vSphere Plugin

Bacula Systems Documentation

Contents

Contents

Important: Remember to read the Best Practices chapter common for all of our hypervisor plugins.

This document aims at presenting the reader with information about various techniques and strategies to backup VMware virtual machine using **Bacula Enterprise vSphere Plugin**. It describes the Plugin, defines the scope of its operations, and presents its main features. It also covers how to restore specific files (*Single File Recovery*) or instantly restore VM (*VM Instant Recovery*) from backups made with Bacula Enterprise vSphere Plugin. Make sure you have familiarized yourself with the *VM discovery automation tool* that simplifies Backup Administrator tasks in adding new VMs or managing decommissioned VMs automatically.

1 Scope

This Plugin is available since **Bacula Enterprise** 8.0, and is not applicable to prior versions of Bacula.

It supports the following versions of vSphere:

- ESX/ESXi 5.0, 5.1, 5.5, 6.0, 6.5, 6.7, 7.0 and 8.0
- vCenter 6.0, 6.5, 6.7, 7.0 (Bacula Enterprise version 12.4.3 and higher) and 8.0 (Bacula Enterprise version 16.0.0 and higher)
- Virtual machines: VM hardware version 7 and higher.

The vSphere Plugin has been tested with and is supported on the following Linux platforms:

- RHEL 64bit versions 7, 8 and 9
- CentOS 64bit version 7 (CentOS Stream is not supported by VMWare)
- Oracle Linux version 8
- SLES 64bit version 12.5

See also:

Go to:

- *Features*
- *Backup Strategies*
- *Installation*
- *Configuration*
- *Operations*
- *Limitations*
- *Troubleshooting*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

2 Features

The vSphere Plugin provides virtual machine bare metal recovery, while the backup at the guest level simplifies data protection of critical applications.

The Plugin integrates VMware's Changed Block Tracking (CBT) technology to ensure only blocks that have changed since the initial Full, and/or the last Incremental or Differential backup are sent to the current Incremental or Differential backup stream to give you more efficient backups and reduced network load.

See the features list:

- vSphere Storage APIs – Data Protection based online backups
- VSS-based guest snapshots for quiescing VSS-based applications
- Full, Differential and Incremental image-level backups of virtual machines
- Restores complete virtual machine images
- Restores vmdk files to alternate directory
- Supports both TCP/IP and SAN (FC/iSCSI) VMware datastore access
- NBD (Network Block Device), HotAdd, or SAN access
- The vSphere Plugin is compatible with Copy/Migration jobs. In order for Incremental vSphere Plugin backup jobs to be compatible with Copy Jobs, the devices in the destination Storage Daemon **must** have the setting "Maximum Concurrent Jobs = 1". If this option is not set, then restores from Incremental copied vSphere backup jobs may not be possible. Read the blb:MigrationChapter documentation for more information.

Along with the vSphere Plugin, Bacula Systems provides additional packages which allows:

- *Single Item Restore*
- *Instant Recovery*

from vSphere Plugin backups of VMware VMs. It also integrates with the autodiscovery of new VM with the Scan Plugin

See also:

Go back to:

- *Scope*

Go to:

- *Backup Strategies*
- *Installation*
- *Configuration*
- *Operations*
- *Limitations*
- *Troubleshooting*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

3 Backup Strategies

This article aims at presenting possible VM backup strategies with Bacula Enterprise.

3.1 Image Backup with the vSphere Plugin

With the image backup level strategy, the **Bacula Enterprise** vSphere Plugin will save the virtual machine's disks at the *raw* level, in the VMware/vSphere context. For this to work, you don't need a Bacula File Daemon on each guest VM. You only need one FD with the vSphere Plugin installed. It is recommended that this FD be installed on the same machine as a Storage Daemon (SD) so that data from the ESXi servers traverses a network link only once. The vSphere Plugin will contact your VMware ESXi server to read and save the contents of your virtual machine disks using Network Block Device (NBD), HotAdd, or SAN access. When directly accessing a **vmdk** image stored on your **Datastore**, Bacula doesn't need to walk through the Client Filesystem to open/read/close/stat files, so it consumes fewer resources on your ESXi infrastructure than a backup with a File Daemon on each guest machine would. On the other hand, Bacula will also read and save useless data such as swap files or temporary files.

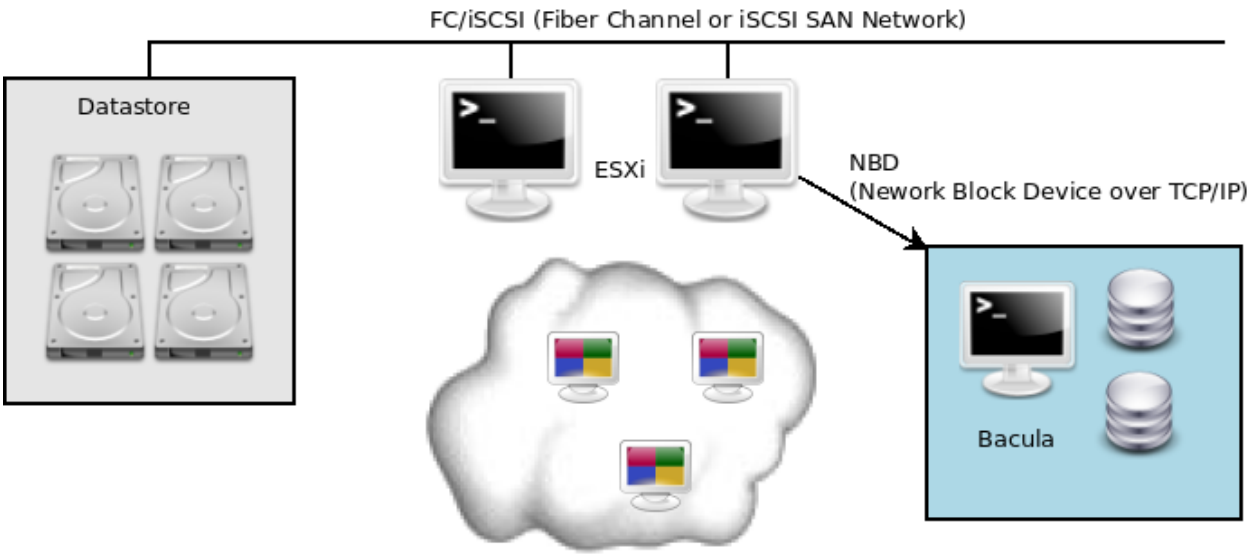

Fig. 1: Backup through TCP Network using NBD (Network Block Device)

When the vSphere Plugin is using the NBD transport method for the backup, the data is streamed to the backup server via the ESXi system's VMkernel port.

The Bacula Enterprise vSphere Plugin can also use your SAN infrastructure to minimize the I/O load on your ESXi servers. Using this method of access, even fewer resources are consumed on the ESXi server, but the data still needs to be read from your datastore, thus it is still possible to experience I/O contention.

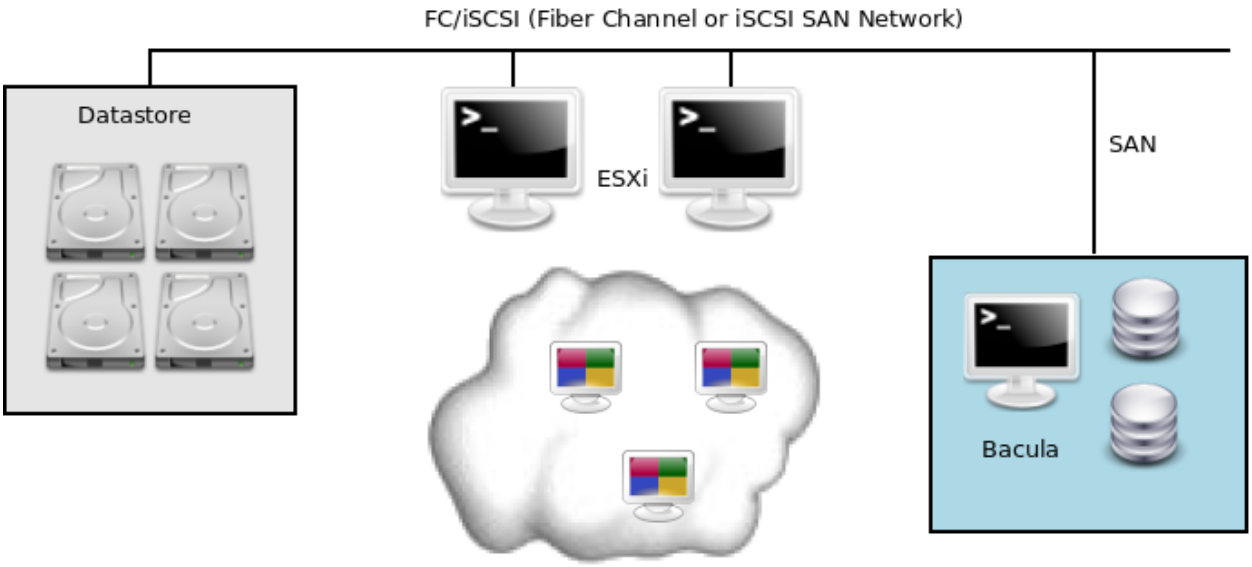

Fig. 2: Backup through SAN Network

When using block differential techniques such as those used by the vSphere Plugin, you need to **ensure** that all Incremental backups are available for restore. If one of your Incremental Jobs is missing at the restore time, Bacula will not be able to create a consistent image. Using the Differential level reduces the number of Jobs that are required for restore, and thus reduces the risk that something might be lost. To avoid losing important Incremental Jobs, you must ensure that your Volume retention periods are long enough to recover all of your data.

See also:

Go back to:

- *Installing Bacula Client on Each Guest*

Go to:

- *Strategies Comparison*

Go back to the *main vSphere Plugin Backup Strategies page*.

Go back to the *main vSphere Plugin page*.

3.2 Installing Bacula Client on Each Guest

This strategy doesn't use the Bacula Enterprise vSphere Plugin, but instead installs a Bacula Enterprise File Daemon (FD) on every virtual machine as if they were physical servers. To optimize the I/O usage on your VMware ESX/ESXi server, you will use **Schedule**, **Priority** and **Maximum Concurrent Jobs** to spread your backup jobs over your backup window. Since all servers use the same set of disks, running all your backup jobs at the same time could create a bottleneck on the disk/network subsystem.

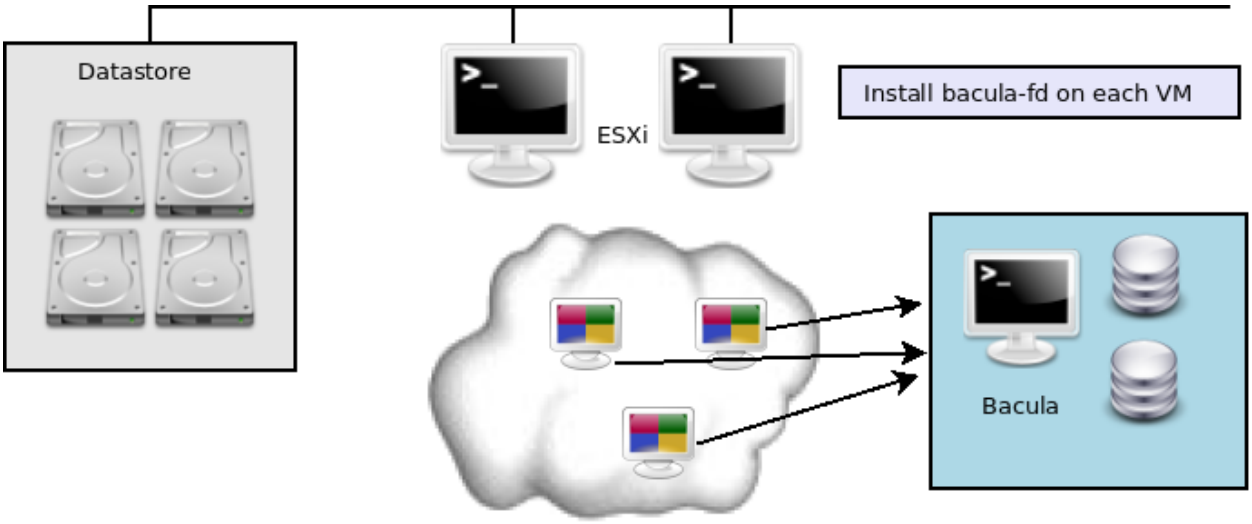

Fig. 3: Installing bacula-fd on each guest

Installing a Bacula Enterprise File Daemon on each virtual machine allows you to manage your virtual servers like physical servers, and to use all of Bacula Enterprise's features such as:

- quick restores of individual files
- checksum of individual files for virus and spyware detection
- Verify Job
- file/directory exclusion (such as swap or temporary files)
- file level compression
- and others.

See also:

Go to:

- *Image Backup with the vSphere Plugin*
- *Strategies Comparison*

Go back to the *main vSphere Plugin Backup Strategies page*.

Go back to the *main vSphere Plugin page*.

3.3 Strategies Comparison

Table 1: Backup Strategies Comparison

Features	Inside Guest	vSphere VADP
Incremental backup	Yes	Yes
FileSet options	Yes	No
Block level backup	No	Yes
Accurate support	Yes	Yes
Speed	Slow	Fast
I/O Load	High	Low
LAN free backup	No	Yes

See also:

Go back to:

- *Installing Bacula Client on Each Guest*
- *Image Backup with the vSphere Plugin*

Go back to the *main vSphere Plugin Backup Strategies page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Scope*
- *Features*

Go to:

- *Installation*
- *Configuration*
- *Operations*
- *Limitations*
- *Troubleshooting*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

4 Installation

This article describes how to install Bacula Enterprise vSphere Plugin.

4.1 Prerequisites

Note: The `Plugin Directory` is set by default, however, it is recommended to double-check if the `Plugin Directory` directive of the `File Daemon` resource in */opt/bacula/etc/bacula-fd.conf* points to the directory where the `vsphere-fd.so` plugin is installed. The standard Bacula plugin directory is: `/opt/bacula/plugins`.

Your File Daemon must have access to the vCenter/vSphere management network, port TCP/443 for API calls, and port TCP/902 for NBD data transfer or direct SAN access to vSphere datastores for SAN transport mode.

vSphere Installation with BIM

In order to install the vSphere Plugin with BIM, install the File Daemon with BIM and choose to install the vSphere Plugin during the FD installation.

Click here for more details on the plugin installation process with BIM.

See also:

See an alternative way of installing the vSphere Plugin - *vSphere Installation with Package Manager*.

