



NetBackup in the Public Cloud

Guidelines for Azure Deployments

This technical paper is designed to assist partners and end users looking to protect workloads deployed in Azure cloud using Netbackup. These guidelines will aid design and implementation of 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.

For more information on Veritas products and solutions, please visit www.veritas.com.

VERITAS™

The truth in information.

TABLE OF CONTENTS

INTRODUCTION 3

NETBACKUP OVERVIEW 3

BUSINESS VALUE 4

WHY ARE CUSTOMERS LEVERAGING THE CLOUD?..... 4

NETBACKUP AND Azure DESIGN OVERVIEW..... 5

Azure CLOUD VERSUS ON PREMISES CONSIDERATIONS 6

USE CASE OVERVIEW..... 6

SENDING DATA FROM ON-PREMISES TO AZURE USING A THIRD-PARTY GATEWAY APPLIANCE 8

BACKUP IN THE CLOUD – AZURE ENABLED NETBACKUP ARCHITECTURES 9

DISASTER RECOVERY USING AZURE 11

CLOUD SIZING AND PERFORMANCE 12

ADDITIONAL ARCHITECTURE REQUIREMENTS..... 20

At-rest..... 20

Storage Costs 21

Compute Costs 21

APPENDIX A - ADDITIONAL INFORMATION 23

INTRODUCTION

The purpose of this technical paper is to provide a technical reference on the capabilities of NetBackup and Microsoft Azure. 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 Microsoft Azure to offer robust backup to and in the cloud. 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.

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

Many Veritas customers are considering Azure cloud as a supplemental data center, a hybrid of both on-prem and cloud, or as a means of completely eliminating the traditional data center. These changes in business model require new strategies to migrate and protect data and workloads. The extensive value of Veritas solutions goes beyond seamlessly protecting the data regardless of the location, to orchestrating the movement of workloads to the cloud. Whether it is a disaster recovery requirement, or the desire to eliminate physical data center management, 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 not having to maintain a datacenter and an expensive DR site. Mid-size customers enjoy having an offsite copy of their data that is built on highly scalable hardware or leverages just-in-time cloud recovery. Large customers with datacenters are identifying workloads that can take advantage of cloud availability and cost while freeing expensive datacenter space for mission critical workloads. Sometimes a customer will need a temporary space for a workload and instead of ramping up a new rack of disks in a datacenter, they temporarily leverage space at a cloud provider to avoid the additional cost of purchased datacenter hardware. Cloud subscription models works very well for these sorts of projects with highly scalable, simple to use models.

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

Cloud also addresses the aspects of hardware maintenance and updates. New firmware for arrays, for example, is required on a regular basis causing risk and downtime to install. Similarly replacing or upgrading hardware impacts the environment as a customer would need to manage these in the datacenter. Alternatively, in the cloud these requirements are taken care of by the cloud provider and are invisible to the customer.

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

NETBACKUP AND Azure DESIGN OVERVIEW

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

NetBackup has created an Azure Virtual Machine (VM) option that allows for quick provisioning of NetBackup instances. The NetBackup Azure VM can then be used with Page Blob storage. Azure Page Blob Storage (Managed Disks) is used as the disk storage for these machines which uses block-based storage. In a Page Blob storage instance, the Azure VM template can be selected to deploy a NetBackup Master or Media server with customized CPU, RAM, disk and IP address as needed by the customer. This provides the necessary storage requirements for NetBackup's storage optimization features like traditional MSDP, Advanced Disk and CloudCatalyst.

This methodology is simply a "virtual environment" similar to the other virtual environments that can be created using on-prem virtualization vendors. Virtual machines are spun up as needed and managed like physical machines. They can also be used within an Azure Resource Manager (ARM) Template that uses a JSON scripting system to automate NetBackup instance creation. Using JSON, deployment of NetBackup, either large or small, has been simplified to filling out a short deployment form. The NetBackup ARM template is available in the Azure Marketplace for easy deployment.

NETBACKUP AND CLOUD CONNECTIVITY

NetBackup can utilize Azure in several different ways depending on the needs of the customer. As outlined in various places in this document NetBackup can utilize Block Blob Storage Hot, Cool or Archive tiers like a regular disk pool or utilize the NetBackup CloudCatalyst to efficiently send deduplicated data to Block Blob Storage. If a customer has resources in the cloud, NetBackup installed in the cloud protects these resources in a similar manner to protecting physical resources in a datacenter. This avoids the cost and performance impact of traversing data back to the datacenter for backups.

Cloud instances are managed using a web-based UI and systems in the Cloud can be further monitored and managed using other connectivity methods such as RDP.

NETBACKUP AND CLOUD RESTORE OPTIONS

Restores of data in the cloud are as simple as in a local data center. The backup admin has full use of the UI to recover data. The selected type of storage (Hot, Cool or Archive) affects the restore performance, Hot tier recovery within the cloud being similar to that of on-prem recovery in a datacenter. A section at the end of this document goes over the restore process of basic Azure storage and includes screen shots. It is beyond the scope of this document to detail all restore scenarios.

