

Veritas NetBackup™ for Microsoft Azure Stack Administrator's Guide

Release 9.0

VERITAS™

Veritas Microsoft Azure Stack Guide

Last updated: 2020-12-11

Legal Notice

Copyright © 2020 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION 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 DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

NB.docs@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Chapter 1	Introduction	7
	Protecting Microsoft Azure Stack VMs using NetBackup	7
	Backing up Microsoft Azure Stack VMs	9
	Restoring Microsoft Azure Stack VMs	10
	Managed disk VM	11
	NetBackup for Microsoft Azure Stack terminologies	11
Chapter 2	Pre-requisites for configuring the Microsoft Azure Stack plug-in for NetBackup	13
	Operating system and platform compatibility	13
	License for Microsoft Azure Stack plug-in for NetBackup	14
	Settings for Backup host in IPV6 setup	14
	Time Sync with Azure Stack	14
	About deployment of NetBackup to protect Microsoft Azure Stack	14
Chapter 3	Configuring NetBackup and Microsoft Azure Stack	15
	Overview of configuring NetBackup and Microsoft Azure Stack	15
	Managing backup hosts	17
	Whitelisting a backup host on NetBackup master server	17
	Adding a Microsoft Azure Stack custom role to provide access permissions to NetBackup administrator	18
	Configuring the Microsoft Azure plug-in using the <code>azurestack.conf</code> configuration file	22
	Whitelisting the configuration file path on NetBackup master server	24
	Creating a file that contains Microsoft Azure Stack credentials	25
	Configuring proxy settings for communication with Microsoft Azure Stack AAD authentication	28
	Adding Microsoft Azure Stack credentials in NetBackup	29
	Creating a BigData policy for Microsoft Azure Stack using the NetBackup Policies utility	31
	Stale Snapshot cleanup	32

Chapter 4	Performing backups and restores of Microsoft Azure Stack	33
	About backing up Microsoft Azure virtual machines	34
	About restoring Microsoft Azure Stack virtual machines	34
	About the restore scenarios for Microsoft Azure Stack VMs from the BAR interface	35
	Considerations for Microsoft Azure Stack VM restore and recovery	37
	Using the BAR interface to restore an Microsoft Azure Stack VM at the same location	39
	Using the <code>bprestore</code> command to restore Microsoft Azure Stack VM at the same location	41
	Using the BAR interface to restore an Microsoft Azure Stack VM to different location	43
	Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location	45
	Conversion of Unmanaged disk VM to Managed disk VM	52
	Conversion of Unmanaged disk VM backed up using old plugin to Managed disk VM	52
	Using the <code>bprestore</code> command to restore Microsoft Azure VM with modified metadata to an alternate location	53
	Using the <code>bprestore</code> command to restore Microsoft Azure Stack VM with modified metadata to an alternate region	56
Chapter 5	Troubleshooting	61
	About NetBackup for Microsoft Azure debug logging	61
	Known limitations for Microsoft Azure protection using NetBackup	62
	Backup fails with error 6662	63
	Backup fails with error 6661	63
	Backup fails with error 6646	63
	Backup fails with error 6629	64
	Backup fails with error 6626	64
	Backup fails with error 6630	64
	Restore fails with error 2850	65
	Backup fails with error 1	65
	Adding Azure Stack credentials to NetBackup fails with error 9101	65
	Adding Azure Stack credentials to NetBackup fails with error 7610	66

Introduction

This chapter includes the following topics:

- Protecting Microsoft Azure Stack VMs using NetBackup
- Backing up Microsoft Azure Stack VMs
- Restoring Microsoft Azure Stack VMs
- Managed disk VM
- NetBackup for Microsoft Azure Stack terminologies


Protecting Microsoft Azure Stack VMs using NetBackup

You can use NetBackup and NetBackup Parallel Streaming Framework (PSF) to protect your Azure Stack VMs.

The following diagram provides an overview of how Microsoft Azure Stack VMs are protected by NetBackup.

Also, review the definitions of terminologies. See “NetBackup for Microsoft Azure Stack terminologies” on page 11.

Figure 1-1 Architectural overview



As illustrated in the diagram:

- The VMs are backed up in parallel streams wherein, NetBackup fetches blob storage data of VHDs. Each backup host fetches data associated with one or multiple VMs. In case of multiple backup hosts, sets of VMs are distributed to each backup host. The job processing is accelerated due to multiple backup hosts and parallel streams.

Note: One VHD's data is not fetched parallelly by multiple backup hosts.

- The communication between the Microsoft Azure Stack and the NetBackup is enabled using the NetBackup plug-in for Microsoft Azure Stack. For this release, the plug-in is available separately and must be installed on all the backup hosts.
- For NetBackup communication, you need to configure a BigData policy wherein, you need to use `Application_Type=azurestack`, and add the related backup hosts.
- You can configure a NetBackup media server, client, or master server as a backup host. Also, depending on the number of VMs, you can add or remove backup hosts. You can scale up your environment easily by adding more backup hosts.

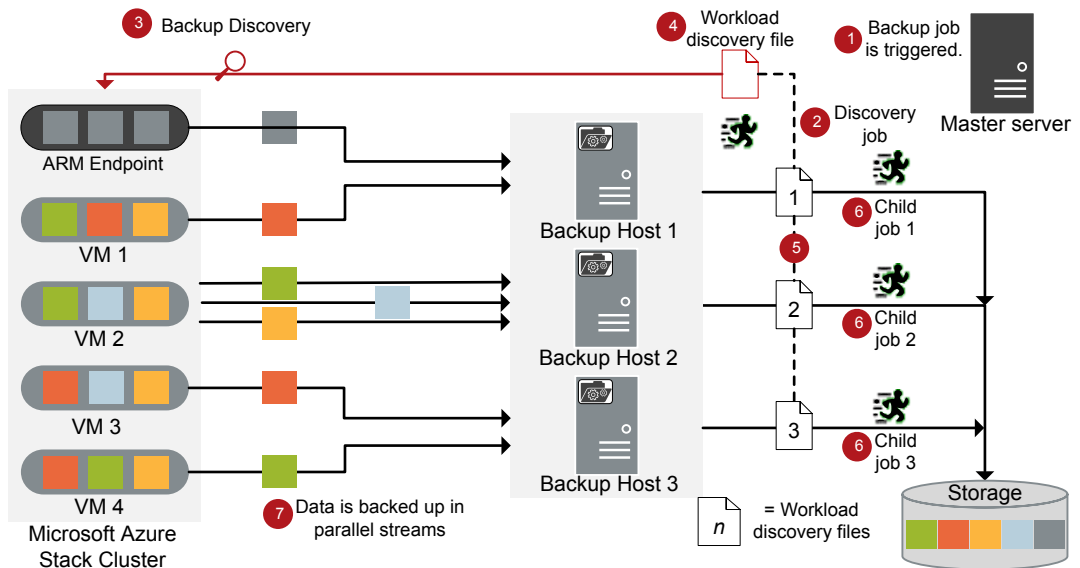
It is recommended that you use a NetBackup media server or a client as a backup host.

- The NetBackup Parallel Streaming Framework enables agentless backup wherein the backup and restore operations run on the backup hosts. There is no agent footprint on the Microsoft Azure Stack VMs. Also, NetBackup is not affected by the Microsoft Azure Stack upgrades or maintenance.

Backing up Microsoft Azure Stack VMs

The following diagram provides an overview of the backup flow:

Figure 1-2 Backup flow



As illustrated in the diagram:

1. A scheduled backup job is triggered from the master server.
2. Backup job for Microsoft Azure Stack is a compound job. When the backup job is triggered, first a discovery job runs.
3. During discovery, the first backup host connects with the Azure Resource Manager (ARM) Endpoint and performs a discovery to get details of VMs and their associated metadata that needs to be backed up.

4. A workload discovery file is created on the backup host. The workload discovery file contains the details of the data that needs to be backed up from the different VMs.
5. The backup host uses the workload discovery file get the details of data that it will backup. Individual workload discovery files are created for each backup host.
6. Individual backup jobs are executed for each backup host. As specified in the workload distribution files, data is backed up.
7. Data blocks are streamed simultaneously from different VMs to multiple backup hosts. The number of parallel streams is the same as the number of backup hosts.

To increase the number of streams, set the `NO_OF_BACKUP_STREAMS_PER_NODE` configuration parameter. For more information, refer to [Configuring the Microsoft Azure plug-in using the `azurestack.conf` configuration file.](#)

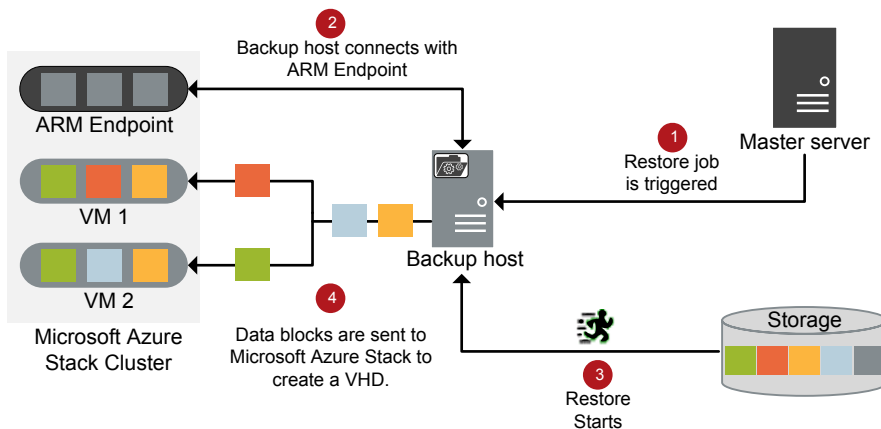
The compound backup job is not completed until all the child jobs are completed.

Restoring Microsoft Azure Stack VMs

For restore only one backup host is used.

The following diagram provides an overview of the restore flow.

Figure 1-3 Restore flow



As illustrated in the diagram:

1. The restore job is triggered from the master server.
2. The backup host connects with the Azure Resource Manager (ARM) Endpoint (source client). Backup host is the destination client.
3. The actual data restore from the storage media starts.
4. Data blocks are sent to Microsoft Azure Stack to create a VHD. After the VHD is created, VM is created and instantiated.

Managed disk VM

- Managed disks are block-level storage volumes that are managed by ARM. With managed disks, you must specify the disk size, the disk type, and provision the disk. The rest is handled automatically.
- Management APIs for Managed Disks are available in the compute provider. The Azure Stack plug-in uses REST APIs to:
 - Create Managed Disk Snapshot
 - Export Snapshot (Grant Access)
 - Unexport (Remove Access)
 - Delete Snapshot
- Single BigData policy can protect both Managed and Unmanaged disk VMs.
- During restore it is possible to convert unmanaged disk VM to managed disk.
- Storage account is required during restore of managed disk to create temporary `vhd` file in the target subscription where VM or the disk is being restored. This is due to lack of required API for creating Managed Disk directly during restore.

NetBackup for Microsoft Azure Stack terminologies

The following table defines the terms you will come across when using NetBackup for protecting Microsoft Azure stack.