Go back to the *main Installation page*.

Go back to the *main vSphere Plugin page*.

Installation with Package Manager

Due to VMware vSphere library dependencies, vSphere Plugin packages are available for a limited number of platforms (check the list *here*). Make a request from your Customer Portal to access the vSphere Plugin.

Installation can be performed using apt-get or yum/dnf with the package names below if you have added the vSphere Plugin to your Bacula repositories. Otherwise, download the packages and install them manually as below:

```
rpm -ivh bacula-enterprise-vsphere-vixdisk*.rpm
```

```
rpm -ivh bacula-enterprise-vsphere-16.*.rpm
```

These packages will ensure that your **Bacula Enterprise** version is compatible with the vSphere plugin and will install the `vsphere-ctl`, `vddk` and `vsphere-fd` programs.

```
# cd /opt/bacula
/opt/bacula # ls plugins/vsphere-fd.so bin/vsphere-ctl* bin/vddk
-rwxr-xr-x 1 root root  551890 10 nov.  15:13 plugins/vsphere-fd.so
-rwxr-xr-x 1 root root 3551890 10 nov.  15:13 bin/vsphere-ctl.jar
-rwxr-xr-x 1 root root    4096 10 nov.  15:13 bin/vsphere-ctl
-rwxr-xr-x 1 root root 3551890 10 nov.  15:13 bin/vddk
```

See also:

See an alternative way of installing the vSphere Plugin - *vSphere Installation with BIM*.

Go back to the *main vSphere Plugin Installation page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Scope*
- *Features*
- *Backup Strategies*

Go to:

- *Configuration*
- *Operations*
- *Limitations*
- *Troubleshooting*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

5 Configuration

The following chapter presents the information on how to configure the vSphere Plugin and how to configure a vSphere backup job.

5.1 vSphere Plugin Configuration

The following chapter presents the information on how to configure the vSphere Plugin.

Prepare Your ESXI or vCenter Environment

The following article gathers the necessary information on how to prepare your ESXI or vCenter environment.

Enabling Block Level Incremental Backup

Note: Changed Block Tracking (CBT) is not supported when the virtual hardware version is 6 or earlier, or when the virtual disk is attached to a shared virtual SCSI bus.

For CBT to identify altered disk sectors since the last change ID, the following items are required:

- The host must be ESX/ESXi 4.0 or later.
- The virtual machine owning the disks to be tracked must be hardware version 7 or later.
- I/O operations must go through the ESX/ESXi storage stack. So NFS is supported, as is RDM in virtual compatibility mode, but not RDM in physical compatibility mode. VMFS is supported, whether backed by SAN, iSCSI, or local disk.
- CBT must be enabled for the virtual machine (see below).
- Virtual machine storage must not be (persistent or non-persistent) independent disk, meaning unaffected by snapshots.

For CBT to identify disk sectors using Full backup, the following items are required:

- The virtual disk must be located on a VMFS volume, backed by SAN, iSCSI, or local disk.

- The virtual machine must have zero (0) snapshots when CBT is enabled, for a clean start.

When using "Thick Provisioned Eager Zeroed" disks, the VMware CBT will report all blocks as "used" during the Full backup.

For virtual machines that do not support CBT, the vSphere Plugin will always perform a Full backup of the virtual disks.

To check if a virtual disk has CBT enabled, open the vSphere Client, select a **powered-off** virtual machine **without snapshots**.

- Right-click the virtual machine and click Edit Settings.
- Click the Options tab.
- Click General under the Advanced section and then click Configuration Parameters. The Configuration Parameters dialog opens.
- Click Add Row.
- Add the `ctkEnabled` parameter and then set its value to `true`.
- Click Add Row, add `scsi0:0.ctkEnabled`,and set its value to `true`.

Note: `scsi0:0` in `scsi0:0.ctkEnabled` indicates the SCSI device assigned to the hard disk that is added to the virtual machine. Every hard disk added to the virtual machine is given a SCSI device that appears similar to `scsi0:0`, `scsi0:1`, or `scsi1:1`.

Note: VMWare FAQ articles may help:

- https://kb.vmware.com/s/article/1020128
- http://kb.vmware.com/kb/1033816

During the first Full backup, the vSphere Plugin will try to enable CBT automatically. To disable this feature, use the `keep_cbt` option in the Plugin command string.

```
Plugin = "vsphere: keep_cbt host=guest1"
```

No snapshots must exist on a virtual machine at the time of enabling CBT on it.

If you notice large Full backup jobs despite small real disk usage, contact your Bacula Systems support team to assess the situation and possibly shrink the backup size by resetting CBT.

Detecting when CBT Is Not Available

When the CBT option is not available for a disk, the `vsphere-ctl*log` file may contain the following notice:

```
com.cybozu.vmbkp.soap.SnapshotManager, getChangedBlocksOfDisk]
      com.vmware.vim25.FileFault
```

When the vSphere Plugin receives this notice, it will automatically do a Full backup of the disk image.

See also:

Go to:

- *Enabling SAN Access*

- *Check User Permissions*
- <QueryInformationAboutvSphereEnvironment>

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Enabling SAN Access

To use the SAN transport method, your backup server where the vSphere Plugin is installed should have access to all LUNs which are exported to your ESXi server(s). The `multipathd` tools will avoid problems with multiple device paths to SAN devices.

Once your SAN LUNs are visible to your backup server as `/dev/sda`, `/dev/sdb`, … The vSphere Plugin will open each LUN to get the UUID and compare it with what the ESXi server is providing. For example, when using iSCSI, the `lsscsi` command will display them as:

```
% lsscsi
[5:0:0:0]    disk    IET      VIRTUAL-DISK     0     /dev/sdb
```

You can verify that SAN transport is used during backup by using the `debug` option in the vSphere Plugin command line. If SAN transport mode is used, the `vddk` trace file will have an entry showing "Opened san://…" like in the following example:

```
% grep san: /opt/bacula/working/vsphere/<moref>/<seq>/0.log
DISKLIB-LIB   : Opened "san://3-snapshot-18[datastore2] test/test_2.vmdk...
```

When the SAN mode is not available, the vSphere Plugin will automatically switch to the `NBD` transport mode.

See also:

Go back to:

- *Enabling Block Level Incremental Backup*

Go to:

- *Check User Permissions*
- <QueryInformationAboutvSphereEnvironment>

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Check User Permissions

Attention: Starting with version 12.8.0 of Bacula Enterprise, it is possible to use bconsole to query vSphere and see if a user has all required permissions to perform backups and restores.

Example


```
[root@localhost bin]# ./bconsole
Connecting to Director localhost:9101
1000 OK: 10002 localhost-dir Version: 12.6.1 (05 March 2021)
Enter a period to cancel a command.
*.query client=localhost-fd plugin="vsphere: server=vcenter_192_168_0_8"␣
↪parameter=permissions
missing=VirtualMachine.Provisioning.DiskRandomAccess
missing=VirtualMachine.Provisioning.DiskRandomRead
missing=VirtualMachine.Provisioning.FileRandomAccess
missing=VirtualMachine.Provisioning.GetVmFiles
missing=VirtualMachine.State.CreateSnapshot
missing=VirtualMachine.State.RemoveSnapshot
missing=VirtualMachine.State.RenameSnapshot
missing=VirtualMachine.State.RevertToSnapshot
missing=VirtualMachine.Interact.PowerOff
missing=VirtualMachine.Interact.PowerOn
...
```

If the list is not empty, the listed permissions must be configured properly.

Note: In addition to the above `bconsole` command, you can also use the `vsphere-ctl` command to check the permissions of the current user on the vCenter system and diagnose issues if any:

```
/opt/bacula/bin/vsphere-ctl query list_missing_permissions
```

The following privileges can be allocated to a role and assigned to a Bacula user to perform vStorage backups and restores. These are the minimum required permissions that have been found to be sufficient in the tests performed by Bacula Systems for a basic vSphere environment.

This list may change in the future. The permissions are best propagated downwards from the root of the vSphere level. Additional privileges might be required if advanced features are in use.

Set the following permissions in your vSphere/vCenter environment:

Privilege Level	Permissions
Datastore	• Allocate space • Browse Datastores • Configure Datastores • Low level file operations • Remove File • Update virtual machine Files
Distributed Virtual Switch	• Host operation
Folder	• Create Folder
Global	• Cancel Task • Disable Methods • Enable Methods • Licenses • Log Event • Manage Custom Attributes • Set Custom Attributes • Settings
Host: Configuration	• Advanced Settings • Storage Partition Configuration
Host: Local Operations	• Create Virtual Machine • Delete Virtual Machine • Reconfigure virtual machine
Network	• Assign Network
Resource	• Assign Vapp to resource pool • Assign Virtual Machine to resource pool • Query Vmotion
Tasks	• Create task • Update task
vApp	• Add virtual machine • Assign virtual machine • Create • Export • Import • vApp application configuration • vApp instance configuration • vApp resource configuration • View OVF Environment
Virtual Machine: Configuration	• Add Existing Disk • Add New Disk • Add or Remove Device • Advanced • Change CPU Count • Change Resource

See also:

Go back to:

- *Enabling Block Level Incremental Backup*
- *Enabling SAN Access*

Go to:

- <QueryInformationAboutvSphereEnvironment>

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main Configuration page*.

Go back to the *main vSphere Plugin page*.

Query Information about vSphere Environment

Attention: New in version 16.0.12

Using a similar mechanism to what was described in the previous section about user permisisons, it is possible to query VMware to get different kinds of information. Below some examples:

```
// List networks
.query client=my-fd plugin="vsphere:" parameter=network

// List resource pools
.query client=my-fd plugin="vsphere:" parameter=pool

// List datastore
.query client=my-fd plugin="vsphere:" parameter=datastore

// List current configuration
.query client=my-fd plugin="vsphere:" parameter=config_list

// Check some configuration section
.query client=my-fd plugin="vsphere: sectionname=vsphere" parameter=config_check
```

See also:

Go back to:

- *Enabling Block Level Incremental Backup*
- *Enabling SAN Access*
- *Check User Permissions*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

See also:

Go to:

- *Connect to ESXi or vCenter Server*

- *Test vSphere Configuration*

Go back to the *vSphere Plugin Configuration page*.

Go back to the *main Configuration page*.

Go back to the *main vSphere Plugin page*.

Connect to ESXi or vCenter Server

Automated Configuration

The article describes the recommended automated configuration using `vsphere-ctl config` command. If you have access to BWeb and would like to use it to configure the connection to VMware infrastructure and your backup jobs, skip ahead to *BWeb VMware Center Module*.

> **Attention: New in version 16.0.8**
>
> The automated configuration is available with version 16.0.12.

Set up the vSphere configuration in the console:

```
[root@localhost bin]# ./vsphere-ctl config create
Enter ESXi/vCenter url: 192.168.0.15
Enter user: administrator@vsphere.local
Enter password:
Connecting to "https://192.168.0.15/sdk"...
OK: successful connection
OK: user has all necessary permissions to perform backups and restores
Select the ESXi host that contains the VMs you wish to backup:
    1) 192.168.0.8
    2) 192.168.0.26
Select host: 1
Computing thumbprint of host "192.168.0.8"
OK: thumbprint for "192.168.0.8" is␣
↪04:24:24:13:3C:AD:63:84:A1:9F:E5:14:82:7D:5C:31:25:A8:FA:89
OK: added entry [vcenter_192_168_0_8] to ../etc/vsphere_global.conf
```

It's also possible to list configurations that are already present:

```
[root@localhost bin]# ./vsphere-ctl config list
[esxi_192_168_0_26]
     root_directory = /opt/bacula/working/esxi_192_168_0_26
     server = 192.168.0.26
     thumbprint = A0:6D:75:53:9E:30:85:BB:99:63:6E:33:C4:B8:64:E9:06:AD:BF:CF
     url = https://192.168.0.26/sdk
     username = root
[esxi_192_168_0_8]
     root_directory = /opt/bacula/working/esxi_192_168_0_8
     server = 192.168.0.8
     thumbprint = 04:24:24:13:3C:AD:63:84:A1:9F:E5:14:82:7D:5C:31:25:A8:FA:89
     url = https://192.168.0.8/sdk
     username = root
```

(continues on next page)

(continued from previous page)

```
[vcenter_192_168_0_26]
     root_directory = /opt/bacula/working/vcenter_192_168_0_26
     server = 192.168.0.26
     thumbprint = A0:6D:75:53:9E:30:85:BB:99:63:6E:33:C4:B8:64:E9:06:AD:BF:CF
     url = https://192.168.0.15/sdk
     username = administrator@vsphere.local
[vcenter_192_168_0_8]
     root_directory = /opt/bacula/working/vcenter_192_168_0_8
     server = 192.168.0.8
     thumbprint = 04:24:24:13:3C:AD:63:84:A1:9F:E5:14:82:7D:5C:31:25:A8:FA:89
     url = https://192.168.0.15/sdk
     username = administrator@vsphere.local
```

Existing configuration entries can be validated with:

```
[root@localhost bin]# ./vsphere-ctl config check vcenter_192_168_0_26
Connecting to "https://192.168.0.15/sdk"...
OK: successful connection
Computing thumbprint of host "192.168.0.26"
OK: local thumbprint matches server thumbprint
Checking user privileges...
OK: user has all necessary privileges
```

Finally, you can delete configuration entries with:

```
[root@localhost bin]# ./vsphere-ctl config delete --entry vcenter_192_168_0_8
OK: deleted entry "vcenter_192_168_0_8".
```

See also:

Go to:

- *Manual Configuration*
- *Using Multiple vSphere Servers*
- *Obscure vSphere Password*
- *vsphere_global.conf Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Manual Configuration

If you have access to BWeb and would like to use it to configure the connection to VMware infrastructure and to configure your backup jobs, skip ahead to *BWeb VMware Center Module*.

Attention: New in version 12.8.0

Starting with version 12.8.0 it is possible to use the config command described in *Automated Configuration* to interactively create a configuration in the command-line. Otherwise continue from here to configure the plugin directly with your text editor.

The vSphere network access to your ESXi or vCenter server is configured in `/opt/bacula/etc/vsphere_global.conf `.

```
% cat /opt/bacula/etc/vsphere_global.conf
[vsphere]
        username = root
        password = vspherepassword
        server = 192.168.0.1
        url = https://192.168.0.1/sdk
        thumbprint = 34:F5:0F:10:82:59:EF:2D:DB:96:CC:5B:C4:66:33:83:DC:91:AF:09
        root_directory = /opt/bacula/working/vsphere
```

Click *here* (recommended to open in a new tab) to see all available directives for the `vsphere_global.conf` file.