Azure 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 Premises:** If you have specific IOPs requirements you can select the appropriate hardware to meet those needs.
 - **Azure Cloud:** You can make your selection of Page Blob volumes based on IOPs requirements. There is a cost associated with guaranteed IOPs.
- **Peer Link Limits:**
 - **On Premises:** You can have as many peer-to-peer links as required.
 - **Azure Cloud:** There are fixed limits on the number of VNET to VNET peer links that are allowed.
- **Storage targets**
 - **On Premises:** Systems typically write to block storage, deduplication devices or MSDP pools.
 - **Azure Cloud:** Storage targets are typically Page Blob/Managed Disk or Block Blob. Page Blob is generally more expensive while Block Blob is at a lower cost and more scalable.

USE CASE OVERVIEW

There are several different use cases with NetBackup and Azure, several of which 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 FROM ON-PREMISES TO AZURE BLOCK BLOB 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 object storage target such as available in Microsoft Azure or Microsoft Azure GovCloud

BACKUP TO THE CLOUD – STANDARD OBJECT STORAGE

NETBACKUP TO AZURE CLOUD

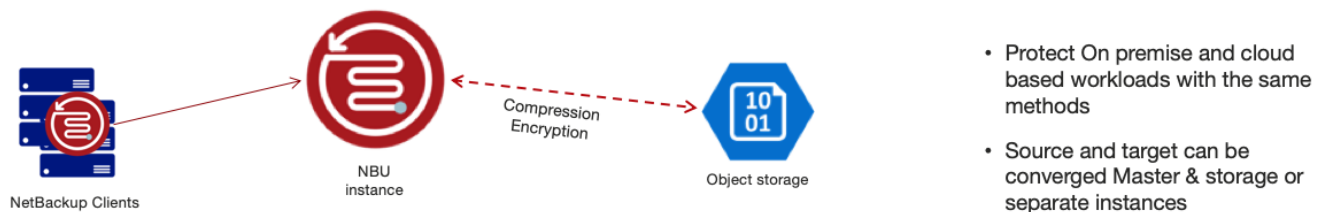


Figure 1- Using Standard object storage

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 Azure Storage pricing page:

<https://azure.microsoft.com/en-us/pricing/details/storage/>

BACKUP FROM ON-PREMISES TO AZURE BLOCK BLOB STORAGE WITH CLOUDCATALYST DEDUPLICATION

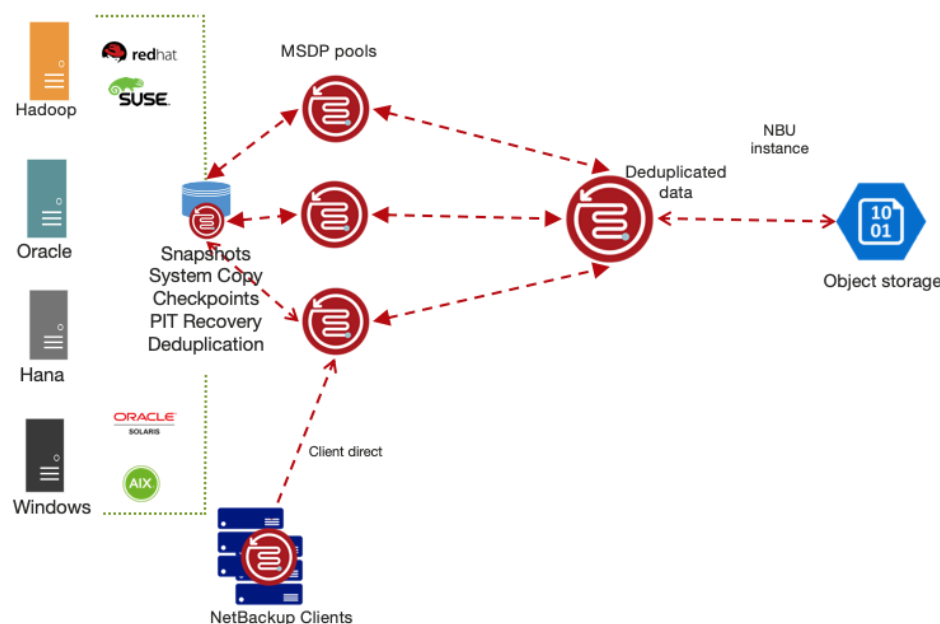
The NetBackup CloudCatalyst solution combines the performance and flexibility of NetBackup with powerful data deduplication technology to better leverage the cloud for storing backups for disaster recovery or long-term data retention. By ensuring backup data remains optimized while in transit to the 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 as a purpose-built appliance, a virtual appliance, or as a build-your-own (BYO) software solution. CloudCatalyst allows customers send backup data to cloud object storage in deduplicated form. CloudCatalyst can process optimized backup images from existing MSDP volumes or directly from a client for transfer to an Azure Block Blob Storage target. The hot and cool tiers of Azure Block Blob Storage have been certified for use with NetBackup CloudCatalyst as well.