Table 1-1 NetBackup terminologies

Terminology	Definition
Compound job	<p>A backup job for Microsoft Azure Stack is a compound job.</p> <ul style="list-style-type: none"> ■ The backup job runs a discovery job for getting information of the data to be backed up. ■ Child jobs are created for each backup host that performs the actual data transfer. ■ After the backup is complete, the job cleans up the snapshots on the Microsoft Azure Stack and is then marked complete.
Discovery job	<p>When a backup job is executed, first a discovery job is created. The discovery job communicates with the ARM Endpoint and gathers information of VMs and associated VHDs. At the end of the discovery, the job populates a workload discovery file that NetBackup then uses to distribute the workload amongst the backup hosts.</p>
Child job	<p>For backup, a separate child job is created for each backup host to transfer data to the storage media.</p>
Workload discovery file	<p>During discovery, when the backup host communicates with the ARM Endpoint, a workload discovery file is created. The file contains information about the VMs and associated VHDs.</p>
Parallel streams	<p>The NetBackup parallel streaming framework allows multiple VMs to be backed up using multiple backup hosts simultaneously.</p>
Backup host	<p>The backup host acts as a proxy client. All the backup and restore operations are executed through the backup host.</p> <p>You can configure media servers, clients, or a master server as a backup host.</p> <p>The backup host is also used as destination client during restores.</p>
BigData policy	<p>The BigData policy is introduced to:</p> <ul style="list-style-type: none"> ■ Specify the application type. ■ Allow backing up distributed multi-node environments. ■ Associate backup hosts. ■ Perform workload distribution.

Pre-requisites for configuring the Microsoft Azure Stack plug-in for NetBackup

This chapter includes the following topics:

- Operating system and platform compatibility
- License for Microsoft Azure Stack plug-in for NetBackup
- Settings for Backup host in IPV6 setup
- Time Sync with Azure Stack
- About deployment of NetBackup to protect Microsoft Azure Stack

Operating system and platform compatibility

For backup host (media server or NetBackup Appliance) as per your requirement:

- Red Hat Enterprise Linux (RHEL) 7.4 and later are supported
- Red Hat Enterprise Linux (RHEL) 8.0 and later versions, requires additional package of “compat-openssl10” to be installed.

For more information, refer to the NetBackup Compatibility List at the following location:

https://www.veritas.com/support/en_US/article.100040093

License for Microsoft Azure Stack plug-in for NetBackup

For the licensing requirements to run the backup and restore operations using the Microsoft Azure Stack plug-in for NetBackup, refer to the following page:

How to use NetBackup plug-ins and agents: download, install, and availability information

More information is available on how to add licenses.

See the NetBackup Administrator's Guide, Volume I

Settings for Backup host in IPV6 setup

Backup host on which Azure Stack plug-in is installed needs to be dual stack:

Ensure that Azure Stack ARM endpoint should be accessible from configured Backup host from policy.

Backup host should be configured with IPV4 and IPV6 addresses.

Time Sync with Azure Stack

Date and time for NetBackup master server, media server, and backup host must be in sync with Azure Stack to avoid backup and restore failures.

After the time sync is done, restart the NetBackup services on the backup host.

About deployment of NetBackup to protect Microsoft Azure Stack

- If you have deployed a multinode Microsoft Azure Stack cluster, you can deploy the NetBackup servers and the backup host outside your cluster and then configure the connection.
See "Overview of configuring NetBackup and Microsoft Azure Stack" on page 15.

Configuring NetBackup and Microsoft Azure Stack

This chapter includes the following topics:

- Overview of configuring NetBackup and Microsoft Azure Stack
- Managing backup hosts
- Adding a Microsoft Azure Stack custom role to provide access permissions to NetBackup administrator
- Configuring the Microsoft Azure plug-in using the azurestack.conf configuration file
- Creating a file that contains Microsoft Azure Stack credentials
- Adding Microsoft Azure Stack credentials in NetBackup
- Creating a BigData policy for Microsoft Azure Stack using the NetBackup Policies utility
- Stale Snapshot cleanup

Overview of configuring NetBackup and Microsoft Azure Stack

The following table lists the steps for configuring NetBackup for Microsoft Azure Stack that are required for authentication:

Table 3-1 Steps for configuring NetBackup for Microsoft Azure Stack

Steps	Component	Details
1	Backup host	<p>Create a backup host and whitelist a NetBackup client if you want to use it as a backup host.</p> <p>For more information, refer to:</p> <ul style="list-style-type: none"> ■ See “Managing backup hosts” on page 17. ■ See “Whitelisting a backup host on NetBackup master server” on page 17.
2	Custom NetBackup role in Microsoft Azure Stack	<p>Create a custom role in Microsoft Azure Stack for NetBackup to backup and restore VMs.</p> <p>For more information, refer to:</p> <p>See “Adding a Microsoft Azure Stack custom role to provide access permissions to NetBackup administrator” on page 18.</p>
3	<ul style="list-style-type: none"> ■ Microsoft Azure Stack credentials file ■ Microsoft Azure Stack plug-in configuration file 	<ul style="list-style-type: none"> ■ Create a file that contains the Azure stack credentials on the master server. See “Creating a file that contains Microsoft Azure Stack credentials” on page 25. ■ Configure the Microsoft Azure Stack plug-in using a configuration file and whitelist the configuration file path. For more information, refer to: <ul style="list-style-type: none"> ■ See “Configuring the Microsoft Azure plug-in using the <code>azurestack.conf</code> configuration file” on page 22. ■ See “Whitelisting the configuration file path on NetBackup master server” on page 24. ■ Add Microsoft Azure Stack credentials to NetBackup to establish communication and protect the data. For more information, refer to: See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.
4	BigData policy	<p>Creating a BigData policy for Microsoft Azure Stack.</p> <p>For more information, refer to:</p> <p>See “Creating a BigData policy for Microsoft Azure Stack using the NetBackup Policies utility” on page 31.</p>

Managing backup hosts

A backup host acts as a proxy client which hosts all the backup and restore operations for Microsoft Azure Stack. In case of Microsoft Azure Stack plug-in for NetBackup, backup host performs all the backup and restore operations without any separate agent installed on the Microsoft Azure Stack.

The backup host must be a RHEL 7.4 or later computer. NetBackup supports only the RHEL platform for a backup host.

Consider the following before adding a backup host:

- For backup operations, you can add one or more backup hosts.
- For restore operations, you can use only one backup host.
- A master, media, or client can perform the role of a backup host.

Note: It is recommended that you use a NetBackup media server or a client as a backup host.

- When using multiple backup host, make sure that all backup hosts are communicating with the media server.
- Azure Stack identity providers
 - For the Azure Active Directory (AAD) identity provider, all backup hosts require connectivity to <https://login.microsoftonline.com>, Azure Resource Manager endpoints, or Azure blob storage endpoints, which require ports 80 and 443 for communication.
 - For the Active Directory Federation Services (ADFS) identity provider, all backup hosts require connectivity to Azure Resource Manager endpoints, Azure blob storage endpoints, or ADFS endpoints, which require ports 80 and 443 for communication.

You can add a backup host while configuring **BigData** policy using the NetBackup Administration Console.

See “Creating a BigData policy for Microsoft Azure Stack using the NetBackup Policies utility” on page 31.

Whitelisting a backup host on NetBackup master server

To use the NetBackup client as a backup host, you must whitelist it. Perform the whitelisting procedure on the NetBackup master server .

Whitelisting is a security practice used for restricting systems from running software or applications unless these have been approved for safe execution.

To whitelist a backup host on NetBackup master server

- ◆ Run the following command on the NetBackup master server:
 - For UNIX

```
bpsetconfig -h masterserver
bpsetconfig> APP_PROXY_SERVER = clientname1.domain.org
bpsetconfig> APP_PROXY_SERVER = clientname2.domain.org
bpsetconfig>
UNIX systems: <ctl-D>
```
 - For Windows

```
bpsetconfig -h masterserver
bpsetconfig> APP_PROXY_SERVER = clientname1.domain.org
bpsetconfig> APP_PROXY_SERVER = clientname2.domain.org
bpsetconfig>
Windows systems: <ctl-Z>
```

This command sets the `APP_PROXY_SERVER = clientname` entry in the backup configuration (`bp.conf`) file.

For more information about the `APP_PROXY_SERVER = clientname`, refer to the *Configuration options for NetBackup clients* section in *NetBackup Administrator's Guide, Volume I*

Veritas NetBackup Documentation

Adding a Microsoft Azure Stack custom role to provide access permissions to NetBackup administrator

NetBackup requires access to Azure Stack subscriptions to protect them. You must create a custom user in Active Directory for NetBackup and grant the user the role to access the subscriptions. You can either give a co-owner role to the user or you can create a custom role with permissions that are required for backup and recovery. An Azure Stack administrator as a subscription owner can create the custom role for a subscription.

The minimum permissions that NetBackup requires are as follows:

- Microsoft.Compute/virtualMachines/*

Adding a Microsoft Azure Stack custom role to provide access permissions to NetBackup administrator

- Microsoft.Network/networkInterfaces/*
- Microsoft.Network/networkSecurityGroups/join/action
- Microsoft.Network/networkSecurityGroups/read
- Microsoft.Network/publicIPAddresses/join/action
- Microsoft.Network/publicIPAddresses/read
- Microsoft.Network/publicIPAddresses/write
- Microsoft.Network/virtualNetworks/read
- Microsoft.Network/virtualNetworks/subnets/read
- Microsoft.Network/virtualNetworks/subnets/join/action
- Microsoft.Resources/subscriptions/resourceGroups/read
- Microsoft.Storage/storageAccounts/read
- Microsoft.Storage/storageAccounts/listKeys/action

To create a custom role, complete the following steps:

- 1** For Active Directory Federation Services (ADFS) Create a user or service principal named `nbu_azst` in the Active Directory from the Active Directory Users and Computers dialog box from Microsoft Management Console.

For Microsoft Azure Active Directory (Azure AD) Create the service principal from the Microsoft Azure Active Directory Users dialog box.

Complete the following steps on a Windows computer that has PowerShell for Azure Stack.

For more information, refer to <https://docs.microsoft.com/en-us/azure/azure-stack/azure-stack-powershell-install>.

Adding a Microsoft Azure Stack custom role to provide access permissions to NetBackup administrator

- 2** Create a new text file `rbac_NBU_role.json` and add the following script in the file:

```
{
  "Name": "NBU BnR Role",
  "IsCustom": true,
  "Description": "Let's you perform backup and recovery of VMs",
  "Actions": [
    "Microsoft.Compute/virtualMachines/*",
    "Microsoft.Compute/Disks/read",
    "Microsoft.Compute/Disks/write",
    "Microsoft.Compute/Disks/beginGetAccess/action",
    "Microsoft.Compute/Disks/endGetAccess/action",
    "Microsoft.Compute/Snapshots/*",
    "Microsoft.Network/networkInterfaces/*",
    "Microsoft.Network/networkSecurityGroups/join/action",
    "Microsoft.Network/networkSecurityGroups/read",
    "Microsoft.Network/publicIPAddresses/join/action",
    "Microsoft.Network/publicIPAddresses/read",
    "Microsoft.Network/publicIPAddresses/write",
    "Microsoft.Network/virtualNetworks/read",
    "Microsoft.Network/virtualNetworks/subnets/read",
    "Microsoft.Network/virtualNetworks/subnets/join/action",
    "Microsoft.Resources/subscriptions/resourceGroups/read",
    "Microsoft.Resources/Resources/read",
    "Microsoft.Storage/storageAccounts/read",
    "Microsoft.Storage/storageAccounts/listKeys/action"
  ],
  "NotActions": [],
  "AssignableScopes": [
    "/subscriptions/subscription_ID_1",
    "/subscriptions/subscription_ID_2"
    .
    .
    .
  ]
}
```

Note: Ensure that you add the required subscriptions under the `AssignableScopes` field so that the custom role is created with those subscriptions.