To get the `thumbprint` value, you can copy it from the ESXi console screen. Hit F2 then log in. The thumbprint is displayed in "View Support Information" under the "SSL Thumbprint (SHA1)".

Starting with Bacula Enterprise 8.6, it is also possible to use the `vsphere-ctl` 'thumbprint' command to display the `thumbprint` from the Bacula client.

```
# /opt/bacula/bin/vsphere-ctl thumbprint 192.168.0.1
```

The thumbprint may also be obtained via an ssh session on the ESXi host:

```
# openssl x509 -sha1 -in   \
    /etc/vmware/ssl/rui.crt -noout -fingerprint | cut -d '=' -f 2 "
```

Or on a vCenter server:

```
# openssl x509 -sha1 -in \
    /etc/vmware-vpx/ssl/rui.crt -noout -fingerprint | cut -d '=' -f 2 "
```

See also:

Go back to:

- *Automated Configuration*

Go to:

- *Using Multiple vSphere Servers*
- *Obscure vSphere Password*
- *vsphere_global.conf Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Using Multiple vSphere Servers

You may specify multiple vSphere servers in the `vsphere_global.conf` file. When using this feature, you will need to specify the `server=xxx` option in the Plugin Command line. If you have access to BWeb and would like to use it to configure the connection to VMware infrastructure and to configure your backup jobs, skip ahead to *BWeb VMware Center Module*. You may also use the CLI automated configuration tool starting with version 12.8.0 in *Automated Configuration*.

Click *here* (recommended to open in a new tab) to see all available directives for the `vsphere_global.conf` file.

It is also mandatory to specify a unique root_directory for each section so that information about VMs from one vCenter or ESXi server is not overwritten with information from a different one.

```
% cat /opt/bacula/etc/vsphere_global.conf
[vsphere]
    username = root
    password = vspherepassword
    server = 192.168.0.1
    url = https://192.168.0.1/sdk
    thumbprint = 01:F5:0F:10:82:59:EF:2D:DB:96:CC:5B:C4:66:33:83:DC:91:AF:01
    default_datastore = datastore1
    default_restore_host = esx1
    root_directory = /opt/bacula/working/vsphere

[other]
    username = root
    password = vspherepassword
    server = 192.168.0.2
    url = https://192.168.0.2/sdk
    thumbprint = 02:F5:0F:10:82:59:EF:2D:DB:96:CC:5B:C4:66:33:83:DC:91:AF:01
    default_datastore = abigdatastore
    root_directory = /opt/bacula/working/other
    vddk_backup_transport_mode = san:nbdssl
    vddk_restore_transport_mode = san:nbd

[secure]
    username = root
    hpassword = MTEyOjEyNzoGAwAYFQIVABEDAwcfAhQA
    server = 192.168.0.3
    url = https://192.168.0.3/sdk
    thumbprint = 03:F5:0F:10:82:59:EF:2D:DB:96:CC:5B:C4:66:33:83:DC:91:AF:01
    default_datastore = abigdatastore
    root_directory = /opt/bacula/working/secure
```

The `[vsphere]` section is optional in the `vsphere_global.conf` file. If not present, make sure the *–server* options is always used for backup (or vsphere-ctl operations).

```
% cat /opt/bacula/etc/conf.d/FileSet-other.conf
FileSet {
  Name = FileSet-other
  Include {
    Plugin = "vsphere: server=other"
  }
}

% /opt/bacula/bin/vsphere-ctl --server other update
1: 1 vm1
2: 2 vm2
```

Click *here* to see all the vSphere Fileset plugin command options.

See also:

Go back to:

- *Automated Configuration*

- *Manual Configuration*

Go to:

- *Obscure vSphere Password*
- *vsphere_global.conf Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Obscure vSphere Password

Starting with the 8.0.3 version of the vSphere Plugin, it is now possible to obscure the vSphere password in the `vsphere_global.conf` file. The obscured password field is called `hpassword`.

Click *here* (recommended to open in a new tab) to see all available directives for the `vsphere_global.conf` file.

The bconsole `@encode` command can be used to generate the obscured password. Note that if the string you want to encode contains "=", you must use the `string=` keyword in the command.

```
# /opt/bacula/bin/bconsole
* @encode vspherepassword
MTEyOjEyNzoGAwAYFQIVABEDAwcfAhQA

* @encode string="passwordwith="
NTMwOjU0Mzpic2FhZX1gdmV7ZnovAA
```

```
# cat /opt/bacula/etc/vsphere_global.conf
[vsphere]
    username = root
    hpassword = MTEyOjEyNzoGAwAYFQIVABEDAwcfAhQA
    server = 192.168.0.1
    url = https://192.168.0.1/sdk
    thumbprint = 01:F5:0F:10:82:59:EF:2D:DB:96:CC:5B:C4:66:33:83:DC:91:AF:01
```

See also:

Go back to:

- *Automated Configuration*
- *Manual Configuration*
- *Using Multiple vSphere Servers*

Go to:

- *vsphere_global.conf Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

vsphere_global.conf Options

Table 2: Sphere Connection Configuration vsphere_global.conf

Option	Required	Default	Info
keep_generation	No	100	Maximum number of Backup between two Full backups
profile_all_vm	No	vsphere_all_vm.prof	Internal filename used to store virtual machine profile information
root_directory	No	/opt/bacula/working/vsphere	Working directory of the vSphere plugin
vddk_path	No	/opt/bacula/bin/vddk	
vddk_backup_transport_mode	No	tadd:nbdssl:nbd	Specify the different transport method to try with the VDDK service during backup. Available with 14.0.1. See the VMWare Disk Transport Library documentation for the list of values.
vddk_restore_transport_mode	No	tadd:nbdssl:nbd	Specify the different transport method to try with the VDDK service during restore Available with 14.0.1. See the VMWare Disk Transport Library documentation for the list of values.
username	Yes		Username allowed to connect to vSphere
password	Yes		Username password allowed to connect to vSphere
hpassword	No		Username obscured password allowed to connect to vSphere. Read more *here*.
timeout	No	60 seconds	Connection timeout (available since version 8.2.9) used to contact the vSphere server in seconds. The timeout with internal file locking is 10x the value (available since version 8.4.15).
thumbprint	Yes		SSL Thumbprint of the vSphere server (required for vSphere 6.0 and above)
server	Yes		vSphere ESXi server used for Backup
url	Yes		vSphere ESXi or vCenter server URL used for SOAP call
default_datastore	No	datastore1	Default datastore for restore
default_restore_host	No		Default ESX server used for restore if multiple ones are available in the vCenter
default_ovf	No		Default OVF description used when current OVF fails to be loaded in VMware (available in version 6.2.3-2 and later).
root_directory	No	/opt/bacula/working/vsphere	Directory used to store internal plugin data
datastore_minimum_space	No		Minimum space to keep on a Datastore. ex: 5GB
datastore_allow_overprovisioning	No	Yes	Allow to restore a VM using Over Provisioning. When set to `no`, the restore process will ensure that all `full` disks can fit on the Datastore.
datastore_refresh_interval	No	600 seconds	Specify the interval used to refresh storage statistics of the Datastore.
nfc_host_port	No	902	Specify the NFC TCP port to contact the ESX server. Available with 8.4.12.
server_port	No	443	Specify the HTTPS TCP port to contact the ESX server. Available with 8.4.12.
checkssl	No	No	Check SSL certificate with vSphere server. Available with 16.0.3.

The `index` feature will generate records in the Catalog to quickly seek to a given block in the backup stream. The granularity of the index can be controlled with the Storage Daemon device directive `MaximumFileIndex`. The default value is 100MB.

Going back to the Configuration chapter

To go back to the main Configuration chapter, click *here*.

See also:

Go back to:

- *Automated Configuration*
- *Manual Configuration*
- *Using Multiple vSphere Servers*
- *Obscure vSphere Password*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Prepare Your ESXI or vCenter Environment*

Go to:

- *Testing vSphere Configuration*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Test vSphere Configuration

To test the vSphere Plugin, you can use the following command as the user that the File Daemon runs as - (usually `root`):

```
# /opt/bacula/bin/vsphere-ctl update
1: 3 squeeze2
2: 4 squeeze.esx
```

The `vsphere-ctl update` command should print a list of all virtual machines that are defined in your ESXi server. If not, check if your credentials in `vsphere_global.conf` are properly set.

Note: When a VM is removed from inventory and then re-imported - even if it is re-imported to the same ESXi host or same vCenter Server - its MoRef will be changed from its previous value. This will require that the `vsphere-ctl update` command is run prior to attempting backup of this VM again. Additionally, a Full backup must be performed when a VM is removed and re-imported to inventory.

If you foresee a use case where VMs are often removed from inventory only to be re-imported at some point, we recommend that the `vsphere-ctl update` command be triggered in an Admin Job's RunScript which is set to run before the normal nightly backups of your vSphere infrastructure.

The `list` command displays information that is detected on the ESXi hosts and datastores.

```
# /opt/bacula/bin/vsphere-ctl list
Display host list available and their datastores:
 esxi.lan
     datastore1
     datastore2
```

(continues on next page)

(continued from previous page)

```
Will now display configured settings for restore:

No default_restore_host defined in vsphere_global.conf file, trying to
get it from vSphere. Will use restore host esxi.lan

No default_datastore defined in vsphere_global.conf file, trying to
get it from vSphere. Will use datastore datastore1
```

See also:

Go back to:

- *Prepare Your ESXI or vCenter Environment*
- *Connect to ESXi or vCenter Server*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

See also:

Go to:

- *BWeb VMware Center Module*
- *vSphere Plugin Backup Job Configuration*

Go back to the *main Configuration page*.

Go back to the *main vSphere Plugin page*.

5.2 BWeb VMware Center Module

BWeb Management Suite's configuration mode includes a "VMware Center" module which is devoted to integrating a Bacula environment with a vSphere environment.

In VMware Center, all vCenter and/or ESXi hosts may be defined and configured. Once configured, these vSphere hosts may be assigned to individual or multiple Bacula Clients having the vSphere Plugin installed.

Once vSphere hosts are assigned to Bacula Clients, vSphere Plugin Jobs and associated FileSets may be automatically configured so that as VMware VMs appear in a VMware environment, new Jobs/FileSets will be automatically added. Additionally, as VMs are removed from a VMware environment, their Jobs/FileSets may be automatically disabled, or completely removed depending on specific requirements. If preferred, Job and FileSet creation and removal may be performed manually instead.

Configuring New vSphere Host

On the BWeb main menu (left panel), expand "Virtual Machines" and click "VMware Center". You will be taken to the page with the "vSphere Hosts" listing, which should currently be empty.

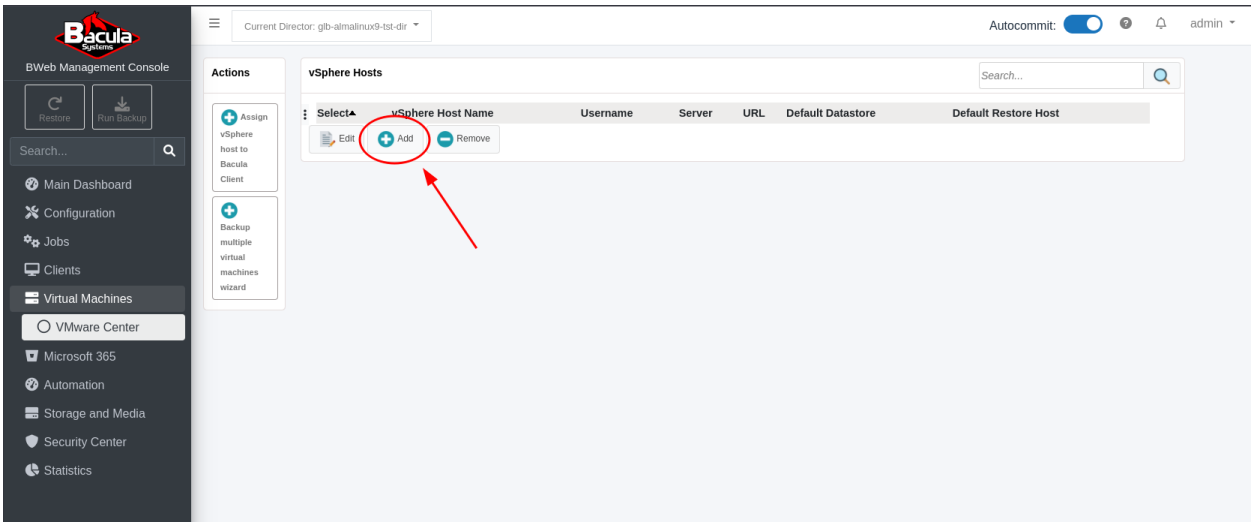

To add a new vCenter server or specific ESXi host to the list, click on the “Add” icon.

Fill in the fields in the pop-up form.

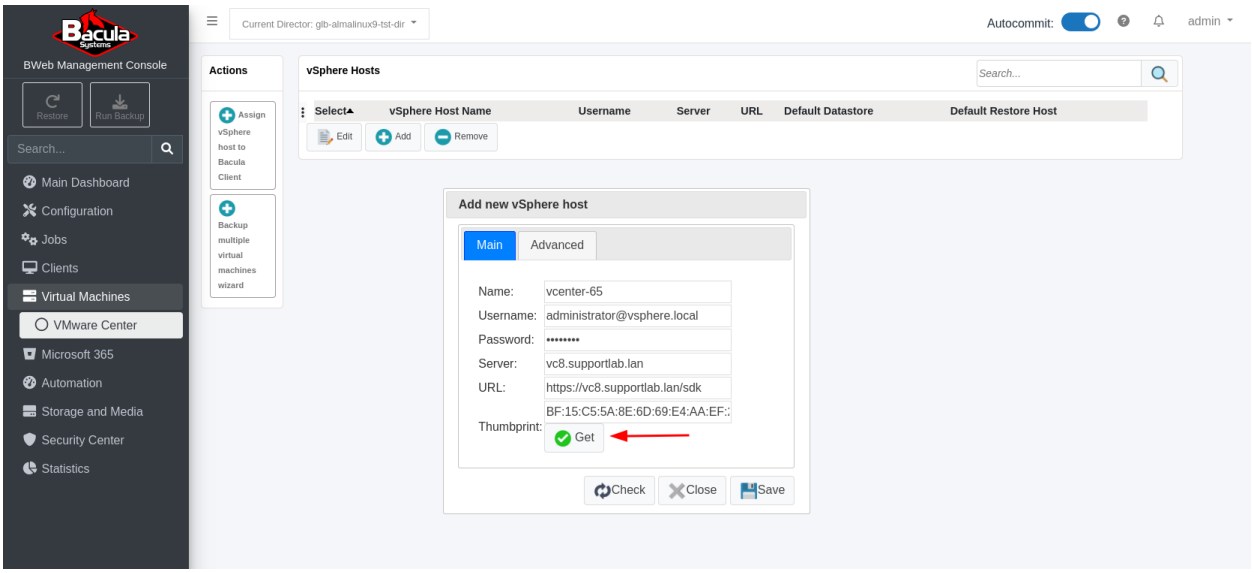

To get the “Thumbprint”, just click on the “Get” icon next to the Thumbprint field. Be sure that the “Server” IP address (or FQDN) is correct. It is important to verify this Thumbprint information now to prevent future issues or potential miscommunications with the wrong vSphere server.