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.

BACKUP TO THE CLOUD – CLOUDCATALYST TO OBJECT STORAGE

NETBACKUP ARCHITECTURE EXTENDED TO THE CLOUD



- Deduplication at the source
- No rehydration of images required between pools
- Automatic protection of all nodes, physical and virtual
- Protect On premise and cloud based workloads with the same methods
- Source and target can be converged Master & storage or separate instances

Figure 2- Sending Data to Azure Using Cloud Catalyst

SENDING DATA FROM ON-PREMISES TO AZURE USING A THIRD-PARTY GATEWAY APPLIANCE

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 performs the deduplication before the changed blocks are forwarded on to the cloud via Azure Storage APIs. This solution will work with any Azure 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 which requires the data to be rehydrated prior to sending to the gateway device.

Sending Data to Azure Using a Third Party Dedupe Appliance

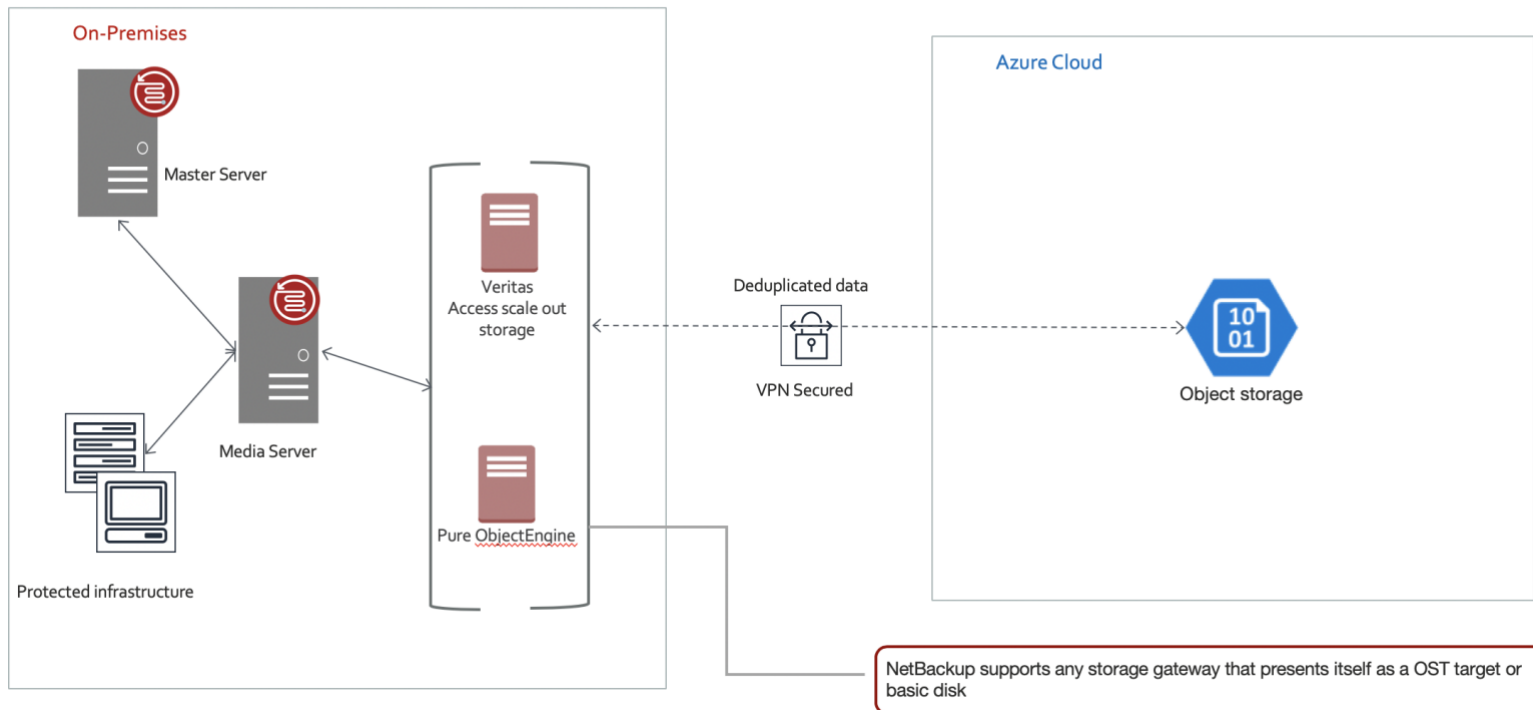


Figure 3- Sending Data to Azure Using a Third Party Dedupe Appliance

BACKUP IN THE CLOUD – AZURE ENABLED NETBACKUP ARCHITECTURES

In addition to sending data to the cloud for DR or tape elimination, 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 running workloads entirely in the cloud is more cost effective and offer the ability to provision VMs with the VM and all storage being in Azure. When protecting Azure based workloads, it is an important cost consideration to minimize data movement to On-Prem by running NetBackup within the cloud as well.