For example, in the file snippet, replace `subscription_ID_1` and `subscription_ID_2` with actual subscription IDs that you have.

3 Run the following commands:

- `Add-AzureRMEnvironment -Name AzureStackAdmin -ArmEndpoint "ArmEndpointValue"`
For example, `Add-AzureRMEnvironment -Name AzureStackAdmin -ArmEndpoint "https://management.local.azurestack.external"`
- `Add-AzureRmAccount -EnvironmentName "AzureStackAdmin"`
- `New-AzureRmRoleDefinition -InputFile "<directory_path>\rbac_NBU_role.json"`

You can use the following ARM endpoints:

- provider subscription
- tenant subscription

4 Open the Microsoft Azure Stack console and complete the following steps:

1. Click **Menu** and open the subscriptions that you want to protect with NetBackup. Click **Access Control (IAM) > Roles** to view the newly created role.
2. From **Subscriptions > Access Control (IAM)**, click **Add**. In the **Select Name** field add `nbu_azst` user (ADFS) or the display name of the service principal (AAD), in the **Type** field select **User**, and in the **Role** field select the newly added role.

5 Add the `nbu_azst` user or service principal to the `tpconfig` command to take backups.

See "Adding Microsoft Azure Stack credentials in NetBackup" on page 29.

Configuring the Microsoft Azure plug-in using the `azurestack.conf` configuration file

NetBackup master server uses the `azurestack.conf` file to save the configuration settings for communication with Microsoft Azure Stack.

You must create the `azurestack.conf` file in the `/usr/opensv/var/global` directory.

Configuration definitions must be in the format of "attribute = value"; the single space before and after the '=' is required.

The options and values are case-sensitive.

Note: You must not provide a blank value for any of the parameters, or the backup job fails.

Here is a sample of the `azurestack.conf` file:

```
VM_STATE = Running
FETCH_STORAGE_KEYS = false
CA_FILE_PATH = //directory_path_system_CA_certificate/certificate_name.crt
VM_SNAPSHOT_IN_DISCOVERY = true
NO_OF_BACKUP_STREAMS_PER_NODE = 1
ENABLE_SNAPSHOT_CLEANUP = 0
SNAPSHOT_CLEANUP_MIN = 1440
SNAPSHOT_FETCH_RETRY_COUNT = 60
ASYNC_SNAPSHOT_SUPPORT = true
SET_PUBLIC_IP_DURING_RESTORE = false
STAGING_STORAGE_ACCOUNTS = headsuptsta, restorestadisks, stafordiskstorages
```

- The possible values for `VM_STATE` are `Running`, `Deallocated`, or `Stopped`.
- The value for `SNAPSHOT_FETCH_RETRY_COUNT` specifies the maximum retries that can happen if a VM snapshot process fails. The value cannot exceed 3.
- The value for `FETCH_STORAGE_KEYS` specifies whether the storage account with access key is required in the Azure Stack credentials file. The value can be either `true` or `false`. If the value is `true`, then you do not specify the storage account with access keys in the credential file. The default value is `true`.
- The value for `CA_FILE_PATH` is the directory path of the system CA certificate and the certificate name. For example, `/etc/pki/tls/certs/ca-bundle.crt`. This directory path is the default path for all system CA certificates.
- The value for `VM_SNAPSHOT_IN_DISCOVERY` defines whether VM disk snapshots are created for all disks that are attached to the VMs that are specified in the backup selection in the backup policy. The default value is `false`, specifying that snapshots of all disks that are attached to a single VM are created at a time during the backup job.
- The value for `NO_OF_BACKUP_STREAMS_PER_NODE` specifies the maximum number of backup streams per backup host. The default value is 1. You can define the value up to 8.
- The value for `ENABLE_SNAPSHOT_CLEANUP` specifies when to clean up stale snapshots. You can use the following values:
 - 0

Do not delete the stale snapshots. This is the default value.

- 1
Delete the stale snapshots after the backup job completes.
- 2
Delete the stale snapshots as part of the next discovery job.
- The value for `SNAPSHOT_CLEANUP_MIN` specifies the time in minutes after which the snapshot is deleted. The default value is 1440 minutes (24 hours). The ideal value is the time difference in minutes between two backup jobs.
- The value for `SNAPSHOT_FETCH_RETRY_COUNT` specifies the retry count for the plug-in to check for the snapshot in a given storage container. The default value is 60 and the values can be between 60 and 120.
- The value for `ASYNC_SNAPSHOT_SUPPORT` specifies that unmanaged disk snapshots are taken using the async REST API. The default value is true. Set the value to false to disable unmanaged disk snapshots using the async REST API.
- The value for `SET_PUBLIC_IP_DURING_RESTORE` specifies whether to use a public IP address during the VM restore. The default value is true that specifies that the VM is restored with a public IP address. Set the value as false, to restore the VM with a private IP address.
- The value for `STAGING_STORAGE_ACCOUNTS` is a comma separated list of storage account names that are already available on Azure Stack and the value is required for only managed disks VM restore or recovery. You can specify one or many storage accounts names for this key. Storage account name is required for restoring VHD blobs for the managed disks conversion operation. As per the target subscription location that is the selected as restore location, Netbackup Azure Stack plug-in automatically picks only one storage account from the list that is provided for this key. This is a one-time activity where you can specify all storage accounts needed for all restore jobs as per the target location selected for each restore job.

Note: Do not add `VM_STATE` in the `azurestack.conf` file if you want to take a backup if all VMs.

Whitelisting the configuration file path on NetBackup master server

After you create the configuration file, you must whitelist the path of the configuration file so that NetBackup lets the backup operation to run successfully. Run the whitelisting procedure on a NetBackup master server.

Whitelisting is a security practice used for restricting systems from running software or applications unless these have been approved for safe execution.

To whitelist the configuration file path:

Run the following command on the NetBackup master server:

1 For UNIX:

```
bpsetconfig -h masterserver_name
bpsetconfig BPCD_WHITELIST_PATH = /usr/opensv/var/global/
```

Exit the command line.

For clustered master server:

```
bpsetconfig BPCD_WHITELIST_PATH = /<shared drive mount
location>/var/global/
```

2 For Windows:

```
bpsetconfig -h masterserver_name
bpsetconfig BPCD_WHITELIST_PATH = <install_dir>\NetBackup\var\global\
```

For clustered master server:

```
bpsetconfig BPCD_WHITELIST_PATH = <Shared
drive>\NetBackup\var\global\
```

Exit the command line.

For more information about `BPCD_WHITELIST_PATH`, see the *Configuration options for NetBackup servers* section in the *NetBackup Administrator's Guide, Volume I*.

Creating a file that contains Microsoft Azure Stack credentials

To communicate with Microsoft Azure Stack, the plug-in must have access to the Microsoft Azure Stack credentials. The credentials must be stored in a file on the NetBackup master server. The credentials are stored in an encrypted format and the plug-in securely accesses the information.

To create a file with the Microsoft Azure Stack credentials on the master server:

- At any location on the master server, created a file with a JSON format. For example, you can create a file named `azurestack.creds` in the `/usr/opensv/var/global/` directory.
- Open the file and add the following content:

```
{
  "IdentityProvider": "ADFS",
  "TenantId": "tenant.domain.com",
  "ClientId": "1950a258-227b-4e31-a9cf-717495945fc2",
  "ClientSecret": "client_secret",
  "AuthResource":
  "https://management.adfs.azurestack.local/metadata/a6ad92e4-5b80-4c88-b84f-a7f25c12ba27",
  "teststorageacl":
  "9ghIt35bQeSvjZxXUPj8LinMs6aXPb2tMFjXVIG6N2v2FO6LRg+HzLz2LX1xR/qRkQYvNPIaE/v+QnUovzaKpQ==",
  "rg1disks540":
  "R6Lu3buXZ4HVtRTrNEHzzJqo2gShjQytfjX1hRkvmfQMVWnvKwMmEt2CUfmh1bxI7JCE0Gh5TKA9r3I88eit2FdA==",
  "StorageAccount3": "asasdlfkjaasdfasdfasdfasdf09sd8fhaopisdmbanpsdf98asdfpusadf====",
  "StorageAccount11": "90asdfasdfasdfasd-98fha-sdf98asdb-fau9bsdf-ay8svfasd==",
  "StorageAccount19": "90asdfasdfasdfasd-98fha-sdf98asdb-fau9bsdf-ay8svfasd==",
  "StorageAccount121": "90asdfasdfasdfasd-98fha-sdf98asdb-fau9bsdf-ay8svfasd==",
  "StorageAccount13": "90asdfasdfasdfasd-98fha-sdf98asdb-fau9bsdf-ay8svfasd==",
  "StorageAccount14": "90asdfasdfasdfasd-98fha-sdf98asdb-fau9bsdf-ay8svfasd==",
  "StorageAccount12": "90asdfasdfasdfasd-98fha-sdf98asdb-fau9bsdf-ay8svfasd=="
  ...
}
```

Note: The `StorageAccount` details are required if `FETCH_STORAGE_KEYS = false` in the `azurestack.conf` file.

Option	Identity Provider	Description
<code>IdentityProvider</code>	AAD and ADFS	Values can be either ADFS (Active Directory Federation Services) or AAD (Azure Active Directory).
<code>TenantId</code>	AAD	Value is the tenant domain. For example, <code>"tenant.onmicrosoft.com"</code> . See the section called "Obtaining the <code>TenantId</code> value for AAD" on page 27.
<code>ClientId</code>	ADFS AAD	Value is <code>1950a258-227b-4e31-a9cf-717495945fc2</code> . Value is the application ID of the service principal that has the NetBackup backup and recovery role for the subscriptions that NetBackup must protect. See the section called "Obtaining the <code>ClientId</code> value for AAD" on page 27.
<code>ClientSecret</code>	AAD	Value is the client secret of the service principal that has the NetBackup backup and recovery role for the subscriptions that NetBackup must protect. See the section called "Obtaining the <code>ClientSecret</code> value for AAD" on page 27.

Option	Identity Provider	Description
AuthResource	AAD and ADFS	<p>Value of the key audiences that is obtained by opening the following URL in a web browser:</p> <p><code>https://management.{region}.{azurestackFQDN}/metadata/endpoints?api-version=2015-01-01</code></p> <p>For example:</p> <p><code>https://management.eng.azurestack.veritas.com/metadata/endpoints?api-version=2015-01-01</code></p> <p>The URL returns a JSON value that is the value of the key audiences.</p>
StorageAccount	AAD and ADFS	<p>The storage account with the access key.</p> <p>If the value of <code>fetchStorageKeys</code> in the <code>azurestack.conf</code> file is <i>false</i>, then you must add this option.</p>

Obtaining the `TenantId` value for AAD

1. Sign in to `https://portal.azure.com`.
2. Open **Azure Active Directory > Properties** and locate the **Directory ID** that is the `TenantId`.

Obtaining the `ClientId` value for AAD

To obtain the `ClientId` value, you must create a new service principal or use an existing service principal.

1. Sign in to `https://portal.azure.com`.
2. Open **Azure Active Directory > App registrations**.
3. In the **Search by name or AppID** field, search for `NBU-ASTK-1` and click the service principal **Display Name** in the results.
4. Use any of the following steps to get the `ClientId`:
 - Open **Settings** and locate and copy **Application ID** that is the `ClientId`.
 - Open **Properties** and locate and copy **Application ID** that is the `ClientId`.