To verify that the information in the form has been filled in correctly for this vSphere host, click on the circular arrows (check) icon, and after a few seconds, the icon should turn into a green check mark as in the image below.

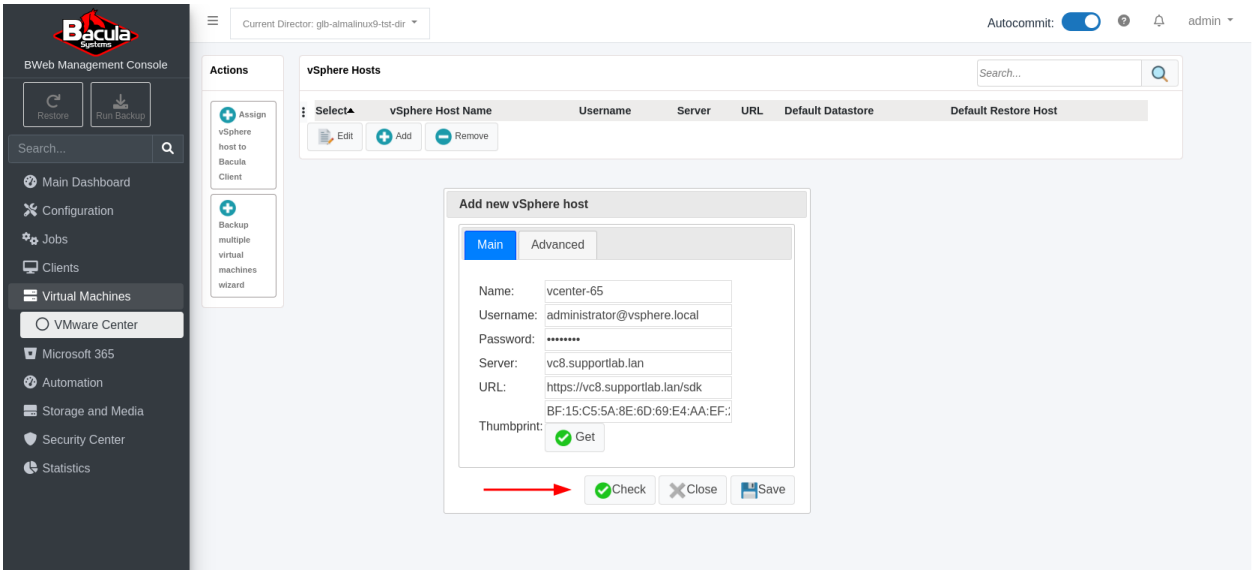

Click “Save” and you will be returned to the “vSphere Hosts” page, which should now be populated with this one (in this case called “vcenter-65”) vCenter host just created.

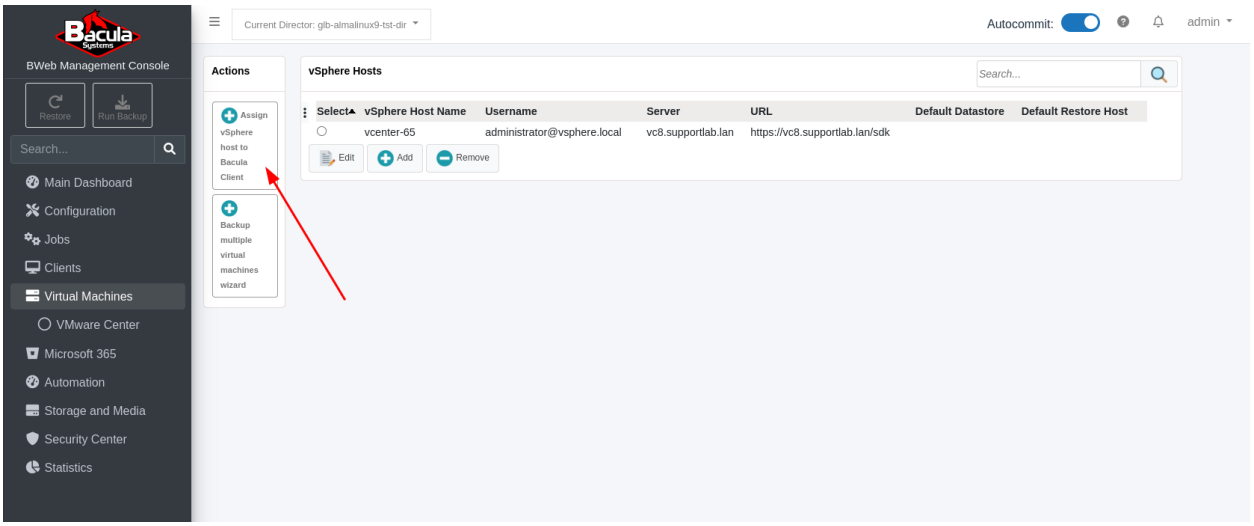

Next, we need to assign this vSphere (vCenter) host which we named “vcenter-65” to a Bacula FD with the vSphere plugin installed. To do this, click the “Assign vSphere host to Bacula Client” button in the upper left corner.

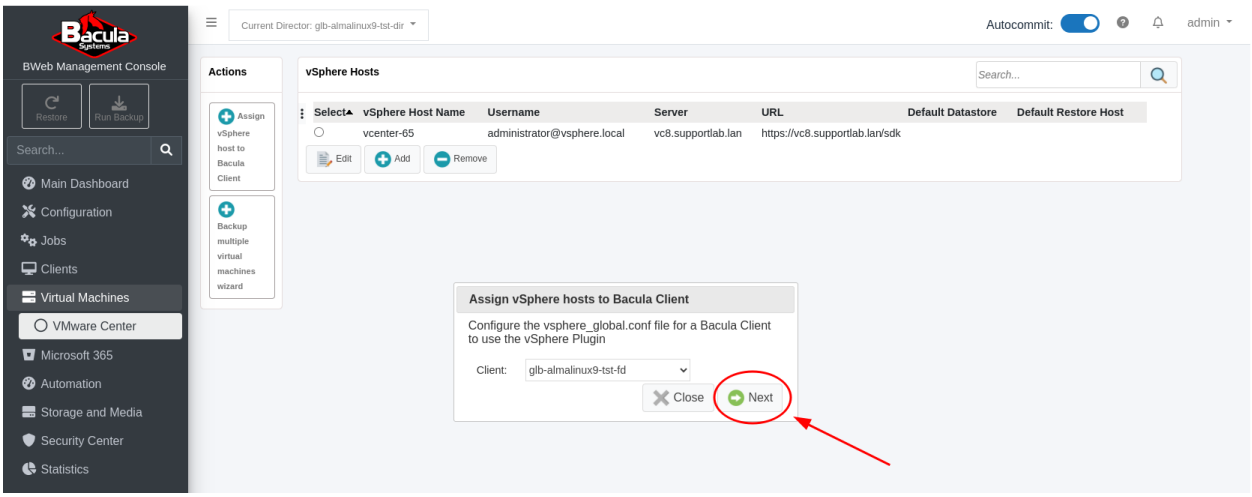

Select the correct Bacula Client from the “Client” drop-down list. In the screenshot, there is only one Bacula Client configured. A fully configured Bacula environment will have all of the Clients listed here to choose from.

Click “Next”, and you will be presented with a dialog box where the vSphere host(s) (shown here as vcenter-65) are shown in the “vSphere servers available” box. To assign this vSphere host to this FD, highlight it, then click the green arrow pointing to the “vSphere servers assigned” box.

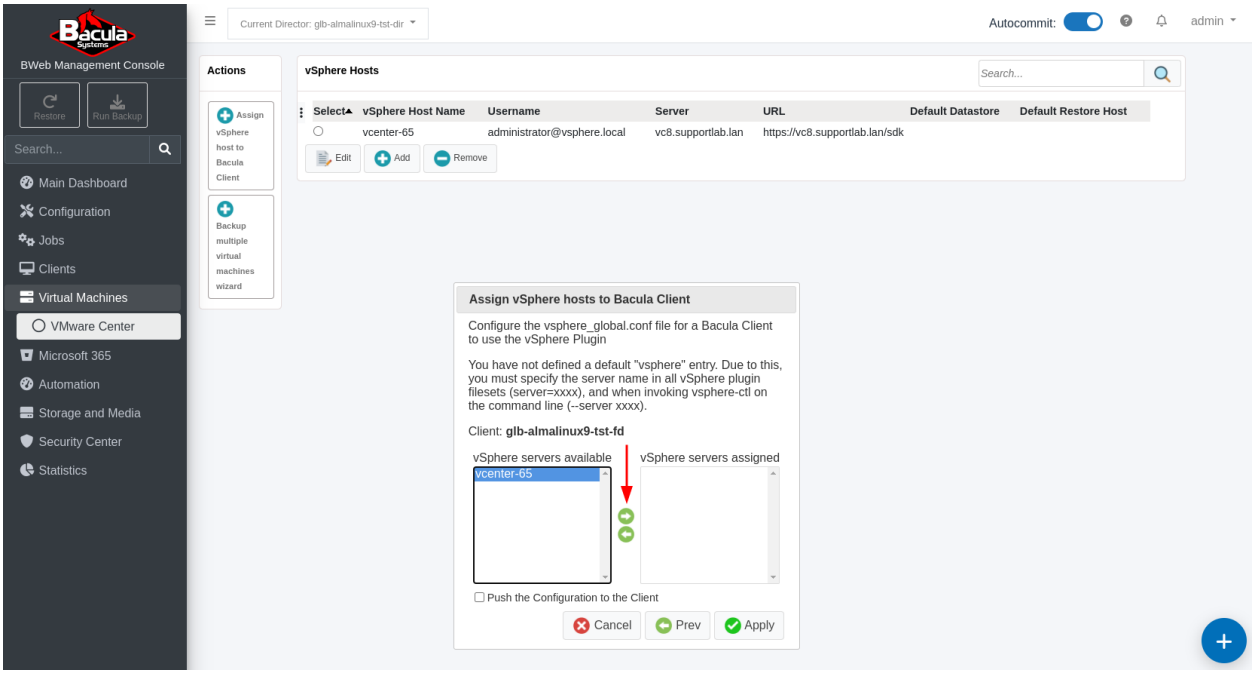

The “vcenter-65” vSphere host will be added to the “vSphere servers assigned” box on the right.

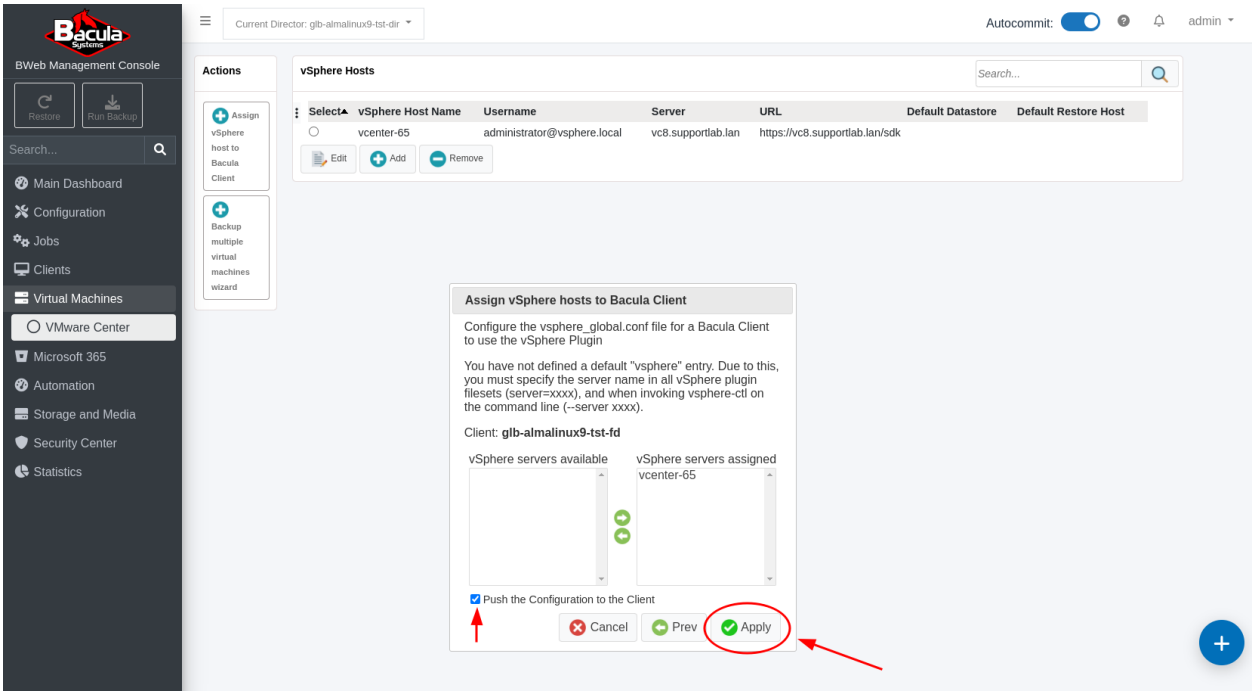

Be sure to check the “Push the Configuration to the Client” check box.

Click “Apply” and you will be taken to the “Push Configuration to the Client” page.