Backups of these workloads are still required to protect from corruption and malicious activity such as ransomware. Azure VMs function similar to a datacenter using a hypervisor environment for VMs so there are built in safeguards to improve data availability, however failure and corruption can still occur. NetBackup in

Azure IaaS works exactly like NetBackup in a datacenter. A NetBackup master and media server can be provisioned from the marketplace using Azure Resource Manager (ARM) or manually deployed in a BYO fashion.

The image below (figure 4) shows a media server running a Media Server Deduplication Pool (MSDP) so that full and incremental backups will be deduped in the Azure storage (Page Blob/Managed Disks) layer. Alternately, Block Blob Storage, such as Premium, Hot, Cool or Archive can be used as a storage target for NetBackup. For optimal storage cost savings, CloudCatalyst can be used to store duplicated data in Block Blob Storage. Each option has benefits depending on the need of the customer and the data type. Azure VM instance types and sizing recommendations are covered later in this document.

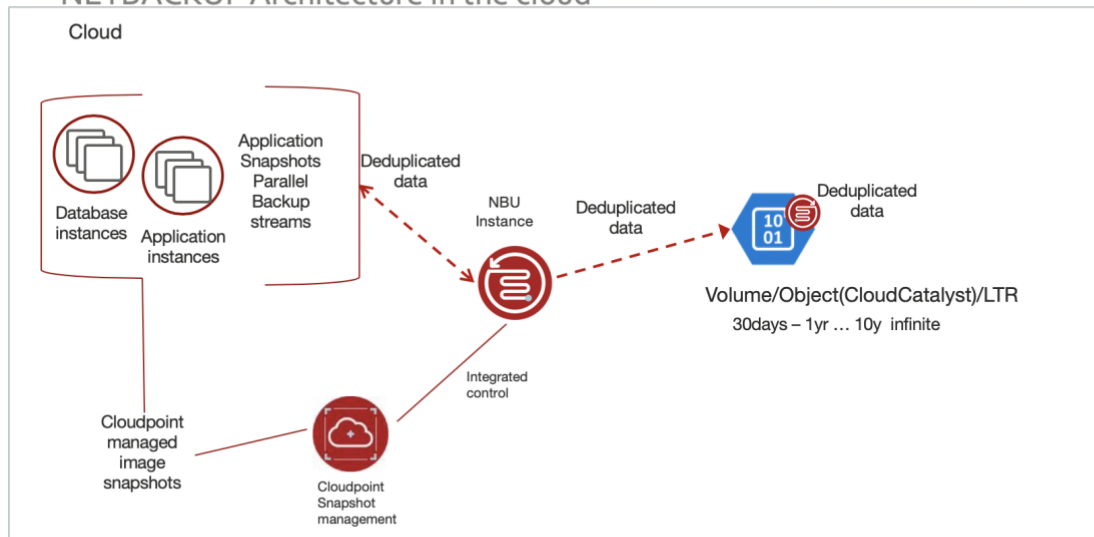
NetBackup offers an Azure VM image and an ARM template which is a template to make simplify and automate Master and/or Media Servers provisioning.

NetBackup is available in the Azure Marketplace:

<https://azuremarketplace.microsoft.com/marketplace/apps/veritas.veritas-netbackup-8-s>

BACKUP IN THE CLOUD – PROTECT WORKLOADS

NETBACKUP Architecture in the cloud



- Fast and easy deployment of NetBackup using marketplace automation
- Using traditional MSDP or add in CloudCatalyst to write to Object storage

Figure 4 - Workloads in Azure VMs

Similar to On-Prem, NetBackup can work in Azure with 3rd party 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 performs the work to dedupe the data before sending it on to Azure Block Blob storage. As stated earlier third-party deduplication

appliances are not compatible with MSDP and will require the data be rehydrated prior to sending to the device.

There are many other configuration options using NetBackup with Azure that can be tailored to the customers' needs, these use cases outline a handful of them.

DISASTER RECOVERY USING AZURE

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 second part is running as a cloud-based service. This setup would be to use the functionality of AIR to provide automation and optimization to deliver the data to the cloud for mobility and disaster recovery. (figure 5).