Obtaining the `ClientSecret` value for AAD

To obtain the `ClientSecret` value, you must create a new service principal or use an existing service principal.

1. Sign in to `https://portal.azure.com`.
2. Open **Azure Active Directory > App registrations > New application registration**.

3. Create an application with the **Name** as `NBU-ASTK-1`.
Select the **Application Type** as **Web App / API**.
Enter the **Sign-on URL** as `https://astk.nbu.com`.
Click **Create**.
4. Open **Azure Active Directory > App registrations**.
5. In the **Search by name or AppID** field, search for `NBU-ASTK-1` and click the service principal **Display Name** in the results.
6. Open **Settings > Keys** and add a new password information as follows and then save:
Description: `Credential_1`
Expires: `Never`
Value: `seedvalue_1`
7. **Value** displayed is the `ClientSecret`. The value is displayed only once. If you close the window, the value is not displayed again.

Configuring proxy settings for communication with Microsoft Azure Stack AAD authentication

If your network requires proxy settings so that the backup hosts can connect to the Internet, use any of the following methods:

- Use standard environment variable (simple configuration) `https_proxy` to specify the proxy URL, port number, username, and password in the following format:
`https_proxy=https://USERNAME:PASSWORD@PROXYIP_HOSTNAME:PROXYPORT`
- If you require a different proxy for the NetBackup Azure Stack plug-in or you do not want to use the `https_proxy` variable, then you can add the following proxy details in the credentials file:

Key	Description
<code>InternetProxyUrl</code>	Specify the proxy URL and port number to connect to the AAD authentication service and to <code>login.microsoftonline.com</code> over the internet. For example, <code>https://myproxyInternet.com:8000</code> .
<code>InternetProxyUsername</code>	Specify the username to authenticate the proxy internet URL, if required.

Key	Description
InternetProxyPassword	Specify the username to authenticate the proxy internet URL, if required.

```

{
  "IdentityProvider": "AAD",
  "TenantId": "tenant.domain.com",
  "ClientId": "1950a007-227b-4e31-a9cf-717495945fc2",
  "ClientSecret": "client_secret",
  "AuthResource": "https://management.adfs.azurestack.local/metadata/a6ad92e4-5b80-4c88-b055-a7f25c12ba27",
  "InternetProxyUrl": "proxy.domain.com:8080",
  "InternetProxyUsrename": "myusername",
  "InternetProxyPassword": "mypassword"
}

```

Adding Microsoft Azure Stack credentials in NetBackup

To establish a seamless communication between Microsoft Azure Stack clusters and NetBackup for successful backup and restore operations, you must add and update Microsoft Azure Stack credentials to the NetBackup master server.

Use the `tpconfig` command to add credentials in NetBackup master server.

For more information about the `tpconfig` command, see the NetBackup Commands Reference Guide.

To add credentials in NetBackup

- 1 Run `tpconfig` command from the following directory paths:

On UNIX systems, `/usr/opensv/volmgr/bin/`

- 2 Run the following command by providing appropriate values for each parameter to add Microsoft Azure Stack credentials:

```

tpconfig -add -application_server_user_id user_ID
-application_type application_type -application_server
application_server_name -password password_of_the_nbu_azst_user
-application_server_conf "/usr/<file_path>/azurestack.creds"

```

- For AAD, NetBackup uses `clientId` and `clientSecret`, so enter the value for `-application_server_user_id` as dummy and `-password` as dummy.

Note: The user you want to add must have co-owner permissions to the subscription you want to protect.

For example,

```
tpconfig -add -application_server_user_id example_user_ID  
-application_type azurestack -application_server  
application_server_name -password password_of_the_nbu_azst_user  
-application_server_conf "/usr/openv/var/global/azurestack.creds"
```

Here, the numeric value 8 can also be specified for the `-application_type` parameter that corresponds to Microsoft Azure Stack .

- 3 Run the `tpconfig -dappservers` command to verify if the NetBackup master server has the Azure credentials added.

For example, here is a sample output:

```
Application Server Host Name:      management.local.azurestack.external  
Application Server Type:          azurestack  
Required Port:                    0  
User of Application Host:         root
```

- 4 After you use `tpconfig` to add the credentials, you can delete the credentials file from the location that you have used in the `tpconfig -add` command..

- 5 Run the following command to update or delete the `tpconfig` credentials:

- Delete

```
tpconfig -delete -application_server_user_id user_ID  
-application_type application_type -application_server  
application_server_name -password password_of_the_nbu_azst_user  
-application_server_conf "/usr/<file_path>/azurestack.creds"
```

- Update

To change the attributes or options in the credentials file, update the credentials and then use the `tpconfig -update` command.

```
tpconfig -update -application_server_user_id user_ID  
-application_type application_type -application_server  
application_server_name -password password_of_the_nbu_azst_user  
-application_server_conf "/usr/<file_path>/azurestack.creds"
```

Creating a BigData policy for Microsoft Azure Stack using the NetBackup Policies utility

Use the following procedure to create a BigData policy with the NetBackup Policies utility.

To create a BigData policy with the NetBackup Policies utility

- 1 In the **NetBackup Administration Console**, in the left pane, expand **NetBackup Management > Policies**.
- 2 On the **Actions** menu, click **New > Policy**.
- 3 Type a unique name for the new policy in the **Add a New Policy** dialog box. Click **OK**.
- 4 On the **Attributes** tab, select **BigData** as the policy type.
- 5 On the **Attributes** tab, select the storage unit for BigData policy type.
- 6 On the **Schedules** tab, click **New** to create a new schedule.

You can create a schedule for a **Full Backup** for your BigData policy. Once you set the schedule, Microsoft Azure data is backed up automatically as per the set schedule without any further user intervention.
- 7 On the **Clients** tab, enter the IP address or the host name of the ARM Endpoint.

You can add the following ARM Endpoints:

 - Provider subscription
 - Tenant subscription
- 8 On the **Backup Selections** tab, enter the following parameters and their values as shown:
 - *Application_Type=azurestack*
The parameter values are case-sensitive.
 - *Backup_Host=IP_address or FQDN*
You can specify multiple backup hosts.
 - Specify assets to backup
 - For all the VMs in a subscription: */Subscription ID*
 - For all the VMS in a resource group: */Subscription ID/Resource Group*
 - For a single VM: */Subscription ID/Resource Group/VM Name*

Note: The directory or folder specified for backup selection while defining BigData Policy with Application_Type = azurestack must not contain space or comma in their names.

- 9 Click **OK** to save the changes.

Stale Snapshot cleanup

- The value for `ENABLE_SNAPSHOT_CLEANUP` in the `azurestack.conf` file specifies when to clean up stale snapshots. You can use the following values:
 - 0
Do not delete the stale snapshots. This is the default value.
 - 1
Delete the stale snapshots after the backup job completes.
 - 2
Delete the stale snapshots as part of the next discovery job.
- The value for `SNAPSHOT_CLEANUP_MIN` in the `azurestack.conf` file specifies the time in minutes after which the snapshot is deleted. The default value is 1440 minutes (24 hours). The ideal value is the time difference in minutes between two backup job.

All the snapshots that are created as a part of backup job are deleted in the same backup job. However, due to some reasons like network failure, power outage, etc. the snapshot can remain in the Azure Stack environment. In such scenarios you can use the `ENABLE_SNAPSHOT_CLEANUP` parameter and set it to 1.

You can identify the snapshots that NetBackup creates using the following naming convention:

OSDisk

```
<Disk UUID_OSDisk_9999_<Snap id>_ BDAZSNAP>
```

For example:

```
/cb9b5b6d-0c60-4ed1-9cd7-7aaf054b899f_OSDisk_9999_1600405027_BDAZSNAP
```

DataDisk

```
<Disk UUID_DataDisk_1_<Snap id>_ BDAZSNAP>
```

For example:

```
/cb9b5b6d-0c60-4ed1-9cd7-7aaf054b899f_DataDisk_1_1600405027_BDAZSNAP
```


Performing backups and restores of Microsoft Azure Stack

This chapter includes the following topics:

- About backing up Microsoft Azure virtual machines
- About restoring Microsoft Azure Stack virtual machines
- About the restore scenarios for Microsoft Azure Stack VMs from the BAR interface
- Using the BAR interface to restore an Microsoft Azure Stack VM at the same location
- Using the `bprestore` command to restore Microsoft Azure Stack VM at the same location
- Using the BAR interface to restore an Microsoft Azure Stack VM to different location
- Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location
- Using the `bprestore` command to restore Microsoft Azure VM with modified metadata to an alternate location
- Using the `bprestore` command to restore Microsoft Azure Stack VM with modified metadata to an alternate region

About backing up Microsoft Azure virtual machines

You can either schedule a backup job or run a backup job manually. See, NetBackup Administrator's Guide, Volume I

For overview of the backup process, See "Backing up Microsoft Azure Stack VMs" on page 9.

The backup process comprises of the following stages:

1. Pre-processing: In the pre-processing stage, the first backup host that you have configured with the BigData policy, triggers the discovery. At this stage, the VMs and associated metadata is discovered for backup.
2. Data transfer: During the data transfer process, one child job is created for each backup host.

About restoring Microsoft Azure Stack virtual machines

Use the **NetBackup, Backup, Archive, and Restore** console to manage restore operations.

Table 4-1 Restoring Microsoft Azure data

Task	Reference
Understanding the restore process	See "Restoring Microsoft Azure Stack VMs" on page 10.
Understanding the restore scenarios	See "About the restore scenarios for Microsoft Azure Stack VMs from the BAR interface" on page 35. See "Considerations for Microsoft Azure Stack VM restore and recovery" on page 37.
Restoring Microsoft Azure Stack VM at the same location	<ul style="list-style-type: none">■ Restore Wizard See "Using the BAR interface to restore an Microsoft Azure Stack VM at the same location" on page 39.■ Command line interface See "Using the <code>bprestore</code> command to restore Microsoft Azure Stack VM at the same location" on page 41.

Table 4-1 Restoring Microsoft Azure data (continued)

Task	Reference
Restoring Microsoft Azure Stack VM to an alternate location	<ul style="list-style-type: none"> ■ Restore Wizard See “Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location ” on page 45. ■ Command line interface See “Using the bprestore command to restore Microsoft Azure VM with modified metadata to an alternate location” on page 53. See “Using the bprestore command to restore Microsoft Azure Stack VM with modified metadata to an alternate region” on page 56.
Migrating Azure Stack unmanaged disk VM to managed disk VM	See “Conversion of Unmanaged disk VM to Managed disk VM” on page 52.
Restoring Azure Stack VM to a different location	See “Using the BAR interface to restore an Microsoft Azure Stack VM to different location” on page 43.