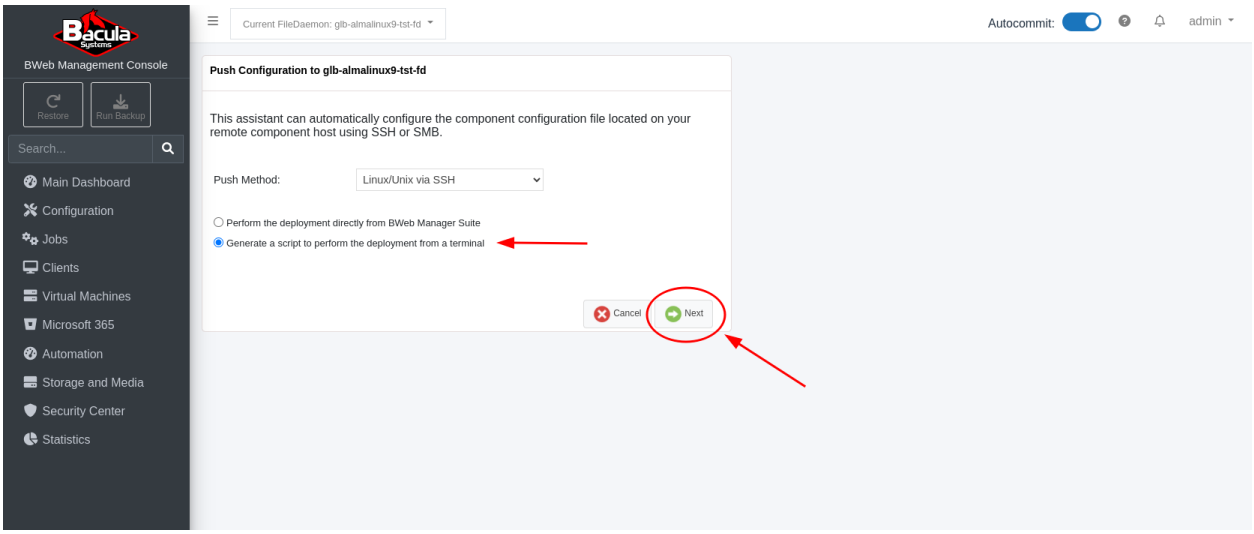

The settings on this page refer to the Bacula Client where the FD is running the vSphere plugin. It is the “Client” system we chose previously, and its name is displayed in the header of this dialog box.

Choose the Push Method, set it to Linux and check “Generate a script to perform the deployment from a terminal”. Click “Next”.

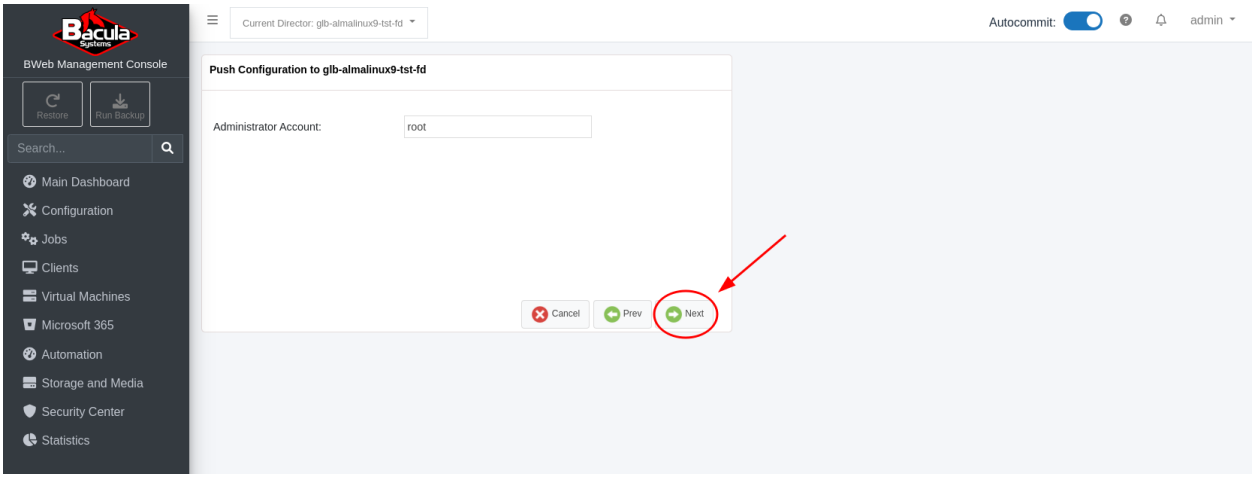

Set the “Administrator Account” to “root” for this Bacula Client FD.

Click “Next” and you will be presented with the “Push Configuration” dialog box where there will be a Linux command line script in the yellow box.

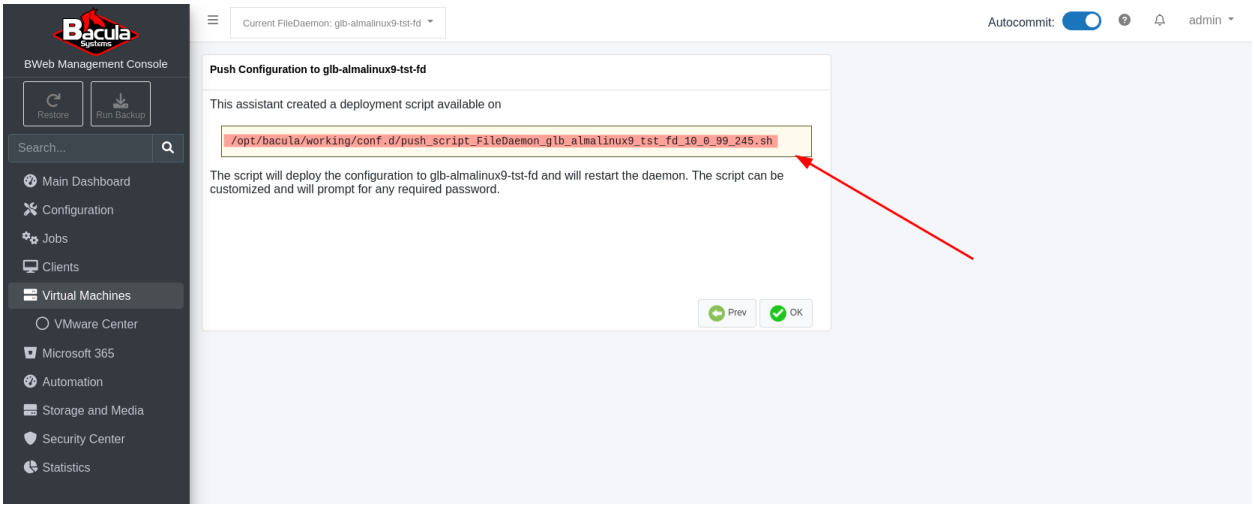

Copy the script path to the clipboard.

Next, ssh into your Bacula Director/BWeb server as root:

```
$ ssh root@10.0.99.245
root@10.0.99.245's password:
Activate the web console with: systemctl enable --now cockpit.socket

Last login: Wed Feb  7 10:27:54 2024 from 10.255.0.6
[root@glb-almalinux9-tst ~]#
```

Paste the command copied previously:

```
[root@glb-almalinux9-tst ~]# /opt/bacula/working/conf.d/push_script_FileDaemon_glb_
↪almalinux9_tst_fd_10_0_99_245.sh
```

Hit <enter> and you will see something similar to the following output. This output shows that the script run on the Bacula Director/BWeb server copies a second script to the Bacula Client and then runs that script on the Client via ssh.

This second script copies a correctly configured file into the "/opt/bacula/etc" directory on the Client and then restarts the Bacula FD:

```
INFO: Execute user script '/opt/bweb/bin/deploy_script_linux.sh'
INFO: Checking required files on 10.0.99.245
INFO: Copy configuration files
The authenticity of host '10.0.99.245 (10.0.99.245)' can't be established.
ED25519 key fingerprint is SHA256:B+W7gybwB/qYcXWWAMLwR1hElPS6gcupykIgMFtk3PU.
This key is not known by any other names
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added '10.0.99.245' (ED25519) to the list of known hosts.
root@10.0.99.245's password:
bacula-push-50944.tar                          100%   10KB   9.1MB/s   00:00
INFO: Extract configuration files
INFO: Backing up original configuration from 10.0.99.245
tar: Removing leading `/' from member names
/opt/bacula/working/vsphere_vcenter-65/
tmp.ou8uyzscYk                                 100%   36    70.9KB/s   00:00
tar: Removing leading `/' from member names
/opt/bacula/etc/vsphere_global.conf
tar: Removing leading `/' from hard link targets
Shared connection to 10.0.99.245 closed.
Shared connection to 10.0.99.245 closed.
tar: Removing leading `/' from member names
/opt/bacula/working/vsphere_vcenter-65/
/opt/bacula/etc/vsphere_global.conf
Shared connection to 10.0.99.245 closed.
INFO: Clean up temporary config files
Shared connection to 10.0.99.245 closed.
INFO: Restarting  service on 10.0.99.245
INFO: Clean up
```

The Bacula Client is now ready to run backup Jobs of VMware VMs managed by the vCenter server that we named "vcenter-65" in our "vSphere Hosts" listing.

At this point, you may manually create Jobs/FileSets to backup one or more VMs using the vSphere plugin on the Client. See the *Creating a vSphere Plugin Backup Job and FileSet in BWeb* section.

Alternatively, by clicking the "Backup multiple virtual machines wizard" button on the "VMware Center" page, you may create an automated configuration whereby Bacula will create a new Job and associated FileSet for each VM managed by this "vcenter-65" vCenter server. These Jobs/FileSets may be based on all VMs, some specific VM names, VM names using wildcards, VMs based on "VM Tags", all VMs on one or more Datastores, or one or more Resource Pools. See the *Creating An Automated vSphere Backup Environment Using BWeb* section for more information about configuring automatic job creation.

When configured using this automated method, Bacula will automatically add a new Job and associated FileSet for each new VM found, and will disable (or remove) the Job/FileSet for any VM that has been decommissioned and is no longer available on this vCenter server.

See also:

Go to:

- *Creating vSphere Plugin Backup Job and FileSet in BWeb*
- *Creating Automated vSphere Backup Environment Using BWeb*
- *Creating Automated vSphere Backup Environment Using Command Line Scripting*

Go back to the *main BWeb VMware Center Module page*.

Go back to the *main vSphere Plugin page*.

Creating vSphere Plugin Backup Job and FileSet in BWeb

Manually Creating vSphere Plugin Backups With BWeb

Before proceeding with this section, be sure to follow *Configuring New vSphere Host article* guide where you will be shown how to add vSphere hosts to the BWeb "VMware Center". The purpose of this section is to demonstrate how to manually create Bacula vSphere Plugin backup Jobs and their associated FileSets to backup your VMware VM(s) with the use of BWeb.

Creating New vSphere Plugin FileSet

On the BWeb main menu, expand "Configuration -> Director" and click "FileSets". You will be taken a page where all of the currently configured FileSets are listed.

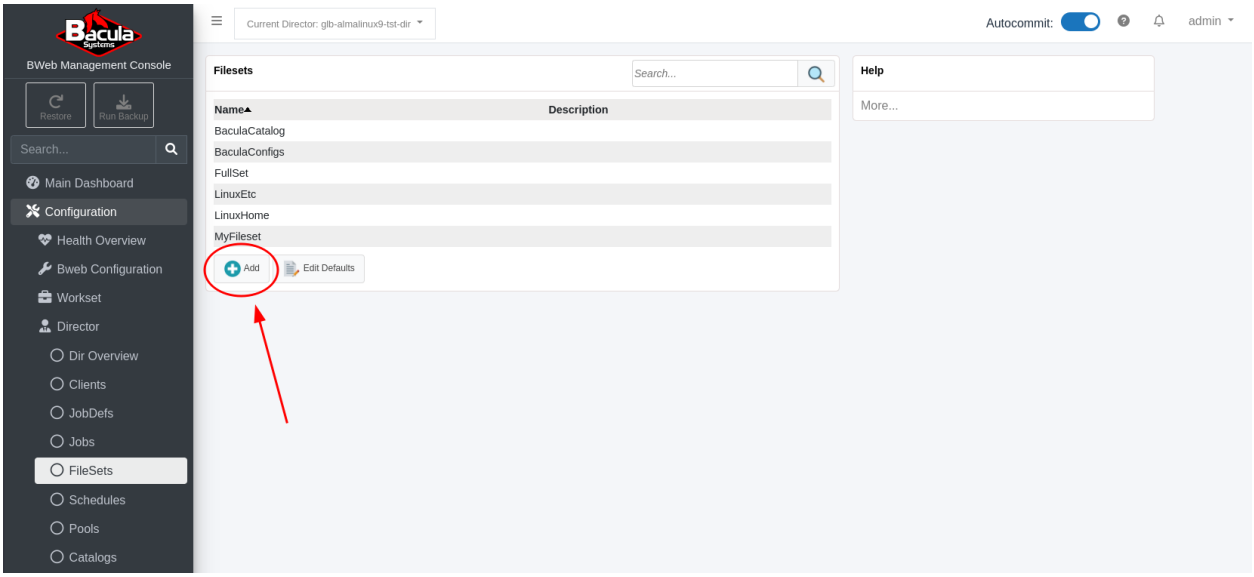

Here, we will start by creating a FileSet which will be used to backup one VMware VM managed by a vCenter server. We will use the 'vcenter-65' vCenter server which was configured in the "BWeb VMware Center" section of this document.

Click the "+" next to the word "FileSets" in the middle dialog box. You will be taken to a form where this new FileSet will be configured.

Fill in the "Fileset Name:" field with an appropriate name. In this example, the name "vcenter-65_vm-test" clearly indicates that the VM being backed up is called 'vm-test' and it is managed by a vCenter server called 'vcenter-65'. Optionally, you may also fill in the "Description:" field.