BACKUP TO THE CLOUD – DISASTER RECOVERY AND LTR

NETBACKUP ARCHITECTURE EXTENDED TO THE CLOUD

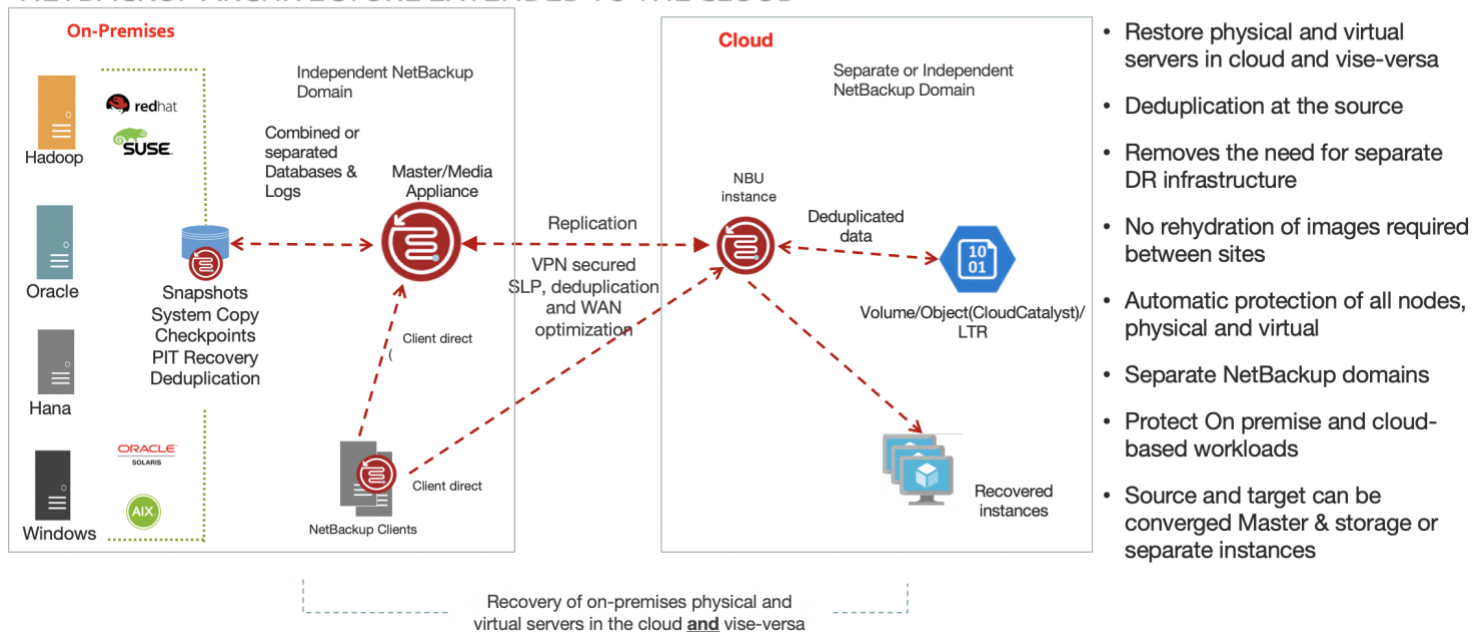


Figure 5 – NetBackup Auto-Image Replication for Cloud Recovery

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 Azure. From there, an AIR process can be used to automatically send data from MSDP in the data center to MSDP or

CloudCatalyst in Azure. The transferred data metadata is encapsulated in the transfer so the import into the Azure NetBackup domain is near instantaneous after the data is copied. Customers have been using this model for global DR protection, such as to move data from a datacenter in San Francisco to a datacenter in London, for quite a while. Leveraging this technology for a cloud target is no different for NetBackup it is just another AIR target.

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

Veritas offers additional solutions, such as Veritas Resiliency Platform (VRP), which automates workload migration into the cloud and can integrate with 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. VRP can also orchestrate workload recovery. Figure 6 provides an example of how NetBackup and VRP can be used to recover workloads into Azure.

