a VPC

Virtual Private Cloud

Your VPCs

Subnets

Route Tables

Internet Gateways

Egress Only Internet Gateways

DHCP Options Sets

Elastic IPs

Endpoints

Endpoint Services

NAT Gateways

Peering Connections

Security

Network ACLs

Security Groups

Type ⓘ	Protocol ⓘ	Port Range ⓘ
Custom TCP Rule	TCP	13724
Custom TCP Rule	TCP	1557
Custom TCP Rule	TCP	13705
Custom TCP Rule	TCP	8443
Custom TCP Rule	TCP	13713
Custom TCP Rule	TCP	13717
Custom TCP Rule	TCP	13716
Custom TCP Rule	TCP	18782
Custom TCP Rule	TCP	13783
Custom TCP Rule	TCP	13711
Custom TCP Rule	TCP	13702
Custom TCP Rule	TCP	13701
Custom TCP Rule	TCP	13785
Custom TCP Rule	TCP	18720

These inbound security group settings will allow all NetBackup Master Server, Media Server, and Clients to communicate with each other on the network.

TAGGING NETBACKUP RESOURCES

Tagging your AWS resources is very important process as it allows for more precise filtering, searching, and reporting. In relation to data protection, tagging all NetBackup resources will help identify all resources needed to provide NetBackup data protection to the targeted AWS infrastructure.

There are two ways to tag NetBackup resources—when launching an EC2 instance or post deployment.

Tagging during launch

1. Using the EC2 Launch Instance Wizard, under **Add Tags** specify the Key of **Name** and value of the desired NetBackup server name.

Make sure **Instance** and **Volumes** are check marked.

Step 5: Add Tags

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. A copy of a tag can be applied to volumes, instances or both. Tags will be applied to all instances and volumes. [Learn more](#) about tagging your Amazon EC2 resources.

Key	Value	Instances	Volumes
Name	NetBackup Master Server	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

2. If deploying NetBackup using the CloudFormation Template, the template will automatically create a tag for the Instance name.

Application Parameters

NetBackupRole
Install NetBackup as Master or Media server, or for Cloud Recovery

Master

NBUMasterServerName
Use only lowercase(a-z),digits(0-9),minus sign(-) and period(.) for NetBackup master server name

NBUMediaServerName
Use only lowercase(a-z),digits(0-9),minus sign(-) and period(.) for NetBackup media server name (not required for Cloud Recovery)

3. If additional Tags are needed, they can be created under **Configure stack options**.

Configure stack options

Tags

You can specify tags (key-value pairs) to apply to resources in your stack. You can add up to 50 unique tags for each stack. [Learn more](#)

Key

Value

Remove

Add tag

Tagging post deployment

Adding a tag to NetBackup resources already in use can be done at any time.

- 1. Select the NetBackup EC2 resource that requires a tag.
- 2. Select the Tags tab and select **Add/Edit Tags**.

VPCInstance IDInstance Type

i-02e390bd91b5a47...

m5a.4xlarge

Instance: i-02e390bd91b5a47a5Public DNS: ec2-3-81-35

DescriptionStatus ChecksMonitoringTags

Add/Edit Tags

- 3. Specify a **Key** (such as name) and **Value** for this tag.
You can add up to 50 tags per resource. You can hide or show the column for each tag as well.

Add/Edit Tags [X]

Apply tags to your resources to help organize and identify them.

A tag consists of a case-sensitive key-value pair. For example, you could define a tag with key = Name and value = Webserver. [Learn more](#) about tagging your Amazon EC2 resources.

Key	Value	
host name	rwmaster	[X] Show Column
name	NetBackup master	[X] Show Column

Create Tag Cancel Save

- Once saved, the Name column will now display the name provided as well as any additional tags created.

	Name	VPC	host name
<input checked="" type="checkbox"/>	NetBackup master		rwmaster

ROTATING AWS ACCESS KEYS FOR CLOUDCATALYST

NetBackup utilizes a CloudCatalyst server to send deduplicated data to AWS S3 buckets. CloudCatalyst utilizes IAM roles and access keys to make API calls to the AWS S3 services. If AWS access keys are configured, they are stored encrypted with the CloudCatalyst server. If security policies dictate that access keys be changed, each CloudCatalyst server configured to use the access key will need to be updated.