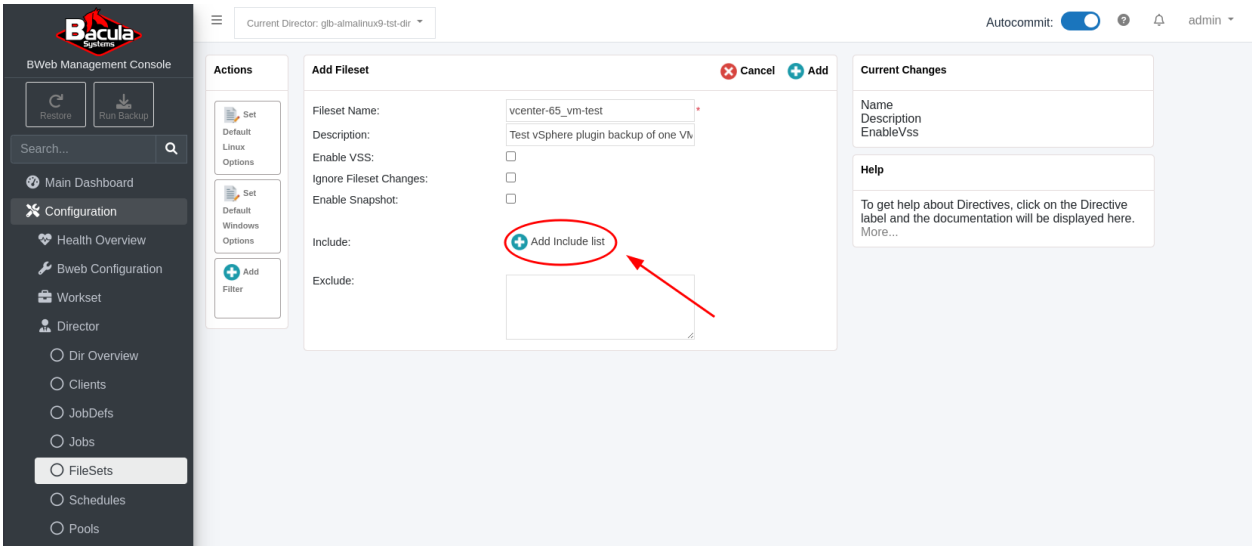

Click "+ Add include list" and you will be taken to the "Configure Fileset" dialog box. In this dialog box, select the client from the drop-down menu. Click the "vSphere" checkbox.

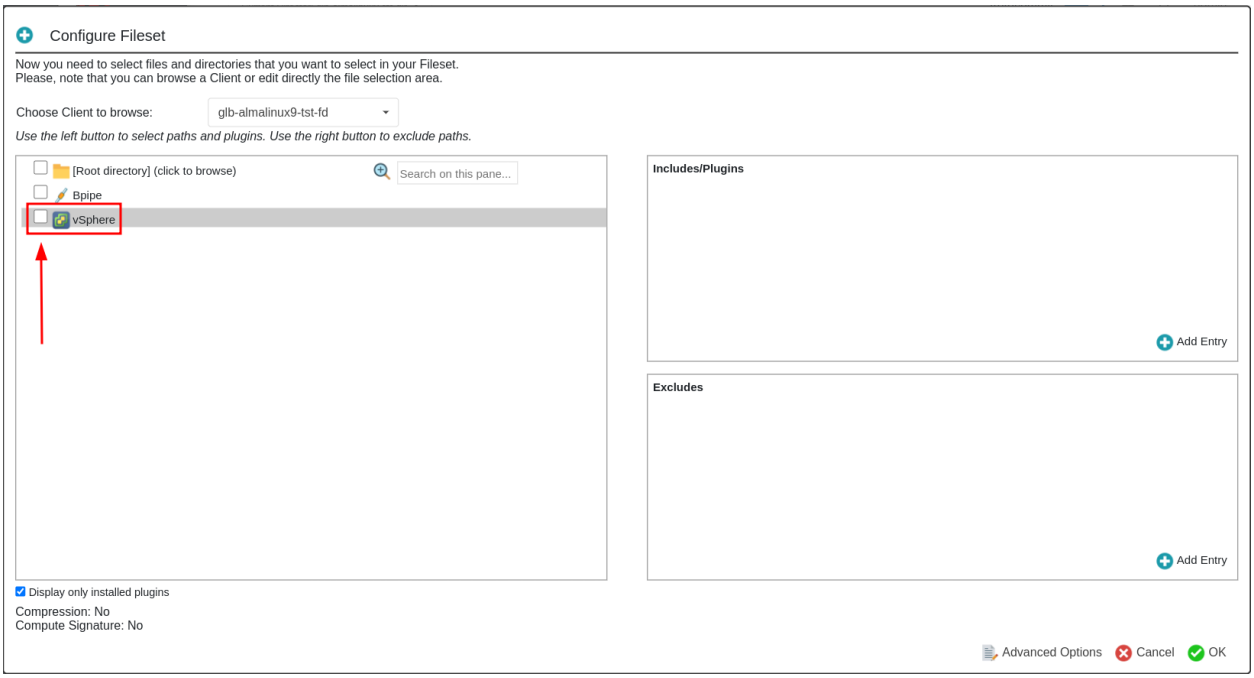

You will be taken to the "Plugin Configuration" dialog box. Here, after refreshing clicking the circular arrows (1), choose the 'vcenter-65' server from the "Server:" drop-down menu. Again, refresh the "Virtual Machine" drop-down menu (2) and click the down arrow. You will see a listing of all of the VMs managed by this vCenter server.

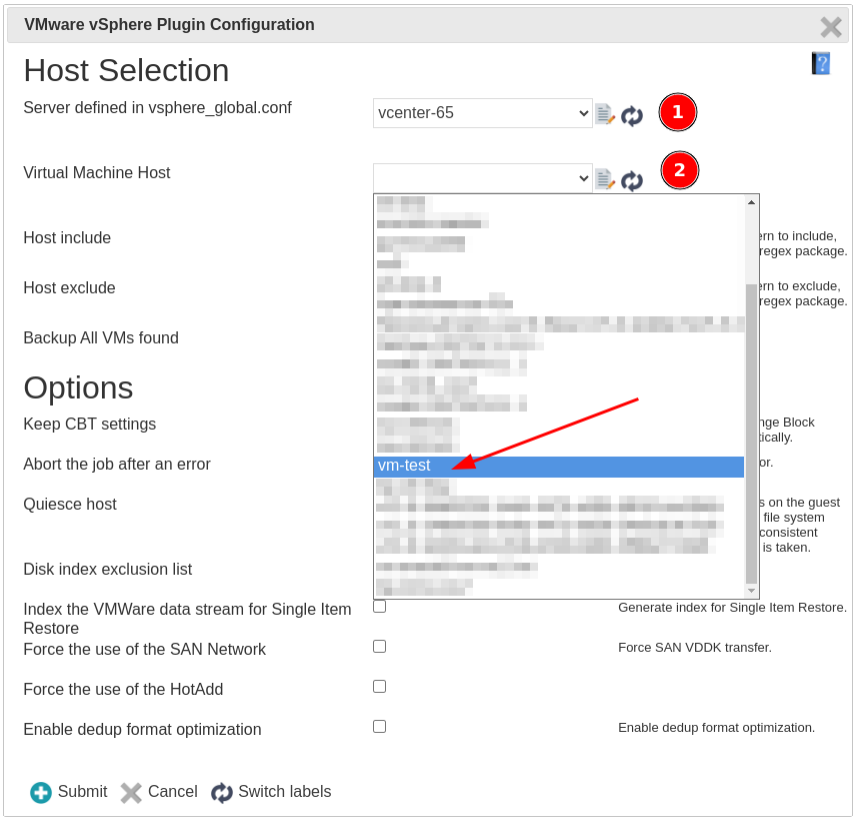

Since we are configuring a new FileSet and Job to backup one specific VM managed by this vCenter server, we will not use the “Host Include” and “Host Exclude” fields. These fields can be used used to backup or exclude VMs based on wildcards like: debian9-* or www-*.

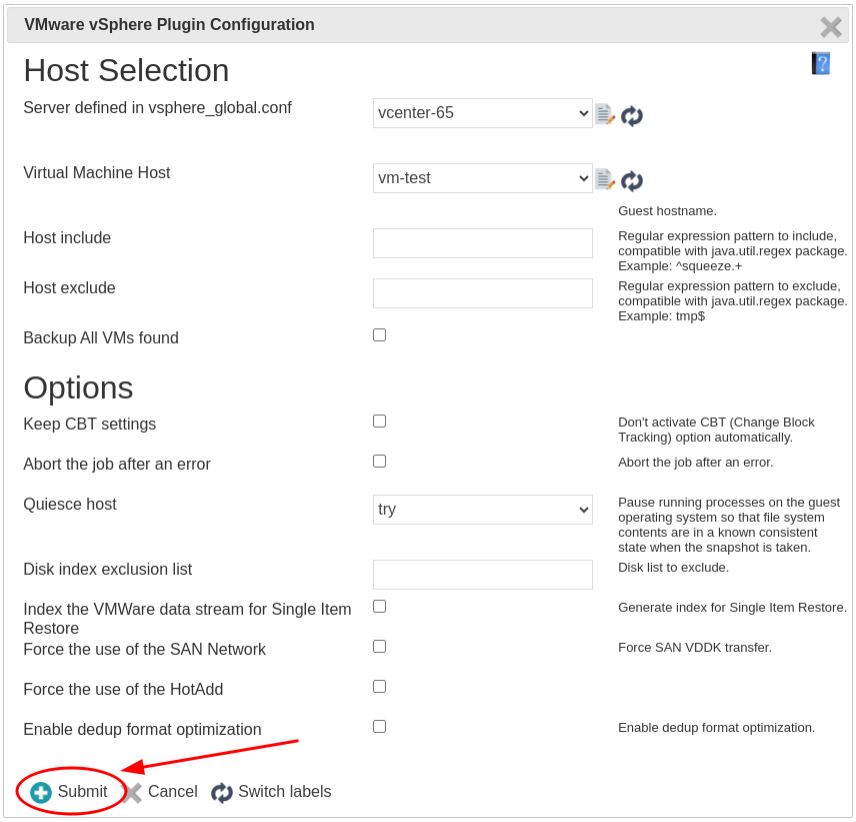

In this example, we will select the VM named “vm-test”, then click “Submit” at the bottom of the dialog box.

You will be taken to the “Configure Fileset” dialog box where you can see in the “Includes/Plugins” box what will be added to the FileSet’s “Include:” section. This is called the “vSphere plugin command line”. In this case, there are only three simple parameters:

1. The name of the Bacula plugin to use: “vsphere:”
2. The ESXi host or vCenter server to communicate with: “server=vc8.supportlab.lan”
3. The VM to be backed up: “host=vm-test”

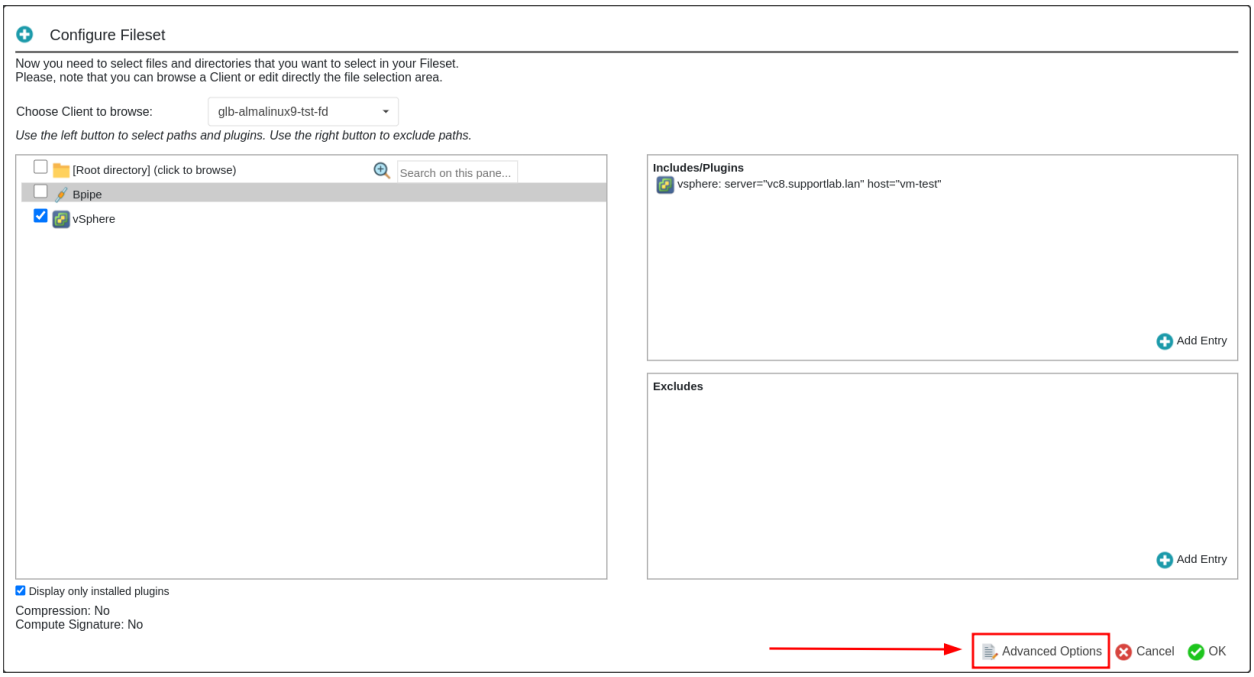

Just click “Advanced Options” here, and you will be taken to the new FileSet’s “Options:” dialog box. Here we will set an SHA1 “Signature” for all files backed up by this vSphere plugin backup FileSet. The rest of the options may be left as-is.