In Cloud Data Recovery Setup

Leveraging NetBackup and VRP

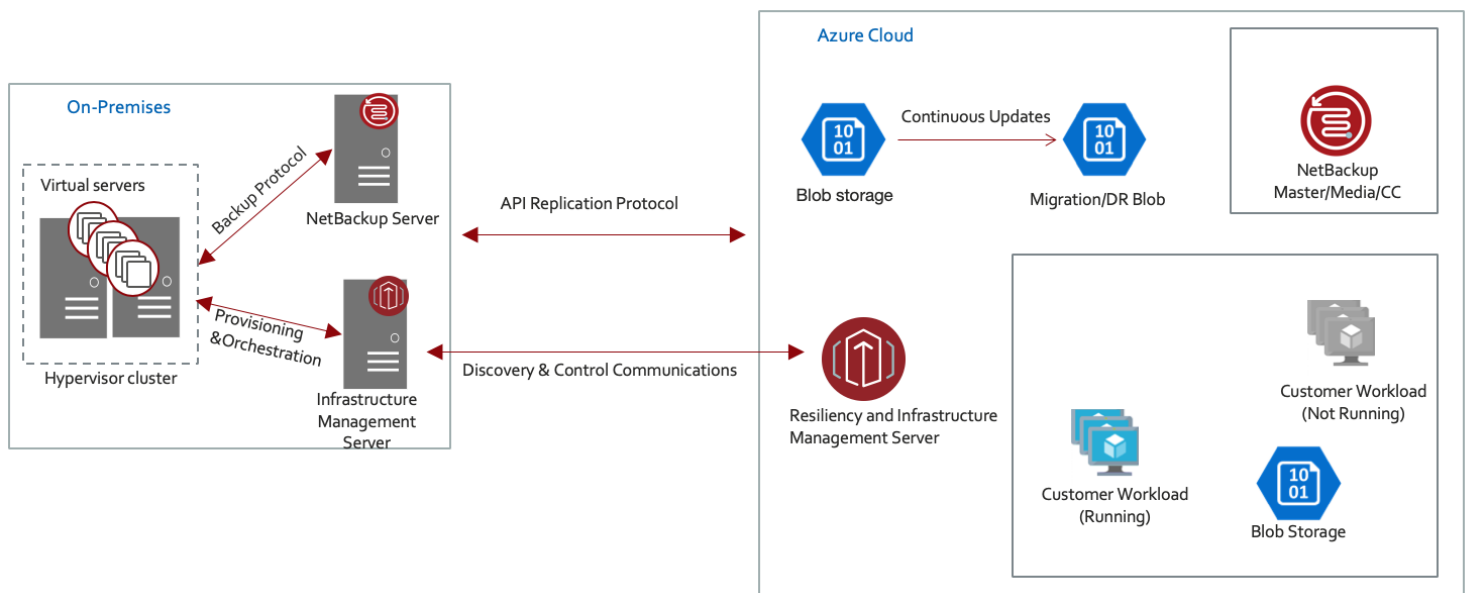


Figure 6 - Cloud DR with NetBackup and VRP

CLOUD SIZING AND PERFORMANCE

Sizing and performance in the Cloud is based on individual customer need thus will vary from customer to customer. Thus, this information is to be used as a guideline and actual deployment may be modified as needed.

To get data to the cloud, customers can utilize any Internet connection if the available bandwidth allows all the data to be transmitted in the target timeframe.

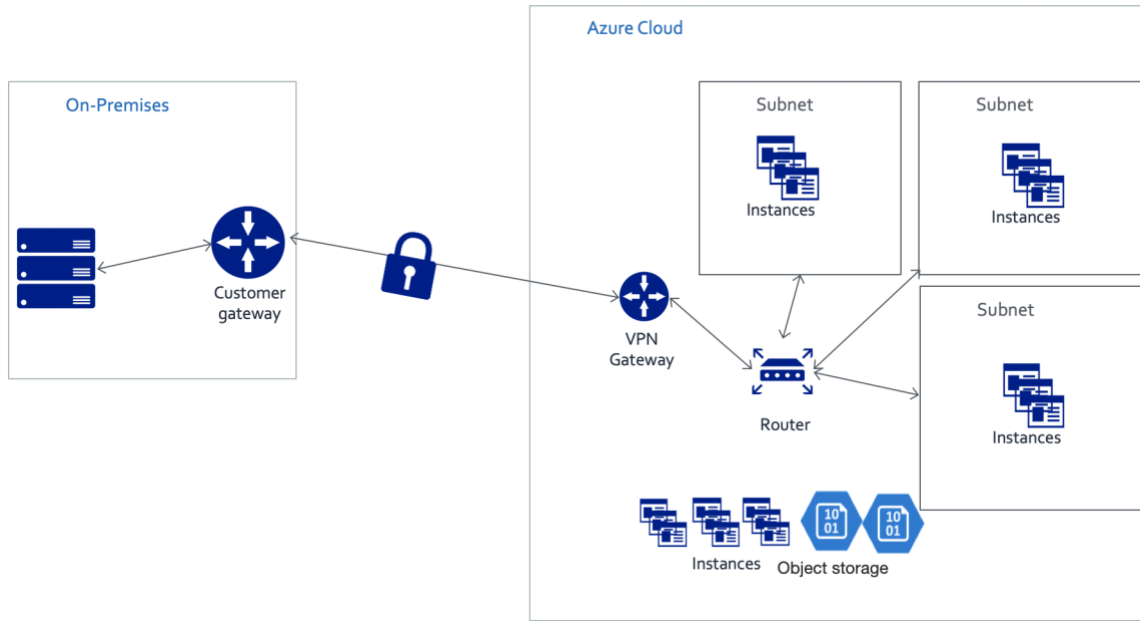


Figure 7 - A general view of cloud infrastructure connected with an on-premises datacenter

With Azure Express route, customers can get a dedicated link to Azure with dedicated high-speed WAN connectivity in the data center. Data can be compressed at the datacenter prior to being sent across the network to Azure or use CloudCatalyst to optimize the data being sent before it is transmitted to the cloud. Bandwidth can also be “throttled” if desired to prevent over-saturation of the network pipe.

AZURE INSTANCE MODEL

Azure has a Regional model. Various regions across the globe have been configured. Many regions have Availability Zones, which are multiple datacenters within the region that communicate with each other over high bandwidth, low latency 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 Azure cloud, as well as have the option to select a geographically dispersed region to provide regional DR.

AZURE REGION PAIRS

Each Azure region is paired with another region within the same geography, together making a regional pair. The exception is Brazil South, which is paired with a region outside its geography. Across the region pairs Azure serializes platform updates (planned maintenance), so that only one paired region is updated at a time. In the event of an outage affecting multiple regions, at least one region in each pair will be prioritized for recovery.