To update the server with a new key requires the use of the `tpconfig` command:

- Connect to the CloudCatalyst server using `ssh`.
- Under `/usr/opensv/volmgr/bin` run `tpconfig -dsh -all_hosts` to show the current access key user and the storage server type for Amazon.

```
Storage Server:      amazon.com
User Id:             AKIAT3NJE7RRGNT04XZG

Storage Server Type: PureDisk_amazon_rawd
```

- Run the following command:
`tpconfig -update -storage_server server name -stype server_type -sts_user_id AWS Access Key -password AWS Secret Key`
- Once completed the output of `tpconfig -dsh -all_hosts` will reflect the new access key. Once this change is made, NetBackup will utilize the new access and secret key to write to the configured S3 bucket.

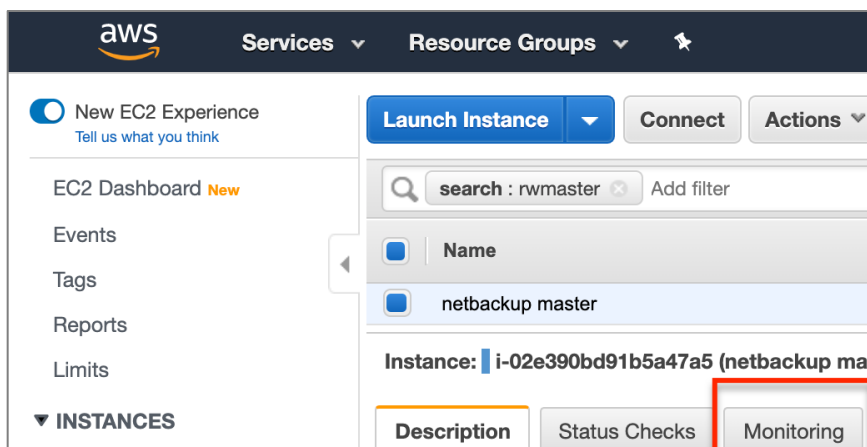
PROTECTING NETBACKUP FROM FAULTS, FAILURES, AND DOWNTIME

Veritas NetBackup utilizes AWS EC2 Service to provide data protection in an AWS VPC. This means that if there is an Availability Zone or Region failure impacting the EC2 service, NetBackup will be affected. Now those type of failures are uncommon, but network connectivity can be interrupted, EC2 instances can go offline, EBS volumes can become corrupted, so choosing the right deployment model that optimizes disaster recovery (see section on disaster recovery) is important to plan for potential failure.

Using alarms and monitoring for the NetBackup environment, the long-term risk of downtime or failure from AZ, instance, or application faults can be detected and avoided. CloudWatch alarms can also be used to alert users when data written to an S3 bucket exceeds a defined threshold. To accomplish this, create CloudWatch alarms for System Status and Instance Status Check failures for NetBackup Server EC2 instances.

To create these alarms, perform the following steps:

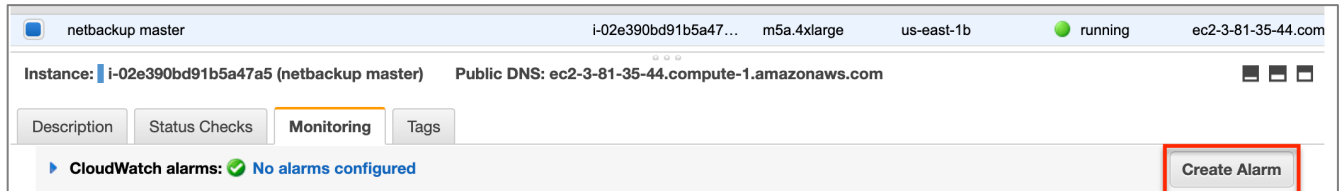
- Select the NetBackup EC2 Instance and then select the **Monitoring** tab.