About the restore scenarios for Microsoft Azure Stack VMs from the BAR interface

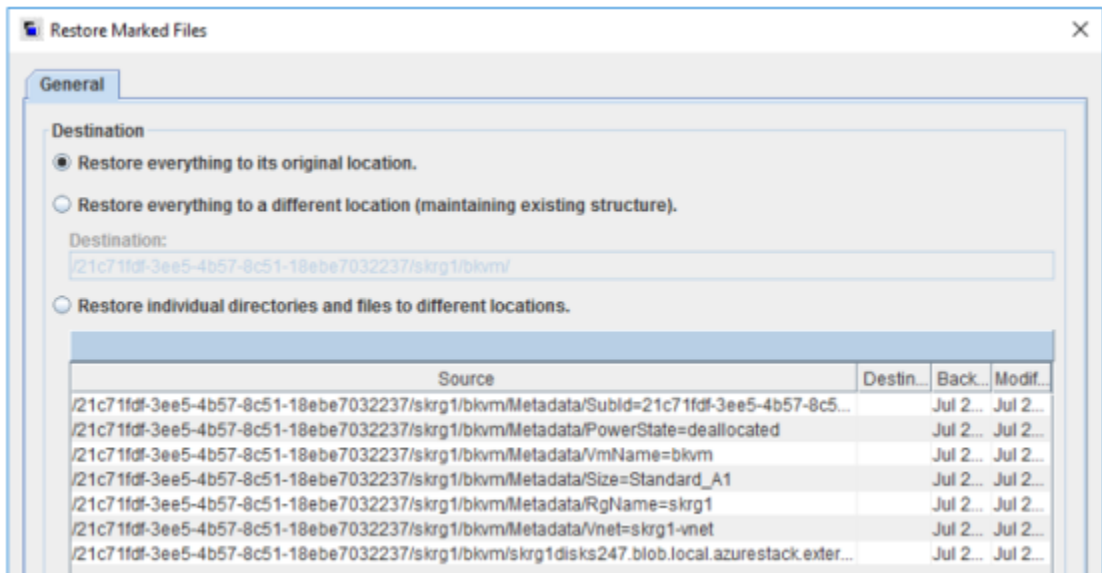
For restoring Microsoft Azure Stack VMs from the **Backup, Archive, and Restore** interface, the following scenarios are possible:

Table 4-2 Options for VM restore

Scenario	Option from the Restore Marked Files dialog box
Restore the Microsoft Azure Stack VM with the existing configuration to the same location (subscription ID and resource group)	Restore everything to its original location.
Restore the Microsoft Azure Stack VM with the existing configuration to an alternate location (subscription ID and resource group)	Restore everything to a different location (maintaining existing structure).
Restore the Microsoft Azure Stack VM with modified configuration (includes VM metadata and location)	Restore individual directories and files to different locations.

The options are available after you enter the details in the **Backup, Archive, and Restore** interface and proceed to the **Restore Marked Files** dialog box.

Figure 4-1 Restore options from the Restore Marked Files dialog box



Restore workflows

The following diagrams give an overview of the restore workflow:

Figure 4-2 Restore the Azure Stack VMs to the original location

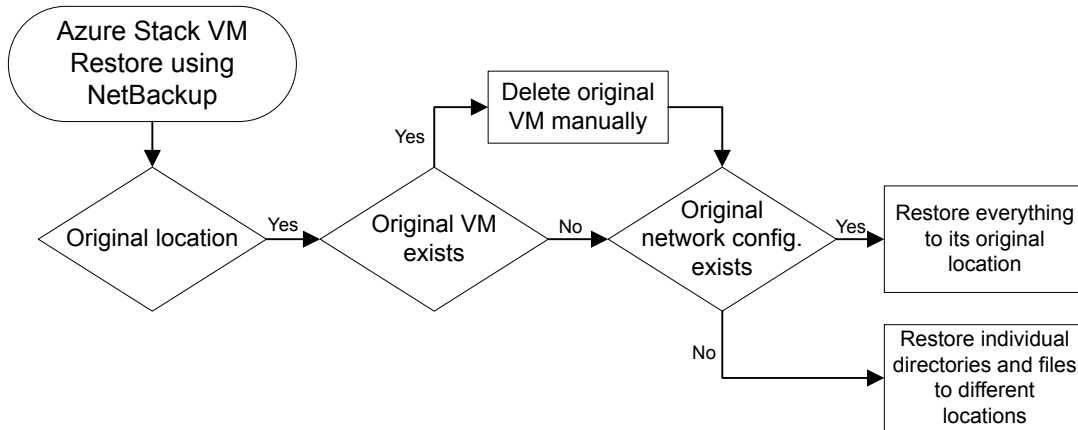
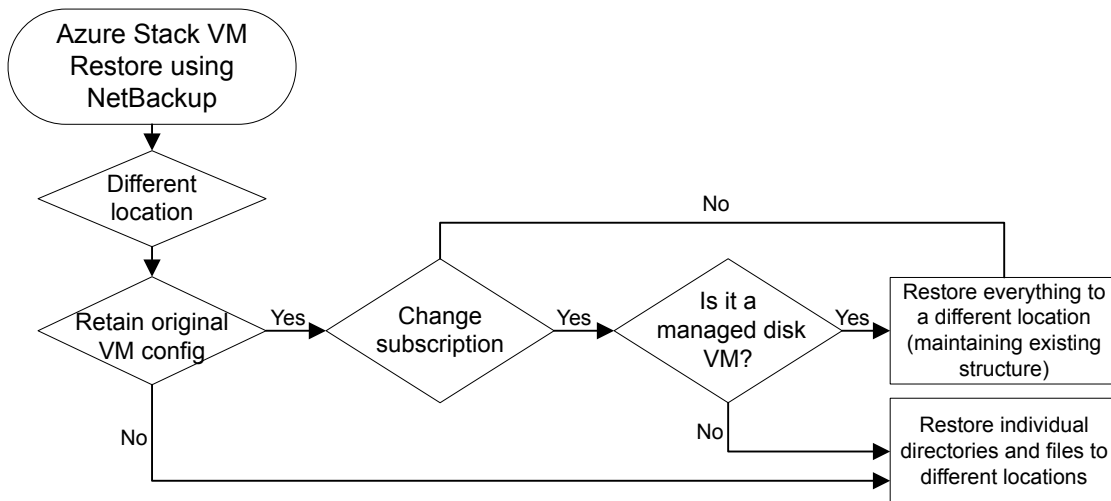


Figure 4-3 Restore the Azure Stack VMs to a different location



Considerations for Microsoft Azure Stack VM restore and recovery

- NetBackup triggers the VM data restore process and if the operation is successful, NetBackup displays the status as successful. Use the Azure Stack portal to monitor the VM creation process.

- If a VM recovery operation fails, you must remove the resources that are created during the restore manually. The resources can include IP address, NIC, OS, and Data disks. Resources can be identified by suffix `-restorenic-`, `-restoreip-`, `-RESTORE-` followed by timestamp.
- You cannot restore a VM with the same name to its original location if the VM still exists. Also Vnet, NSG of the backed up VM needs to exist for VM recovery to succeed.
- To recover a VM, the NetBackup role must have access to the specified subscription and resource group.
- NetBackup can recover the following VM properties:
 - Tags
 - OS Boot Diagnostic Settings
- For any other properties or configuration settings, you must apply them manually after the recovery is done.
- During a recovery, the host name does not change and it remains the same as the backed-up VM. You must log on to the VM and use the OS commands to change the host name.
- When you restore at the original location, a new network configuration is created. One NIC is created and is attached to that virtual network to which the VM was connected during the backup. This step results in the change of the MAC and IP addresses.
- When you want to update the configuration during a VM recovery operation, you must select all the metadata files by double clicking one of the metadata files. You can rename files corresponding to configuration you want to change.
- When you want to update the configuration during a VM recovery operation, you can specify the resource group or networks security group that belong to a different resource group than the VM as follows:

```
Vnet=<ResourceGroup_Name>/<virtual_network_Name>  
Nsg=<ResourceGroup_Name>/<NetworkSecurityGroup_Name>
```

If `ResourceGroup_Name` is not specified and the virtual network or `NetworkSecurityGroup` name is same as the backed-up VM, virtual network or `NetworkSecurityGroup` of the backup time is used during the recovery operation. Otherwise, the specified virtual network is considered to be present in the same resource group as the VM.
- Conversion of unmanaged disk VM to managed disk VM is supported but not vice versa i.e. managed disk VM cannot be converted to unmanaged disk VM.

- Restore from imported images is not supported through BAR GUI. One needs to use `bprestore` command for doing restore.
- Alternate region restore is supported through `bprestore` command only.

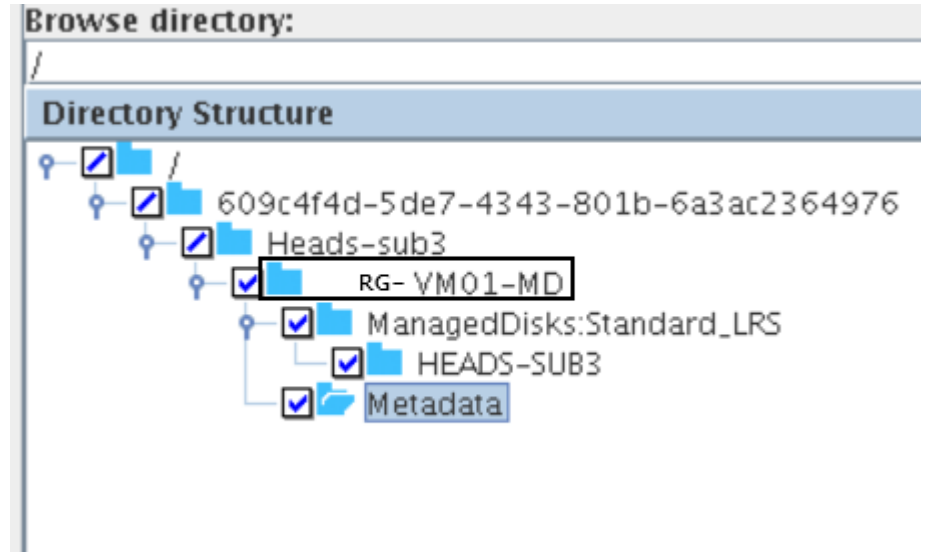
Using the BAR interface to restore an Microsoft Azure Stack VM at the same location

This topic describes how to use the NetBackup Admin console's BAR interface to restore Microsoft Azure Stack on the same Microsoft Azure Stack.

To use the NetBackup Admin console's BAR interface to perform a restore

- 1 Open the **Backup, Archive, and Restore** interface.
- 2 On the **Specify NetBackup Machines and Policy Type** wizard, enter the source and destination details for restore.
 - Specify the Microsoft Azure Application Endpoint as the source for which you want to perform the restore operation.
From the **Source client for restores** list, select the required Application server.
 - Specify the backup host as the destination client.
From the **Destination client for restores** list, select the required backup host.
 - On the **Specify NetBackup Machines and Policy Type** wizard, enter the policy type details for restore.
From the **Policy type for restores** list, choose **BigData** as the policy type for restore.
Click **Ok**.
- 3 Select the appropriate date range to restore the complete data set.
- 4 In the **Browse** directory, specify the root directory ("/") as the path to browse.
- 5 From the File menu (Windows) or Actions menu (UNIX), choose **Specify NetBackup Machines and Policy Type**.
- 6 Go to the **Backup History** and select the backup images that you want to restore.
- 7 In the **Directory Structure** pane, expand the **Directory**.
All the subsequent files and folders under the directory are displayed in the **Contents of Selected Directory** pane.

- 8 In the **Contents of Selected Directory** pane, select the check box for the Microsoft Azure VMs that you want to restore.



- 9 Click **Restore**.
- 10 In the **Restore Marked Files** dialog box, select the destination for restore as per your requirement.
 - Select **Restore everything to its original location** to restore your files to the same location where you performed your backup.

Note: For more information about the restore scenarios, See “About the restore scenarios for Microsoft Azure Stack VMs from the BAR interface” on page 35.

- 11 Click **Start Restore**.
- 12 Verify that the VM gets restored and instantiated.
- 13 After the VM is restored, open the **Microsoft Azure Stack Admin Portal** and assign the VM network interface to the required network security group.