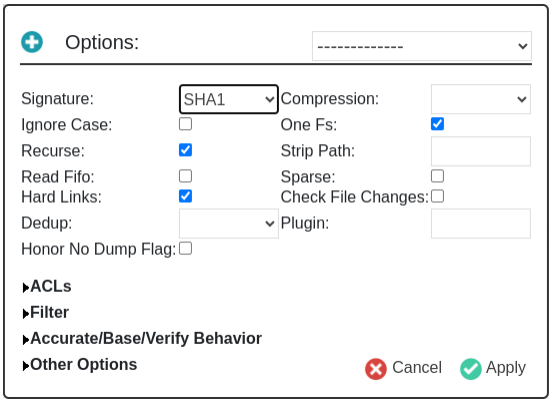

Click “Apply”, and you will be taken back to the “Add Fileset” dialog where you can see the newly created FileSet with the vSphere plugin command line and the optional SHA1 “Signature” which was added to the FileSet:

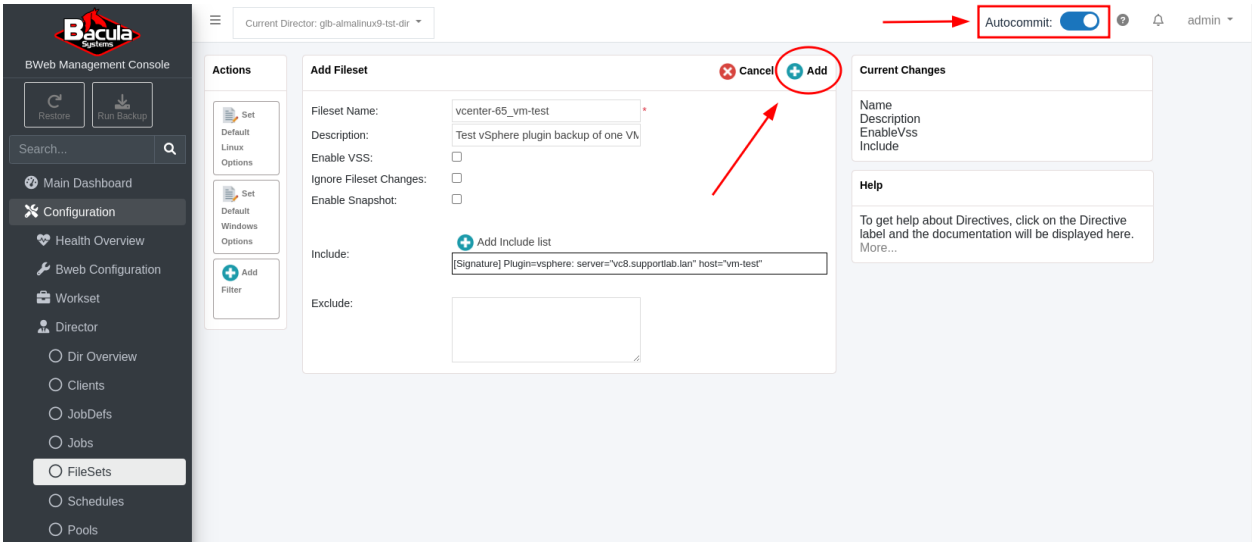

Make sure to toggle “Autocommit” in the top-right corner and click “+ Add” at the top of the “Add Fileset” dialog box. You will be taken to the full listing of all configured FileSets where the new ’vcenter-65_vm-test’ FileSet should be listed:

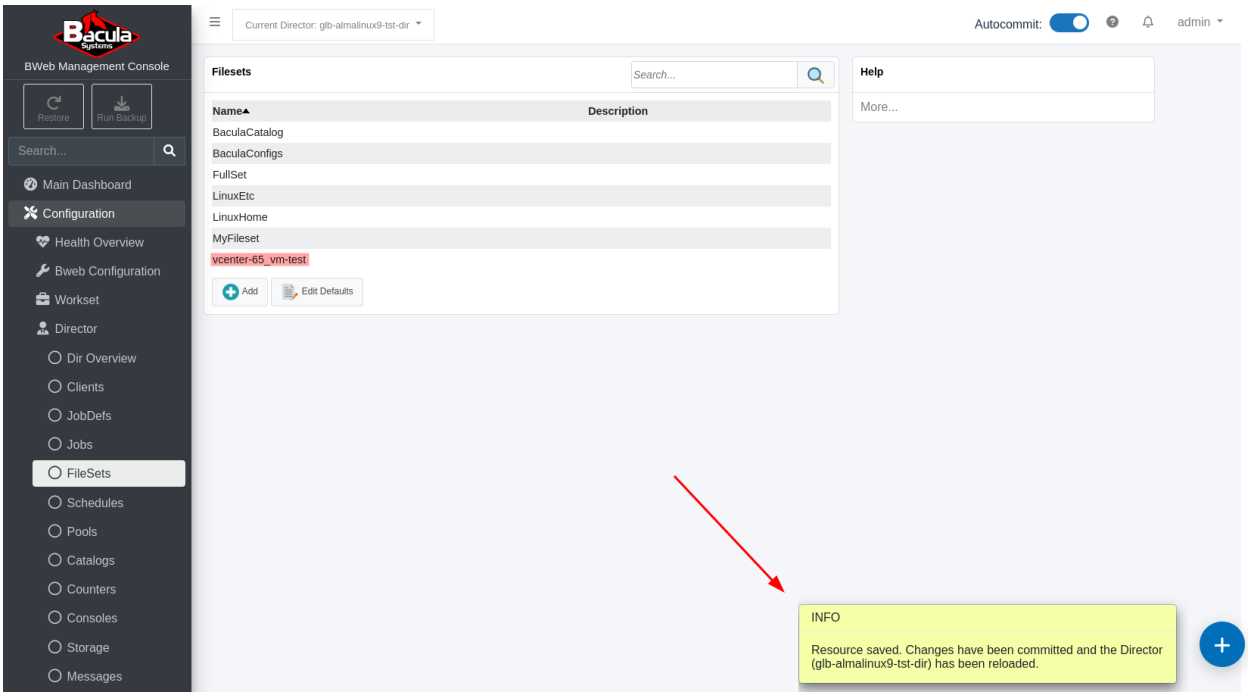

A pop-up window in the bottom-right corner will inform about the successful operation.

Add New Backup Job Which Uses New 'vcenter-65_vm-test' FileSet

On the BWeb main menu, expand "Configuration -> Director" and click "Jobs". You will be taken a page where all of the currently configured Jobs are listed.

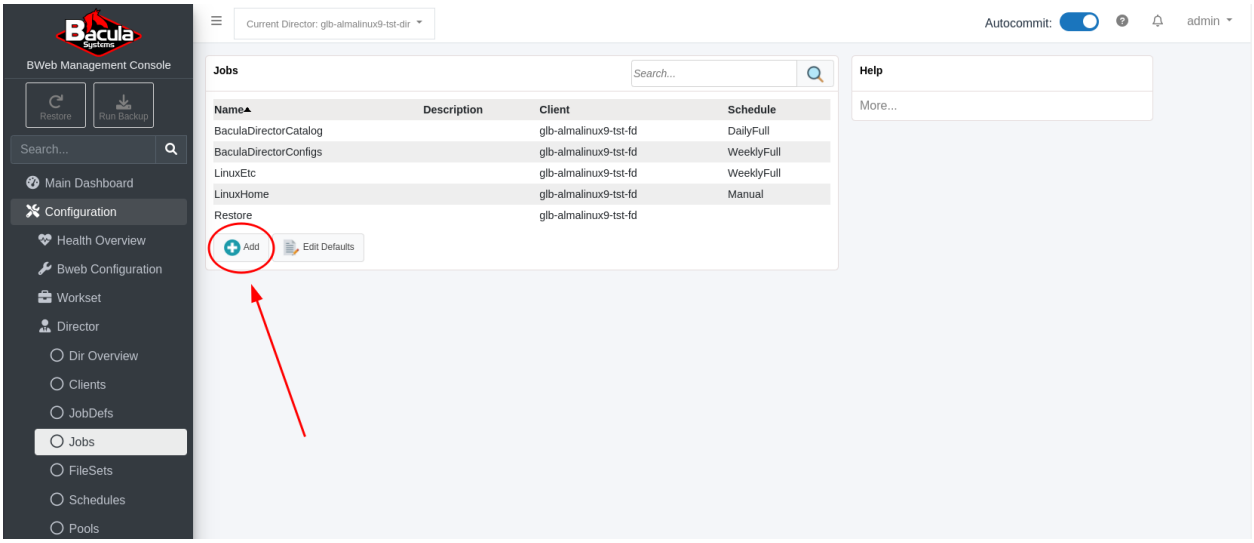

Click the "+ Add" at the bottom of the dialog box. You will be taken to a form where this new Job will be configured.

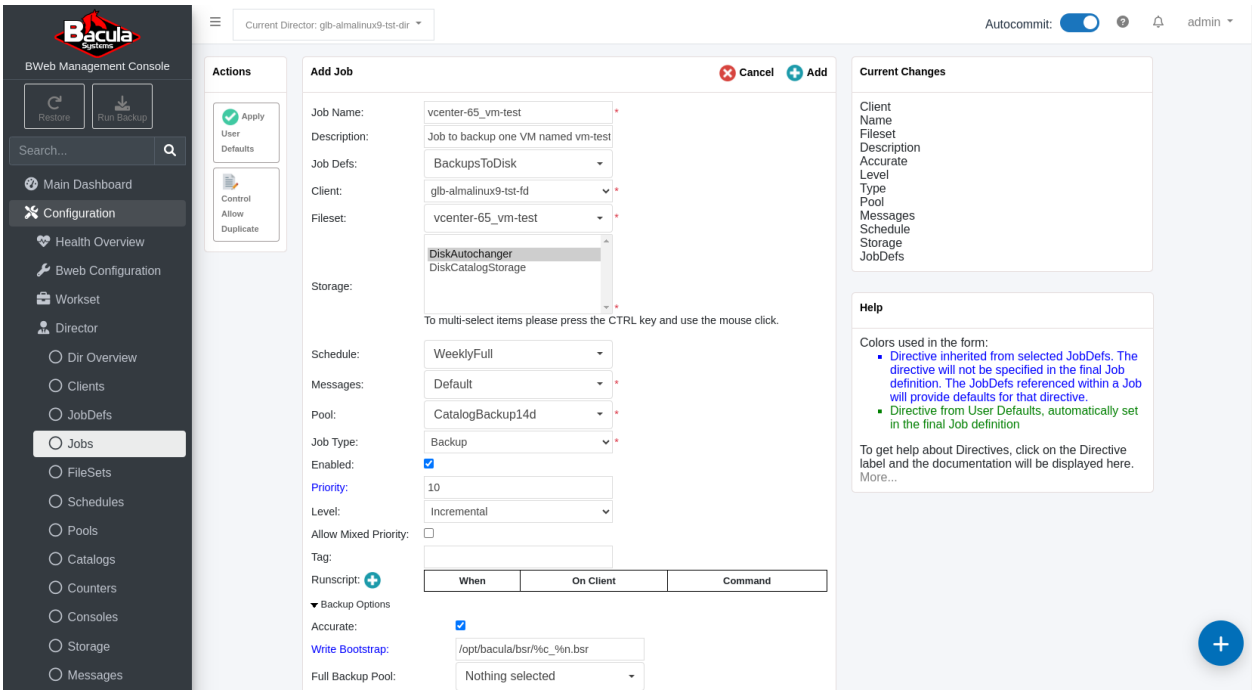

In this dialog box, the "Job Name:" field is filled in with a name that matches the Fileset we have previously configured. This is not required but can be a good practice to implement that makes it easy to understand what this Job does.

You may choose a pre-configured "JobDef" from the "JobDefs:" drop-down menu. This will automatically populate several of the other fields with some default settings inherited from the selected JobDef. Take a look at all of these and verify that they are OK for your needs. All the defaults from the JobDef can be overridden here.

NOTE: Since we are configuring a vSphere plugin backup job, the "Accurate:" option must be enabled. To display this option, click on "Backup Options" to open and expand this normally hidden section.

Click the checkbox next to “Accurate:”, then click the “+ Add” at the top of this dialog box to save this new Job. You should be taken back to the list of all configured jobs where this new job should be listed.

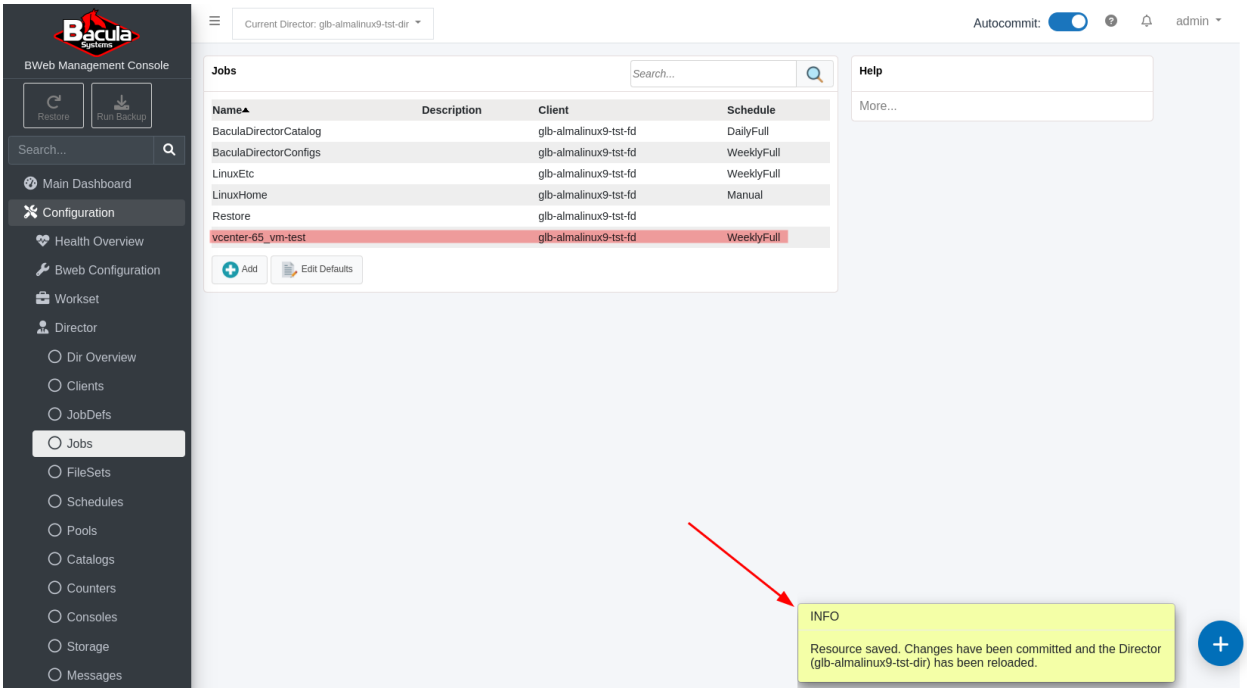

At this time, the new vSphere VM backup Job is available to be run manually, or via a schedule.

Running New vSphere Plugin Job

On the BWeb main menu, click “Run Backup” in the upper-left corner. In the “Job name:” drop-down menu, you should be able to select your new VMware backup job. Follow the steps in the manual job run wizard to run this job.

See also:

Go back to:

- *Configuring New vSphere Host*

Go to:

- *Creating Automated vSphere Backup Environment Using BWeb*
- *Creating Automated vSphere Backup Environment Using Command Line Scripting*

Go back to the *main BWeb VMware Center Module page*.

Go back to the *main vSphere Plugin page*.

Creating Automated vSphere Backup Environment Using BWeb

The purpose of this wizard is to create an automated configuration whereby all VMs managed by a particular ESXi host or vCenter server are backed up. As new VMs are added, they will automatically be backed up, and as VMs are decommissioned, they will be disabled (or removed) from the backup configuration.

It is possible to exclude VMs from this automatic configuration based on their names, or VM tags, etc., but the steps to do this will come after creation of the default “Backup all VMs” configuration is complete.

On the BWeb main menu, expand “Virtual Machines” and click “VMware Center”. You will be taken to the page with the vSphere Hosts listing, which should currently have at least one vSphere host defined. If there are no hosts in this list, follow the *Configuring New vSphere Host article* to create one and assign it to a Bacula Client first. Click on the “Backup multiple virtual machines wizard” button on the left.