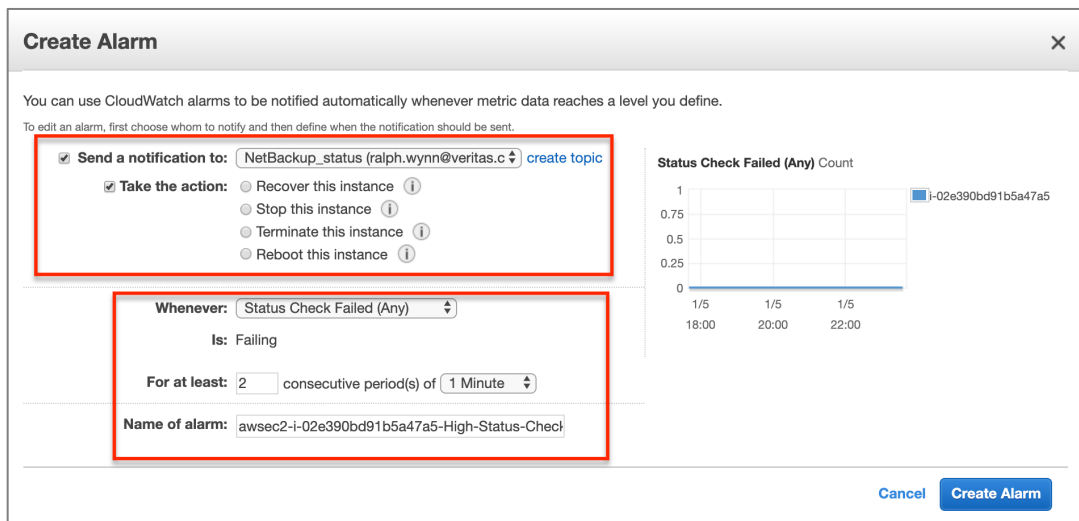
- Click **Enable Detailed Monitoring** so that metrics will be captured every 1 minute.

CloudWatch metrics: Basic monitoring. **Enable Detailed Monitoring**

- On the monitoring table, click **Create Alarm**.



- For the alarm, specify (or create a new topic) to receive alerts of status check failures.
- Select **Take the action** and then select **Reboot this instance** to enable this action whenever any status checks fail.



- If the NetBackup EC2 instance has a hardware (AWS infrastructure) or a software failure (OS, memory), AWS sends a notification and performs the configured action.
- Additional alarms monitoring network packets in and out as well as write operations to attached EBS volumes can be used to detect possible network or storage failures and provide remediation steps.
- Creating a CloudWatch alarm for the S3 Metric of BucketSizeBytes will send a notification when data written to a S3 bucket exceeds a threshold.

NETBACKUP RISK AND AUDIT MANAGEMENT

Veritas NetBackup has built in role-based access control (RBAC) that can limit and control what actions users can take when performing backup, restore, as well as adding clients.

Information on utilizing NetBackup RBAC can be found here:

https://www.veritas.com/content/support/en_US/doc/135031700-135031704-0/index

If additional insight is needed into what API calls NetBackup used to communicate with AWS S3 or what API calls NetBackup is making in the target VPC, the use of CloudTrail can provide the right level of insight. CloudTrail can be used to find out which account deleted a S3 bucket.

Enabling CloudTrail logging for NetBackup resources

To enable CloudTrail logging for NetBackup resources:

1. Create a CloudTrail that captures all create and delete API calls and provides Log Insights.

Create Trail

Trail name*

Apply trail to all regions ☐ Yes ☒ No
Creates the trail in this region and delivers log files for this region

Management events

Management events are records of actions that are performed on or within resources in your AWS account. These are also known as control plane operations. [Learn more](#)

Read/Write events ☒ All ☐ Read-only ☐ Write-only ☐ None ⓘ

Log AWS KMS events ☒ Yes ☐ No ⓘ

Insights events

Insights events are records that capture an unusual call volume of **write management APIs** in your AWS account. Additional **charges** apply. [Learn more](#)

Log Insights events ☒ Yes ☐ No

2. Specify either all buckets or the buckets used by NetBackup and the target S3 bucket for collections.

The screenshot shows the AWS IAM console interface for configuring S3 logging. The 'S3' tab is active. A table lists the S3 buckets to be logged, with one entry: 'rw-forntbu' with a prefix of '/ Prefix (optional)' and permissions for both Read and Write. Below the table, the 'Storage location' section is visible, where 'Create a new S3 bucket' is set to 'No' and the 'S3 bucket*' is set to 'cloudtrailforntbackup'.

- Once created CloudTrail will log all API calls to the target bucket.
It's recommended to utilize AWS Athena to query the CloudTrail logs for actions by the NetBackup IAM user or Access key.

More information on using Athena with CloudTrail can be found here:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/cloudtrail-request-identification.html>

AWS Scheduled Service events

In rare events Amazon can schedule service events for instances. These events can include things like restarting or even retiring an instance. In the event of this scheduled activity, Amazon will notify the controlling user via email ahead of the event with additional details. It is important to make sure that the email address in the account is up to date or there could be an unexpected interruption.