Using the `bprestore` command to restore Microsoft Azure Stack VM at the same location

You can use the `bprestore` command to restore a Microsoft Azure Stack VM in the same resource group.

To restore Microsoft Azure data on the same location as your backup location

- 1 Log on as an Administrator or root user to the NetBackup master server on windows or UNIX system respectively.
- 2 Run the following command on the NetBackup master server by providing appropriate values:

```
bprestore -S master_server -D backup_host -C client -t 44 -X -s  
<bktime> -e <bktime> -L progress_log -f listfile | filenames  
"/subscription ID/resource group/VmName"
```

Where,

```
-S master_server
```

Specifies the name of the NetBackup master server.

```
-D backup host
```

Specifies the name of the backup host.

```
-C client
```

Specifies a configuration server as a source to use for finding backups or archives from which to restore files. This name must be as it appears in the NetBackup catalog.

```
-f listfile
```

Specifies a file (`listfile`) that contains a list of files to be restored and can be used instead of the file names option (`filenames`). In `listfile`, list each file path must be on a separate line.

```
-L progress_log
```

Specifies the name of whitelisted file path in which to write progress information.

```
-t 44
```

Specifies BigData as the policy type.

```
"/subscription ID/resource group/VmName"
```

Specifies the Microsoft Azure Stack VM that you want to restore.

```
-X -s bktime -e date
```

Start and end date for backup image selection. with X option one could specify timestamp instead of human readable format, Refer to `bprestore` command in the Commands Reference Guide.

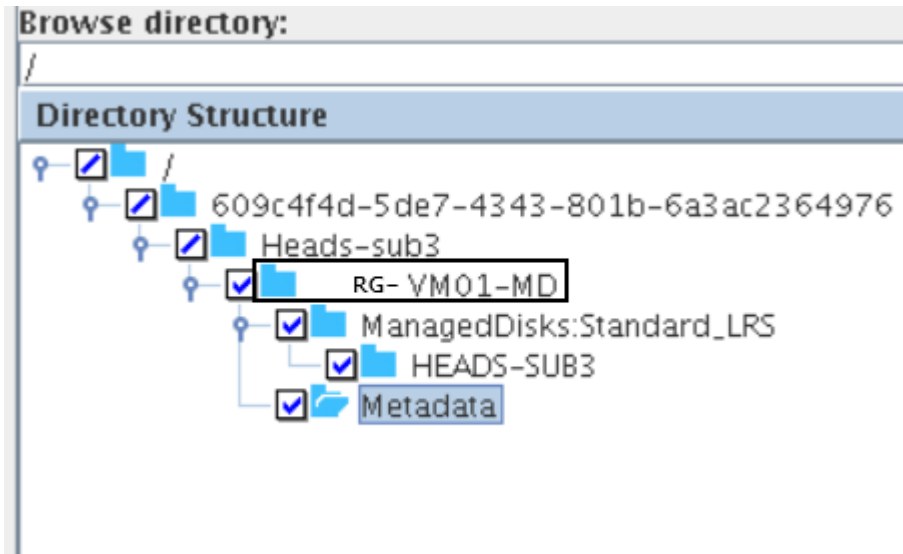
Using the BAR interface to restore an Microsoft Azure Stack VM to different location

This topic describes how to use the NetBackup Admin console's BAR interface to restore Microsoft Azure Stack on the same Microsoft Azure Stack but to different RG or Subscription.

To use the NetBackup Admin console's BAR interface to perform a restore

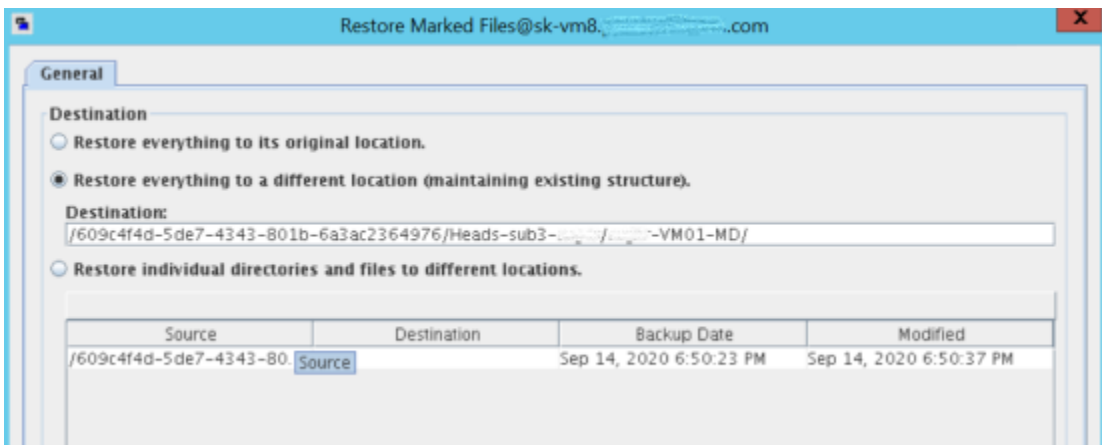
- 1 Open the Backup, Archive, and Restore interface.
- 2 On the Specify NetBackup Machines and Policy Type wizard, enter the source and destination details for restore.
 - Specify the Microsoft Azure Application Endpoint as the source for which you want to perform the restore operation. From the Source client for restores list, select the required Application server.
 - Specify the backup host as the destination client. From the Destination client for restores list, select the required backup host. Restore is faster if the backup host is the media server that had backed up the VM.
 - On the Specify NetBackup Machines and Policy Type wizard, enter the policy type details for restore. From the Policy type for restores list, choose BigData as the policy type for restore. Click Ok.
- 3 Select the appropriate date range to restore the complete data set.
- 4 In the Browse directory, specify the root directory ("/") as the path to browse.
- 5 From the File menu (Windows) or Actions menu (UNIX), choose Specify NetBackup Machines and Policy Type.
- 6 Go to the Backup History and select the backup images that you want to restore.
- 7 In the Directory Structure pane, expand the Directory. All the subsequent files and folders under the directory are displayed in the Contents of Selected Directory pane.

- 8 In the Contents of Selected Directory pane, select the check box for the Microsoft Azure VMs that you want to restore.



- 9 Click Restore.
- 10 In the Restore Marked Files dialog box, select the destination for restore as per your requirement.

Select Restore everything to different location to restore VM to the different RG/Subscription.



Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location

VM path is of following format /<subId>/<RgName>/<VmName>.

This option lets you:

- Changing of the VM Name where restored VM will be restored to original location (i.e. same Subscription and RG) but with different name.
- Changing RG keeping target subscription same. In this case only VMs RG is changed all the configuration including networking configuration of the VM remains the same.
- Changing target Subscription & RG – this is supported only for managed disk VM restore.
 - Network configuration of the managed disk VM restored to different target subscription.
 - If Vnet & NSG exists in the target RG they are used while creating NIC. If No Vnet, NSG found in target RG. Same is searched for in different RG in the target subscription.

Note: For more information about the restore scenarios, See "About the restore scenarios for Microsoft Azure Stack VMs from the BAR interface".

11 Click Start Restore.

12 Verify that the VM gets restored and instantiated.

Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location

NetBackup lets you restore Microsoft Azure Stack VM to another resource group or modify the VM metadata and restore to the same resource group. This type of restore method is also referred to as redirected restores.

This topic describes how to use the NetBackup Admin console's BAR interface to restore an Microsoft Azure Stack VM with modified metadata at an alternate location or another resource group on Microsoft Azure Stack.

To use the NetBackup Admin console's BAR interface to perform a restore

- 1 Open the **Backup, Archive, and Restore** interface.
- 2 On the **Specify NetBackup Machines and Policy Type** wizard, enter the source and destination details for restore.

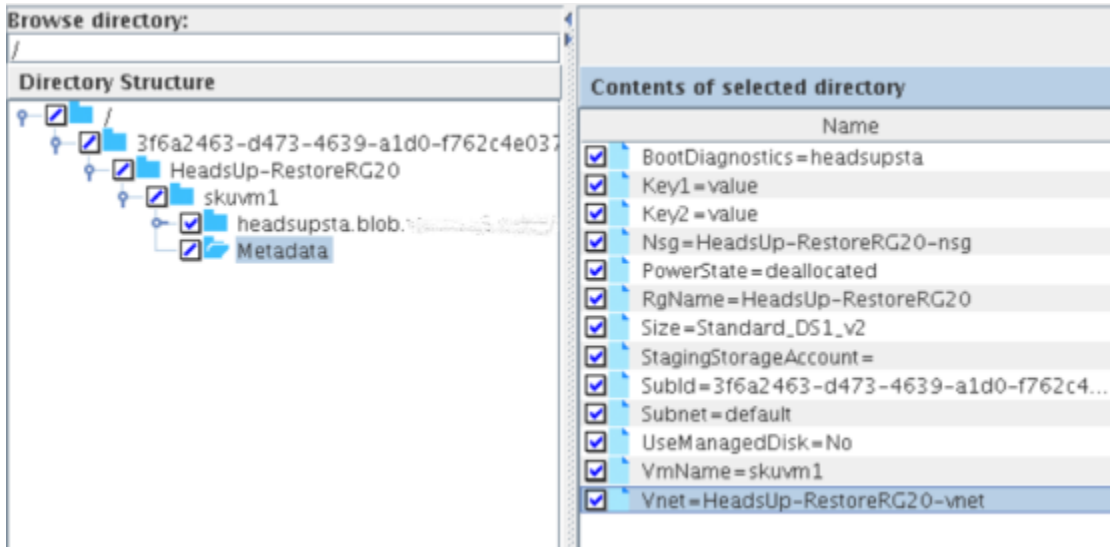
Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location

- Specify the Microsoft Azure Application Endpoint as the source for which you want to perform the restore operation.
From the **Source client for restores** list, select the required Application server.
 - Specify the backup host as the destination client.
From the **Destination client for restores** list, select the required backup host.
 - On the **Specify NetBackup Machines and Policy Type** wizard, enter the policy type details for restore.
From the **Policy type for restores** list, choose **BigData** as the policy type for restore.
Click **Ok**.
- 3 Select the appropriate date range to restore the complete data set.
 - 4 In the **Browse** directory, specify the root directory ("/") as the path to browse.
 - 5 From the File menu (Windows) or Actions menu (UNIX), choose **Specify NetBackup Machines and Policy Type**.
 - 6 Go to the **Backup History** and select the backup images that you want to restore.
 - 7 In the **Directory Structure** pane, expand the **Directory**.
All the subsequent files and folders under the directory are displayed in the **Contents of Selected Directory** pane.