<https://docs.microsoft.com/en-us/azure/best-practices-availability-paired-regions>

Azure allows the selection of a storage account to be Geo-Redundant Storage (GRS) to provide an enhanced level of durability and protect against a region wide disaster.

<https://docs.microsoft.com/en-us/azure/storage/common/storage-redundancy-grs>

AZURE STORAGE OPTIONS

One of the many benefits of the Azure storage model is the ability to quickly add storage to environments. Customers don't pay for the storage until it is utilized. 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, Azure utilizes various types of storage including object (Block Blob across Premium, Hot, Cool and Archive tiers), and block (Page Blob/Managed Disks) depending on the type of use case. Azure Block Blob storage can offer a hierarchical name space (HNS) that access to the data via HDFS APIs (<https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-introduction>) for big data and analytics workloads.

Other options include Azure Files or Azure NetApp Files for high performance file system targets. 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 priced per GB, per transaction and on data egress. Hot storage has the highest cost per GB but the lowest transaction cost, conversely Archive storage is the lowest cost per GB but has a high cost for transactions and data retrieval. The latter is best suited for long term storage, where there is low probability of data retrieval; quite often a tape archive replacement.

Detailed pricing can be seen here <https://azure.microsoft.com/en-us/pricing/details/storage/>

NETBACKUP AZURE INSTANCE SIZING

This example architecture is based on a single NetBackup domain consisting of a NetBackup Master Server, several MSDP Media Servers and a single NetBackup CloudCatalyst Server for Azure storage.

Typically, backups are written directly to each media server's MSDP storage for an immediate copy, then duplicated via CloudCatalyst to Azure Storage. Alternatively, 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 perform deduplication
 - An optional step that can present a local copy for recovery on premises
 - Multiple MDSP pools also help distribute copies across different regions and infrastructures
- NetBackup CloudCatalyst Server
 - The CloudCatalyst Server processes deduplicated writes to Azure Storage. It is a dedicated high-end RedHat server that meets the minimum requirements for CloudCatalyst. It caches the deduplicated backups images from the MSDP media servers or directly from clients while transmitting them to Azure Block Blob Storage.
 - CloudCatalyst can also be a direct backup target, but should be sized larger
- 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 and expected growth. Azure 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 total 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.

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 Azure built on an Azure VM with Azure Block Blob Storage.

Specifications for MSDP Media Server in Azure

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 Azure, 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

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)
- 16 GB RAM Minimal (32GB preferred)
- 1 TB of Disk Storage

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 Azure VM instance can easily be expanded to meet higher requirements that may happen over time. For more information:

<https://docs.microsoft.com/en-us/azure/virtual-machines/windows/resize-vm>

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 Azure, and backup in the cloud workloads only.

Assumptions

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 in Front End Terabytes (FETB):
 - 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

MASTER SERVER RECOMMENDATIONS

Environment Size/ Components	Small	Medium	Large	Extra Large
Master Server	32 GB/8 vCPUs Storage: 500GB Catalog 5GB	64 GB/ 8vCPUs Storage: 500GB Catalog 5GB	64 GB/ 16vCPUs Storage: 500GB Catalog 10GB	122 GB/ 16vCPUs Storage: 500GB Catalog 10GB

MEDIA SERVER DEDUPLICATION POOL RECOMMENDATIONS

Running traditional MSDP in a cloud environment requires specific resources to be available such as 10G network, Managed Disks with required performance etc. The recommendations below have been formulated using Azure kits that addresses MSDP pools of different sizes. These are just recommendations and specific

customer environments may have different needs. Depending on the Azure 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 Dedup target
 - MSDP will be storing all data in managed disks
 - 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:

	Storage	Cores	RAM	Networking	IOPS
10 TB (Small)	1x160 GB SSD 1x16 TB SSD	8	61		
1-20 TB (Small)	1x80 GB SSD 1x16 TB SSD	36	60	10 GB	Provisioned IOPS (SSD)
32 TB (Medium)	1x80 GB SSD 2x16 TB SSD	16	30	10 GB	
	1x160 GB SSD 2x16 TB SSD IOPS – 12,000	8	61		12,000
32-64 TB (Large)	1x80 GB SSD 2-4x16 TB SSD	40	160	10 GB	
	1x80 GB SSD 2-4x16 TB SSD	36	60	10 GB	
	1x160 GB SSD 2x16 TB SSD IOPS – 12,000	8	61	10 GB	12,000
32-96 TB (xLarge)	1x80 GB SSD 2-4x16 TB SSD	40	160	10 GB	
	2x320 GB SSD 2-6x16 TB SSD IOPS -12,000	32	144	10 GB	12,000

Product	Role	Storage	CPUs	RAM(GB)
NBU	CloudCatalyst Minimum	250GB SSD 1+TB Cache (expandable)	4	16GB
	CloudCatalyst Large	500GB SSD 1+TB Cache (expandable)	8	32GB

For NetBackup CloudCatalyst servers, the above table lists a minimum configuration and a large configuration along with the Azure VM instance type and the storage configuration. Customers should 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 1 TB volume represents the local cache volume and mount location required for CloudCatalyst cloud deployments (Figure 8).