When dealing with AWS scheduled service events, follow the instructions from Amazon under [Scheduled Events for your Instances](#).

SUMMARY

Customers are moving to the cloud and a number of cloud providers are moving to the forefront of the cloud megatrend. Amazon Web Services and Veritas have teamed up to create a usable, scalable solution for customers who want a cloud presence. There are multiple paths to the cloud which means that proper planning and research is required to make sure the path taken will yield the expected outcome. In this document, the most common cloud use cases that customers are deploying have been called out. By following some the guidelines called out in this guide for these desired use cases, your cloud journey should be successful.

APPENDIX A – ADDITIONAL INFORMATION

Description	Link
Veritas Information	http://www.veritas.com
NetBackup 8.2 Cloud Administrators Guide	https://www.veritas.com/content/support/en_US/doc/58500769-135186602-0/index
Veritas NetBackup™ Deduplication Guide : v8.2	https://www.veritas.com/content/support/en_US/doc/25074086-136046435-0/index
AWS Marketplace	https://aws.amazon.com/marketplace
Setting up NetBackup CloudCatalyst in AWS	https://www.veritas.com/content/support/en_US/doc/CC_AWS_guide
NetBackup Security and Encryption Guide	https://www.veritas.com/support/en_US/doc/21733320-132525226-0/index
Disaster Recovery for NetBackup CloudCatalyst : v8.2	https://www.veritas.com/content/support/en_US/doc/cloudcatalyst_dr_82
How to configure Amazon Web Services when using I AM Role with NetBackup	https://www.veritas.com/support/en_US/article.100044514
NetBackup CloudCatalyst Technical Whitepaper	https://www.veritas.com/content/dam/Veritas/docs/white-papers/NBA-3.1-Technical-White-Paper-5240-CloudCatalyst-2017-10.pdf
AWS Service Quotas	https://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html
Protecting Amazon S3 objects using NetBackup with s3fs - Whitepaper : v8.2	https://www.veritas.com/content/support/en_US/doc/Amazon_S3_NetBackup_s3fs
How to Resize an EC2 instance	https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html
How to modify an EBS volume	https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-modify-volume.html
NetBackup Deduplication Guide	https://www.veritas.com/support/en_US/doc/25074086-127355784-0/index

AWS Regions and Availability Zones	http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-regions-availability-zones
AWS EC2 Additional Information	https://aws.amazon.com/ec2/
AWS S3 Additional Information	https://aws.amazon.com/s3/
AWS Glacier Storage Additional Information	https://aws.amazon.com/glacier/details/

APPENDIX B – TERMINOLOGY

Amazon Elastic Block Store (Amazon EBS): A storage file system that provides persistent block storage volumes that can be used with Amazon EC2 instances in the AWS Cloud. *However, do not use this storage for CloudCatalyst operations in Amazon EC2.*

Amazon Elastic Compute Cloud (Amazon EC2): A web service that provides secure, resizable compute capacity in the cloud.

Amazon Machine Image (AMI): A file that provides the information necessary to launch an instance (or virtual server) in the Amazon cloud.

Amazon CloudFormation Template (CFT): A CFT allows for automation of deployment of services in AWS

NetBackup Extendable Storage File System (NetBackup ESFS): The NetBackup CloudCatalyst database used for CloudCatalyst operations. CloudCatalyst uses the NetBackup Extendable Storage File System Service (vxesfsd) and its subcomponents to move and manage files in the local cache directory and the cloud.

DISCLAIMER

THIS PUBLICATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS PUBLICATION. THE INFORMATION CONTAINED HEREIN IS SUBJECT TO CHANGE WITHOUT NOTICE.

No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

ABOUT VERITAS TECHNOLOGIES LLC

Veritas Technologies empowers businesses of all sizes to discover the truth in information—their most important digital asset. Using the Veritas platform, customers can accelerate their digital transformation and solve pressing IT and business challenges including multi-cloud data management, data protection, storage optimization, compliance readiness and workload portability—with no cloud vendor lock-in. Eighty-six percent of Fortune 500 companies rely on Veritas today to reveal data insights that drive competitive advantage. Learn more at <http://www.veritas.com>/or follow us on Twitter at [@veritastechllc](https://twitter.com/veritastechllc).



Veritas World
Headquarters

2625 Augustine Drive
Santa Clara, CA 95054

(866) 837-4827

For specific country offices
and contact numbers,
please visit our website.