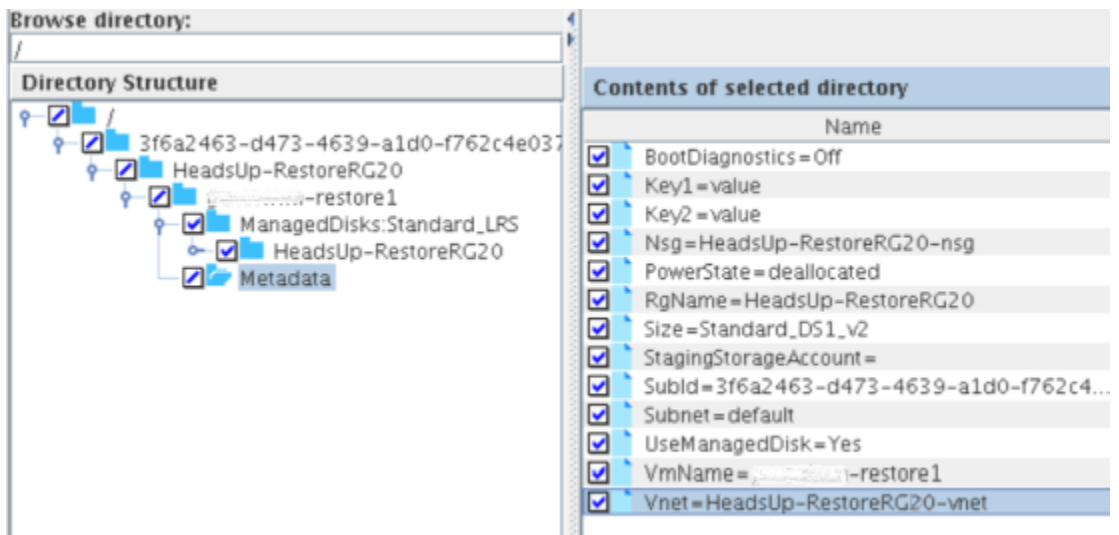
Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location

- 8 Select the VM that you want to restore. Ensure that the storage account directories are selected.

For example:



For Managed disk VM, select all the disks under ManagedDisks.



If you wish to change the disk property, for example, disk RG, disk name select individual disks file.

Note: You must select all the metadata files and rename or modify only the metadata that you want to change.

Select the **Metadata** folder and unselect and select any metadata in the contents of selection column so that all the options appear in the restore marked file dialog box.

Using the BAR interface to restore an Microsoft Azure Stack VM with modified metadata at a different location

- 9 Click the selected **Metadata** directory, and in the **Contents of Selected Directory** pane, select the metadata that you want to modify.

You can modify the following metadata:

Metadata or property	Description	Default value	Valid Value
VmName	Name of the VM.	Name of the VM during backup.	Valid VM Name that is unique in the resource group.
PowerState	State of the VM after restore.	Power state of the VM during backup.	Stopped, Deallocate, or Running
Size	Size of the VM in the Microsoft Azure Stack recommended format. For more information, refer to Azure Stack VM Sizes.	Size of the VM during backup.	Valid VM size.
Vnet	The virtual network that contains the VM.	Vnet of the VM during backup.	Virtual network in the target subscription. If empty value is specified, for example, <code>Vnet</code> , <code>RgName-vnet</code> if exists in target RG for VM is used.
Nsg	The Network Security Group of the VM.	Nsg of the VM during backup.	NSG in the target subscription. If empty value is specified, for example, <code>Nsg</code> , Network security group is not set for the VM.
RgName	Location or the resource group of the Microsoft Azure Stack VM.	Resource Group of the VM during backup.	Resource group that is a part of the target subscription.
Storage Account	The storage account that contains the unmanaged VM disks. This is part of the VHD path and not a separate metadata file.	Storage account of the VM during backup.	Valid storage account that is a part of the target subscription.
SubId	Microsoft Azure Stack subscription ID.	Subscription ID of the VM during backup.	Subscription ID that a NetBackup role can access.

Metadata or property	Description	Default value	Valid Value
Boot Diagnostics	<p>Boot Diagnostics settings.</p> <p>If the boot diagnostics storage account that is specified during restore does not exist in the destination subscription or Azure Stack then boot diagnostics settings are disabled</p>	Boot diagnostics settings of the unmanaged VM during backup.	Valid storage account that is part of the target subscription.
UseManagedDisk	VM has managed disk or unmanaged disks	Managed disks setting for the backed up VM.	<p>Yes for VM with managed disks.</p> <p>No for VM with unmanaged disks.</p>
StagingStorageAccount	Storage account to be used for creating temporary VHD file during a managed disk VM restore.	Empty	Valid storage account that is part of the target subscription.
Storage Account Type	<p>Managed disk storage type. This is part of managed disk path and not Metadata file.</p> <p>If <code>StorageAccountType</code> is not detected during backup. User needs to specify it during restore</p>	Storage Account Type of the backed up managed disk.	Supported Storage Account Type string by Azure Stack, for example, <code>Standard_LRS</code> , <code>Premium_LRS</code> .
Disk RG	Resource Group of the Managed disk. This is part of managed disk path and not Metdata file.	Resource group of the backed up Managed disk	Resource group that is part of the target subscription.

10 Click **Restore**.

- 11** In the **Restore Marked Files** dialog box, select **Restore individual directories and files to different locations**.

Note: For more information about the restore scenarios, See “About the restore scenarios for Microsoft Azure Stack VMs from the BAR interface” on page 35.

For every metadata value that you want to change, select the value, click **Change Selected Destination(s)**, and in the **Destination** field modify the metadata value at the end of the URL.

For example, if you want to change the `VMName`, change:

```
/21c71fdf-3ee5-4b57-8c51-18ebe7032237/SKRG/MSvm15/Metadata/VMName=OldVMName  
to
```

```
/21c71fdf-3ee5-4b57-8c51-18ebe7032237/SKRG/MSvm15/Metadata/VMName=NewVMName
```

Here, `VMName` is the key and `OldVMName` is the value. The metadata and its value have the `Key=Value` format. You must modify the value of all the metadata that you want to change.

Changing managed disk properties

Change

```
/3f6a2463-d473-4639-a1d0-f762c4e0371a/HeadsUp-RestoreRG20/skvm1/  
ManagedDisks:Standard_LRS/HeadsUp-RestoreRG20/  
skvm1_OsDisk_1_be667512b2a44ecfb1ffa43506aaa48c-RESTORE-1598299429
```

to

```
/3f6a2463-d473-4639-a1d0-f762c4e0371a/HeadsUp-RestoreRG20/skvm1/  
ManagedDisks:Premium_LRS/HeadsUp-RestoreRG30/osdisk_1
```

Changing unmanaged disk storage account

Change

```
/3f6a2463-d473-4639-a1d0-f762c4e0371a/HeadsUp-RestoreRG20/skvm1/  
headsupsta.blob.tasz.vxi.vwx.com/vhds/skvm1-UMD-RESTORE-1599155872.vhd
```

to

```
/3f6a2463-d473-4639-a1d0-f762c4e0371a/HeadsUp-RestoreRG20/skvm1/  
stasub2testrgg.blob.tasz2.vxi.vwx.com/vhds/skvm1-UMD-RESTORE-1599155872.vhd
```

Note: For the VM size metadata, specify the modified value in the Microsoft Azure Stack recommended format. The new VM size must be part of your subscription.

For more information, refer to
<https://docs.microsoft.com/en-us/azure/azure-stack/user/azure-stack-vm-sizes>.

- 12 Click **Start Restore**.
- 13 Use the **Azure Stack Admin Portal** to view the VM creation process.

Conversion of Unmanaged disk VM to Managed disk VM

When using BAR GUI, need to use option 3 "Restore individual directories and files to different location" which allows changing the configuration during restore.

Rename following file entries:

```
/d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/Metadata/UseManagedDisk=no
=>
/d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/Metadata/UseManagedDisk=yes

/d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/
nbusrgdisks439.blob.mtcazs.wwtatc.com/vhds/nbu-an-um-vm1120200625233142.vhd
=>
/d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/
ManagedDisks:Standard_LRS/NBU-SamRG2/osdisk
```

Note: Storage account types in lower case is not supported by disk create API and needs to be specified as Standard_LRS, Premium_LRS, etc.

Conversion of Unmanaged disk VM backed up using old plugin to Managed disk VM

When using BAR GUI, need to use option 3 "Restore individual directories and files to different location" which allows changing the configuration during restore.

Rename following file entries:

```
d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/Metadata/Key1=value
=>
/d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/Metadata/UseManagedDisk=yes

/d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/
nbusrgdisks439.blob.mtcazs.wwtatc.com/vhds/nbu-an-um-vm1120200625233142.vhd
=>
/d182ecd4-23c5-4fde-90f3-8951146a385e/NBU-SRG/nbu-an-um-vm11/
ManagedDisks:Standard_LRS/NBU-SRG2/osdisk
```

Note: Staging storage account cannot be specified using `rename` file and needs to be specified as part of plugin configuration file.

Storage account type in lower case is not supported by disk create API and needs to be specified as `Standard_LRS`, `Premium_LRS`, etc.

Using the `bprestore` command to restore Microsoft Azure VM with modified metadata to an alternate location

1. Run the following `bplist` command to list the files:

```
bplist -S master_server -C configuration_server_01 -unix_files
-R 3 -t 44 -X -s <bktime> -e <bktime>
"/21c71fdf-3ee5-4b57-8c51-18ebe7032237/skrg1/bkvm15" > listfile
```

Open the `listfile` in the editor and remove the special characters at the end of the line. Remove all the paths that end with forward slash `/` that corresponds to directories.

2. Run the following command on the NetBackup master server using the modified values for the mentioned parameters in step 1.

```
bprestore -S master_server -D backup_host -C client -R rename_file
-t 44 -L progress log -f listfile | filenames
```

Where,

```
-S master_server
```

Specifies the name of the NetBackup master server.

```
-D backup host
```

Specifies the name of the backup host.

```
-C client
```

Specifies a configuration server as a source to use for finding backups or archives from which to restore files. This name must be as it appears in the NetBackup catalog.

```
-f listfile
```

Specifies a file (`listfile`) that contains a list of files to be restored and can be used instead of the file names option (`filenames`). In `listfile`, list each file path must be on a separate line.

```
-L progress_log
```

Specifies the name of whitelisted file path in which to write progress information.

```
-t 44
```

Specifies BigData as the policy type.

```
-R rename_file
```

Specifies the name of a file with name changes for alternate-path restores.

For example:

```
bprestore.exe -S master_server_01 -D backup_host_01 -C  
configuration_server_01 -t 44 -L "<install_dir>\logs\restore.log"  
-R "<install_dir>\renam_file_path\restore.chg" -f listfile
```

Sample rename files

Convert unmanaged disk to managed disk

```
change  
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/  
headsrq4disks2.blob.vtasazs1.vxi.ver.com/vhds/sr-VM02-UMD20200826032144.vhd  
to  
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/  
ManagedDisks:Standard_LRS/Heads-RestoreRG20/osdisk1
```

Convert unmanaged disk VM to managed disk VM to the same RG

```
change  
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/  
Metadata/VmName=sr-VM02-UMD  
to  
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/  
Metadata/VmName=sr-VM02-UMD-2md2  
  
change  
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/  
Metadata/UseManagedDisk=No  
to  
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/  
Metadata/UseManagedDisk=Yes  
  
change  
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/  
headsuprg4disks2.blob.vtasazs1.vxi.ver.com/vhds/  
sr-VM02-UMD20200826032144.vhd
```

Using the bprestore command to restore Microsoft Azure VM with modified metadata to an alternate location

```
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
ManagedDisks:Standard_LRS/HEADS-SUB3-SR/osdisk1
```

Convert unmanaged disk VM to a managed disk VM with a different subscription

```
change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/Vnet=Heads-sub3-sr-vnet
```

```
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/Vnet=Heads-RestoreRG20-vnet
```

```
change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/VmName=sr-VM02-UMD
```

```
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/VmName=sr-VM02-UMD-2md2
```

```
change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/UseManagedDisk=No
```

```
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/UseManagedDisk=Yes
```

```
change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/SubId=609c4f4d-5de7-4343-801b-6a3ac2364976
```

```
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/SubId=3f6a2463-d473-4639-a1d0-f762c4e0371a
```

```
change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/Size=Standard_DS1_v2
```

```
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/Size=Standard_DS2_v2
```

```
change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/RgName=Heads-sub3-sr
```

```
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/RgName=Heads-RestoreRG20

change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/Nsg=sr-VM02-UMD-nsg
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/Nsg=HeadsUp-RestoreRG20-nsg

change
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/BootDiagnostics=headsuprg4diag398
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
Metadata/BootDiagnostics=headsupsta

change /609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
headsuprg4disks2.blob.vtasazs1.vxi.ver.com/vhds/
sr-VM02-UMD20200826032144.vhd
to
/609c4f4d-5de7-4343-801b-6a3ac2364976/Heads-sub3-sr/sr-VM02-UMD/
ManagedDisks:Standard_LRS/HeadsUp-RestoreRG20/osdisk1
```

Note: Ensure that you have whitelisted all the file paths such as `<rename_file_path>`, `<progress_log_path>` that are not included as a part of NetBackup install path.