Cloud Storage Server Configuration Wizard - NetBackup

Add Storage Server
Select a media server and provide cloud storage service credentials. To be listed below in the media server drop-down list a security certificate must be deployed and NetBackup must be running including the NetBackup CloudStore Service Container (nbcssc).

Cloud storage provider - Microsoft Azure

Service host:

Storage server name:

Media server name:

Deduplication

☒ Enable NetBackup CloudCatalyst

Local cache directory:

Access details for Microsoft Azure account

Storage Account:

Access Key:

If you do not have Microsoft Azure account
[Create an account with Microsoft Azure.](#)

To continue, click Next.

Figure 8 - 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 starting from the client, in transit and at rest. In addition, any data that is sent to another MSDP 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 Azure environment from the media server. This encryption is done at the 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.

Cost Overview

The cost of the cloud will vary depending on what is needed. Simple backups to Azure 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 Azure due to the bandwidth constraints of the network. In addition, Azure 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 Azure 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 Azure. To optimize the data transfer to Azure, NetBackup breaks data down to 64MB of deduplicated data before sending it to the

Configured Azure Storage Account. Each 64MB chunk write or read will incur a GET or PUT request.

NetBackup Storage Lifecycle Policies can help alleviate these costs by automating the movement of data to the cloud to occur only after a time the customer defines as beyond the likely restore period or by automating the removal of the onsite copy to occur only after that period.

For example, if a customer knows data is rarely requested from backups after 2 weeks, they can choose to replicate the data immediately for DR but maintain an onsite copy for 2 weeks to address any restore requests. NetBackup automation will manage the backup, copy and expiration of images as part of the overall backup operation. This model optimizes both onsite storage utilization and minimizes GET cost impacts from frequent restore requests.

Storage Costs

Just as in the above example of costs associated with putting objects into Azure Storage, there is also a cost of using Azure Storage. The costs and available Azure storage types will vary based on region, so be sure to check prices in your intended region when calculating costs. <https://azure.microsoft.com/en-us/pricing/details/storage/blobs/>

Compute Costs

Azure environments where VMs are configured in the cloud using disk storage 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 copying the data from the VM 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. The cost of running these workloads in Azure is part of the customer’s business analysis to determine if moving a workload to the cloud provides a cost benefit. In some cases, these costs can be less than maintaining a data center, in some cases the cost is more which will drive the rate of cloud adoption. That said, the peace of mind knowing the data is highly available and protected can weigh the decision towards cloud. NetBackup’s ability to run natively in the cloud alleviates the PUT/GET penalty of protecting cloud workloads with an On-Prem only solution.

When you deploy an Azure VM the cost of the instance is determined by the hardware type (CPU, memory) disk and utilization.

As noted, cost should not always be the sole determining factor when it comes to a cloud-based solution. There is more to maintaining a datacenter than hard costs. The Azure infrastructure can provide additional benefits that might seem more expensive up front, but the flexibility provided with on-demand storage, uptime guarantees, and staffing may make a move to Azure make sense. Azure is a modern and efficient infrastructure that can support many business needs, and NetBackup will be there to protect the data cloud.

To better understand the cost structure, Azure has created a cloud adoption framework that can be used to determine cost models, business cases and other planning considerations. Customers can input information about a planned deployment and review cost estimates <https://docs.microsoft.com/en-us/azure/architecture/cloud-adoption/>.

Another comprehensive calculator to determine basic cost models can be found here <https://azure.microsoft.com/en-us/pricing/details/virtual-machines/windows/>

SUMMARY

Customers are moving to the cloud and a number of cloud providers are moving to the forefront of the cloud megatrend. Microsoft 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 above guidelines called out in this guide for these desired use case, 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
Azure Marketplace	https://azuremarketplace.microsoft.com/marketplace/
NetBackup Security and Encryption Guide	https://www.veritas.com/support/en_US/doc/21733320-132525226-0/index
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
How to Resize an Azure VM	https://docs.microsoft.com/en-us/azure/virtual-machines/windows/resize-vm
NetBackup Deduplication Guide	https://www.veritas.com/support/en_US/doc/25074086-127355784-0/index
Azure Regions and Region Pairs	https://docs.microsoft.com/en-us/azure/best-practices-availability-paired-regions
Azure Availability Zones	https://docs.microsoft.com/en-us/azure/availability-zones/az-overview
Azure VMs Additional Information	https://azure.microsoft.com/services/virtual-machines/
Azure Blob Storage Additional information	https://azure.microsoft.com/en-us/services/storage/blobs/
Azure Archive Storage additional information	https://azure.microsoft.com/services/storage/archive/
Azure Express Route additional information	https://azure.microsoft.com/services/expressroute/
Azure Security and Compliance additional information	https://www.microsoft.com/trust-center/product-overview

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

www.veritas.com

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