Using the `bprestore` command to restore Microsoft Azure Stack VM with modified metadata to an alternate region

NetBackup lets you restore Microsoft Azure Stack data to another resource group and modify the metadata. This type of restore method is also referred to as redirected restores.

Using the bprestore command to restore Microsoft Azure Stack VM with modified metadata to an alternate region

To perform redirected restore for Microsoft Azure Stack

- 1 Modify the values for *rename_file* and *listfile* as follows:

Parameter	Value
<i>rename_file</i>	<p>Add <code>ALT_APPLICATION_SERVER=</code> entry to specify ARM endpoint of the alternate region</p> <p>To update the <code>VmName</code> metadata, add:</p> <p>change</p> <pre>/21c71fdf-3ee5-4b57-8c51-18ebe7032237/SKRG/MSvm15 /Metadata/VmName=OldVmName</pre> <p>to</p> <pre>/21c71fdf-3ee5-4b57-8c51-18ebe7032237/SKRG/MSvm15/ Metadata/VmName=NewVmName</pre> <p>To change the power state of the VM, add:</p> <p>change</p> <pre>/21c71fdf-3ee5-4b57-8c51-18ebe7032237/SKRG/MSvm15 /Metadata/PowerState=running</pre> <p>to</p> <pre>/21c71fdf-3ee5-4b57-8c51-18ebe7032237/SKRG/MSvm15 /Metadata/PowerState=deallocate</pre> <p>The file paths must start with / (slash).</p> <p>Add a new entry for all the metadata options that you want to modify.</p> <p>Note: For the VM size metadata, specify the modified value in the Microsoft Azure Stack recommended format. The new VM size must be part of your subscription.</p> <p>For more information, refer to Virtual machine sizes supported in Azure Stack.</p>
<i>listfile</i>	List of all the Microsoft Azure files to be restored

Using the `bprestore` command to restore Microsoft Azure Stack VM with modified metadata to an alternate region

- 2 To fetch the credentials information for the alternate Azure Stack:

Add a `tpconfig` entry for the new Azure Stack ARM end point.

Change the name of the generated encrypted file in `/usr/openv/var/global` to match the name of the source client encrypted credentials file.

For example, if `arm-end-point1.conf` is the encrypted source client file and `arm-end-point2.conf` is the encrypted alternate client file. Make a copy of `arm-end-point1.conf` to `arm-end-point1.conf_org` and then copy `arm-end-point2.conf` to `arm-end-point1.conf` before running the `bprestore` command.

- 3** List of files can be obtained by running the following `bplist` command:

```
bplist -S master_server -C configuration_server_01 -unix_files  
-R 3 -t 44 -X -s <bktime> -e <bktime>  
"/21c71fdf-3ee5-4b57-8c51-18ebe7032237/skrg1/bkvm15" > listfile
```

Open the `listfile` in the editor and remove the special characters at the end of the line. Remove all the paths ending with `/` that correspond to directories.

- 4** Run the following command on the NetBackup master server using the modified values for the mentioned parameters in step 1.

```
bprestore -S master_server -D backup_host -C client -R rename_file  
-t 44 -X -s bktime -e bktime -L progress_log -f listfile |  
filenames
```

Where,

`-S master_server`

Specifies the name of the NetBackup master server.

`-D backup_host`

Specifies the name of the backup host.

`-C client`

Specifies a configuration server as a source to use for finding backups or archives from which to restore files. This name must be as it appears in the NetBackup catalog.

`-f listfile`

Specifies a file (`listfile`) that contains a list of files to be restored and can be used instead of the file names option (`filenames`). In `listfile`, list each file path must be on a separate line.

`-L progress_log`

Specifies the name of whitelisted file path in which to write progress information.

`-t 44`

Specifies BigData as the policy type.

`-R rename_file`

Specifies the name of a file with name changes for alternate-path restores.

```
ALT_APPLICATION_SERVER=management.vtsz2.vxi.vs.com
```

```
change /3f6a2463-d473-4639-ald0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/  
to /3f6a2463-d473-4639-ald0-f762c4e0371a/HeadsUp-RestoreRG20/skuvml/Metadata/
```

```
change /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/Nsg=
to /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/Nsg=rs-m

change /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/RgNa
to /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/RgName=r

change /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/
SubId=3f6a2463-d473-4639-a1d0-f762c4e0371a
to /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/
SubId=b326ed58-7537-4c81-b2ac-5b16d6a524b3

change /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/VmNa
to /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/VmName=s

change /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/Vnet
to /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/Metadata/Vnet=rsh

change /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/HUpsta.blob.v
skuvml-UMD-RESTORE-1599155872.vhd
to /3f6a2463-d473-4639-a1d0-f762c4e0371a/HUp-RestoreRG20/skuvml/stasub2testrgg.bl
vhds/skuvml-UMD-RESTORE-1599155872.vhd
```

Troubleshooting

This chapter includes the following topics:

- About NetBackup for Microsoft Azure debug logging
- Known limitations for Microsoft Azure protection using NetBackup
- Backup fails with error 6662
- Backup fails with error 6661
- Backup fails with error 6646
- Backup fails with error 6629
- Backup fails with error 6626
- Backup fails with error 6630
- Restore fails with error 2850
- Backup fails with error 1
- Adding Azure Stack credentials to NetBackup fails with error 9101
- Adding Azure Stack credentials to NetBackup fails with error 7610

About NetBackup for Microsoft Azure debug logging

NetBackup maintains process-specific logs for the various processes that are involved in the backup and restore operations. Examining these logs can help you to find the root cause of an issue.

These log folders must already exist in order for logging to occur. If these folders do not exist, you must create them.

Known limitations for Microsoft Azure protection using NetBackup

The log folders reside on the following directories

- On Windows: `install_path\NetBackup\logs`
- On UNIX or Linux: `/usr/opensv/netbackup/logs`

Table 5-1 NetBackup logs related to Microsoft Azure

Log Folder	Messages related to	Logs reside on
<code>install_path/NetBackup/logs/bpVMutil</code>	Policy configuration	Master server
<code>install_path/NetBackup/logs/nbaapidiscv</code>	BigData framework, discovery, and Microsoft Azure configuration file logs	Backup host
<code>install_path/NetBackup/logs/bpbrm</code>	Policy validation, backup, and restore operations	Media server
<code>install_path/NetBackup/logs/bpbkar</code>	Backup	Backup host
<code>install_path/NetBackup/logs/tar</code>	Restore and Microsoft Azure configuration file	Backup host

For more details, refer to the NetBackup Logging Reference Guide.

Known limitations for Microsoft Azure protection using NetBackup

The following table lists the known limitations for Microsoft Azure protection using NetBackup:

Table 5-2 Known limitations

Limitation	Workaround
Unable to obtain the authentication token after NetBackup upgrades from 8.1.2 to 9.0.	The encryption and the decryption algorithm is updated in NetBackup 9.0. Remove all the existing Microsoft Azure servers from NetBackup and add them again after NetBackup upgrades to 9.0.

Table 5-2 Known limitations (*continued*)

Limitation	Workaround
VMs in a provisioning state other than Succeeded are backed up successfully, but their restore can fail because of the inconsistent VM state at the time of the backup.	If the restore fails for any VM, check the log files and check whether the provisioning state of the VM was Succeeded or something else during the time of the backup. If the provisioning status is anything apart from Succeeded, the VM cannot be restored.

Backup fails with error 6662

Backup fails with the following error:

```
(6662) Unable to find the configuration file.
```

Workaround:

Ensure that you have created a credential file, whitelisted the path to the file, and the file path is specified in the `tpconfig` command.

See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.

Backup fails with error 6661

Backup fails with the following error:

```
(6661) Unable to find the configuration parameter.
```

Workaround:

Verify that the right configuration options are added in the credential file that is specified in the `tpconfig` command.

See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.

Backup fails with error 6646

Backup fails with the following error:

```
(6646) Unable to communicate with the server.
```

Workaround:

Run the backup operation again. The error might be because of the Azure Stack being overloaded.

Backup fails with error 6629

Backup fails with the following error:

```
(6629) Unable to complete the operation. Failed to authorize the user or the server.
```

Workaround:

- Validate the configuration options and the values in the credential file.
- Verify the values when you run the `./tpconfig -dappservers` command.
- Verify the values for the Azure Stack user name and password.

See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.

Backup fails with error 6626

Backup fails with the following error:

```
(6626) The server name is invalid.
```

Workaround:

Verify the name of the ARM endpoint.

Backup fails with error 6630

Backup fails with the following error:

```
(6630) Unable to process the request because the server resources are either busy or unavailable. Retry the operation.
```

Workaround:

- Verify the value of the backup selection from the Azure Stack portal.
- Verify the values of the `AuthResource` in the credentials file for the backup selection.
- Verify that you have added the appropriate ARM endpoint in the backup policy and the credentials file for the backup selection.
- Ensure that you have created a custom role for your Azure Stack subscription.

Run the `tpconfig -update` command after you make changes to the credential file.

See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.

Restore fails with error 2850

Restore fails with the following error:

```
(2850) Restore error.
```

Workaround:

Specify a valid and supported VM size.

Backup fails with error 1

Backup fails with the following error:

```
(1) The requested operation was partially successful.
```

The error details also describe the VHDs that are not backed up.

Workaround:

Ensure that the following parameters are configured properly:

- If `FETCH_STORAGE_KEYS=true`, ensure that the NetBackup administrator has permissions for fetching and accessing storage account and access keys for Azure Stack.
- If `FETCH_STORAGE_KEYS=false`, ensure that you have added required storage accounts with the access keys in the credential file.

Run the `tpconfig -update` command after you make changes to the credential file.

See “Adding a Microsoft Azure Stack custom role to provide access permissions to NetBackup administrator” on page 18.

See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.

Adding Azure Stack credentials to NetBackup fails with error 9101

This error occurs if there is a conflict in the double quotes format provided for the file path in the `tpconfig` command.

For example, `application_server_conf "/usr/opensv/var/global/azure.conf"`

Workaround:

Specify the file path without double quotes or enter the double quotes manually in the command prompt.

See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.

Adding Azure Stack credentials to NetBackup fails with error 7610

This error occurs when there is a formatting error in the credentials file.

Workaround:

Check the syntax or formatting in the credentials file.

Run the `tpconfig -update` command after you make changes to the credential file.

See “Adding Microsoft Azure Stack credentials in NetBackup” on page 29.