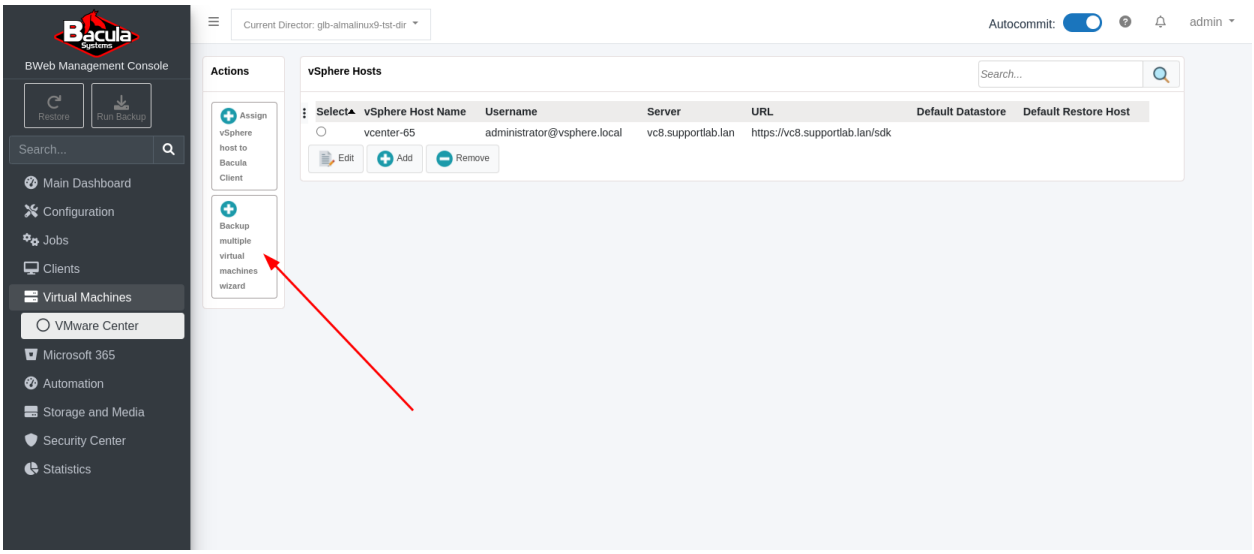

You will be taken to Step 1 of 4 of the Wizard.

Select a vSphere host from the “vCenter/ESXi host” drop-down list.

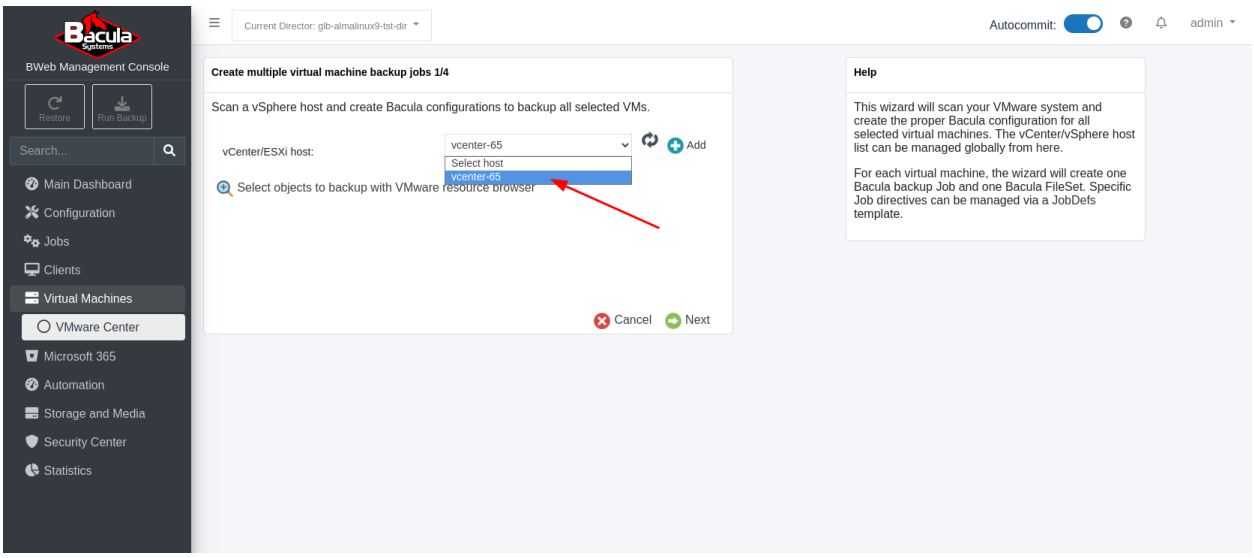

Click the circular arrows icon to obtain information about the VMs and Storages managed by this vCenter server.

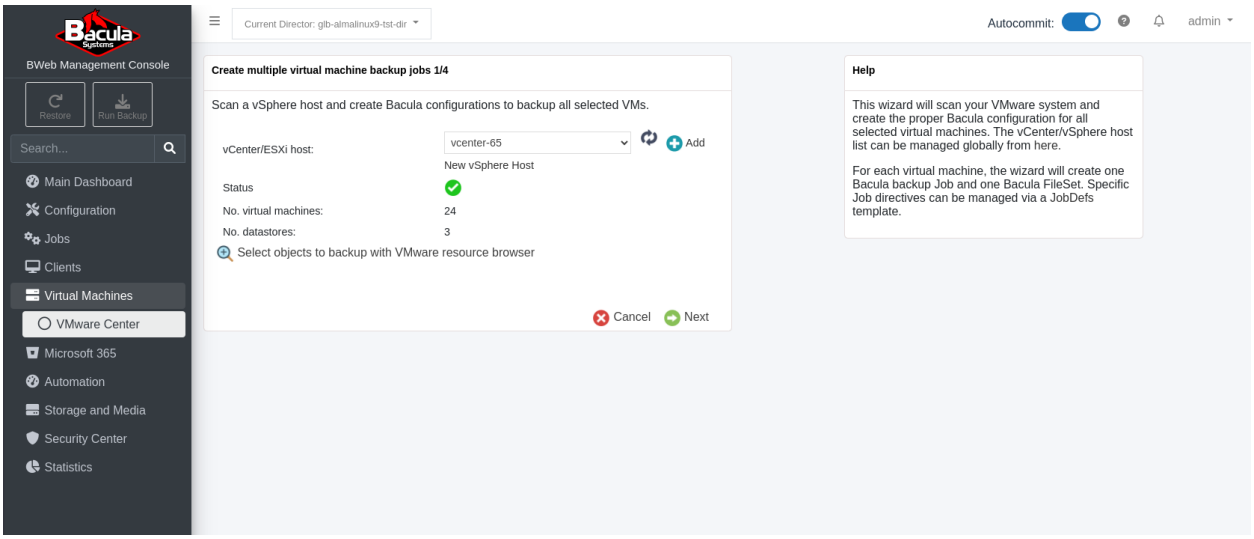

Click “Select objects to backup with VMware resource browser” and you will be presented with a “VMware resource browser” pop-up dialog. Since this “vcenter-65” vSphere host is a vCenter server, there will initially be a top-level “Datacenter” listed. In this case the Datacenter is called “vc8”.

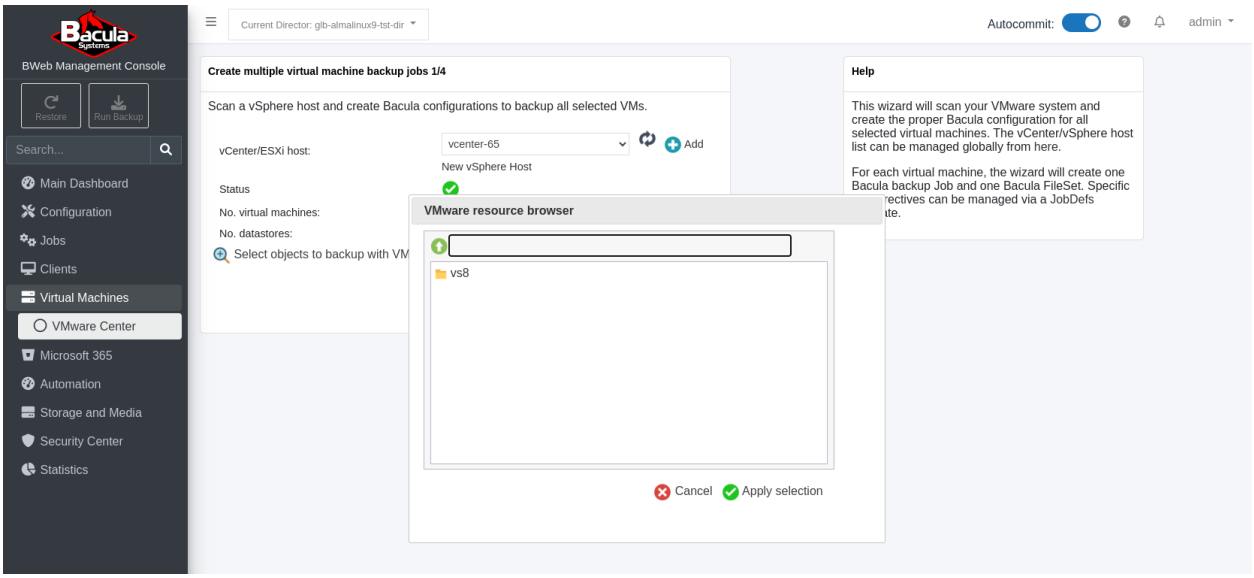

Click on the Datacenter icon to reveal the VM, ResourcePool, and Datastore folders.

In this example we will be focusing on backing up all VMs managed by this vCenter server.

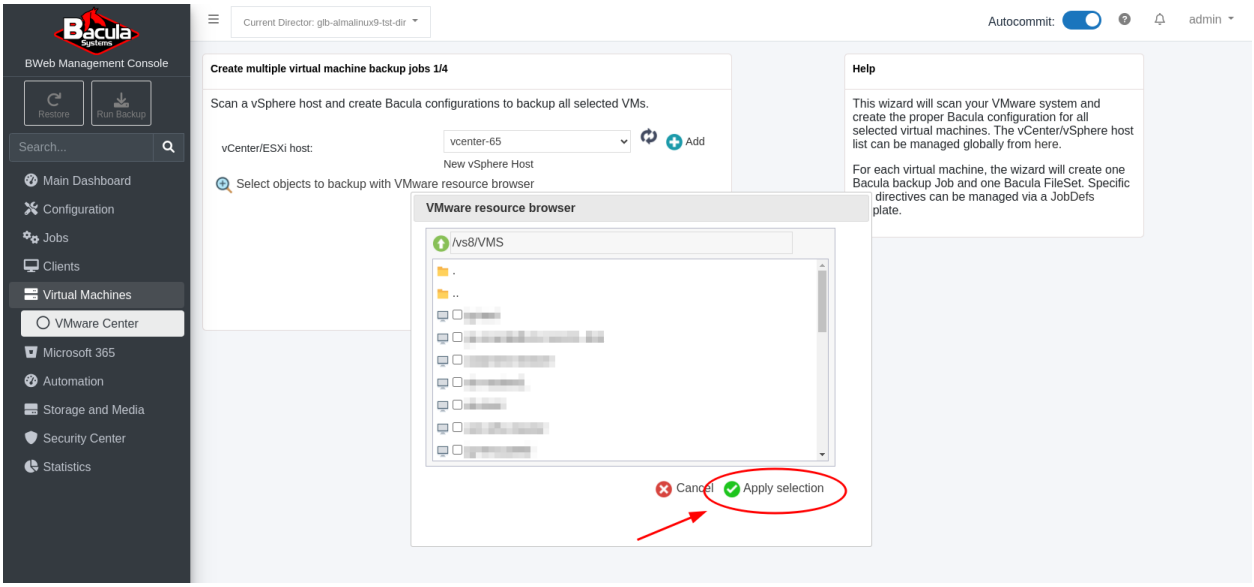

Click on the VMs folder icon to see a list of VMs managed by this vCenter host. Do not check any boxes, and simply click the “Apply selection” green check icon.

You will be taken back to the “Create multiple virtual machine backup jobs 1/4” dialog box. Just click “Next” here.

You will be taken to the step 2/4 dialog box where the options for the FileSets to be created may be set.

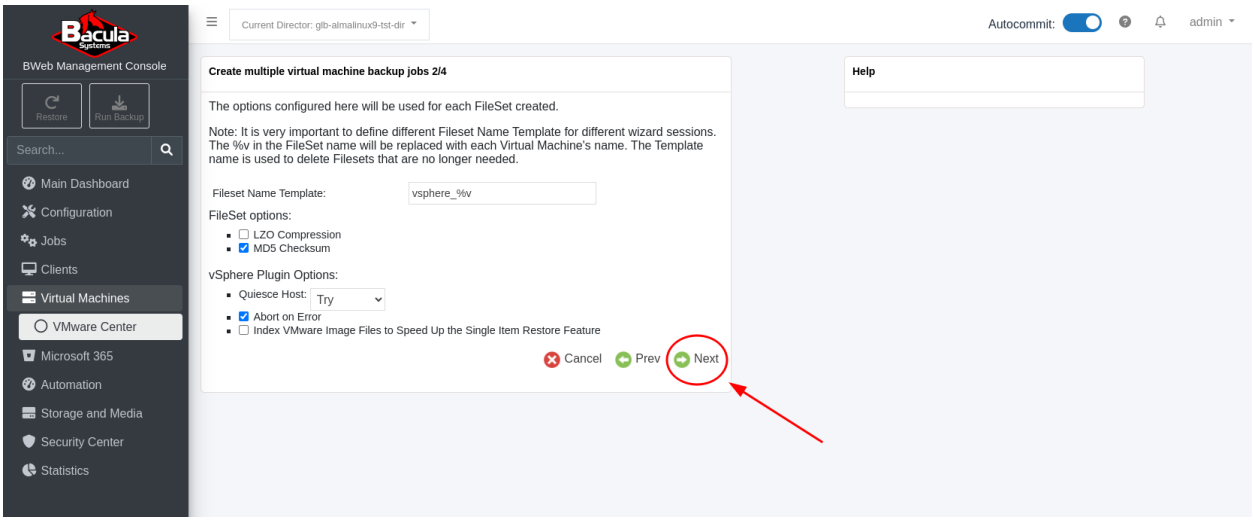

In the “Fileset Name Template:”, the “%v” will be replaced with the name of each VM found.

Leave the settings at their defaults for now and click “Next”.

You will be taken to the step 3/4 dialog box where the options for the Jobs to be created may be set.

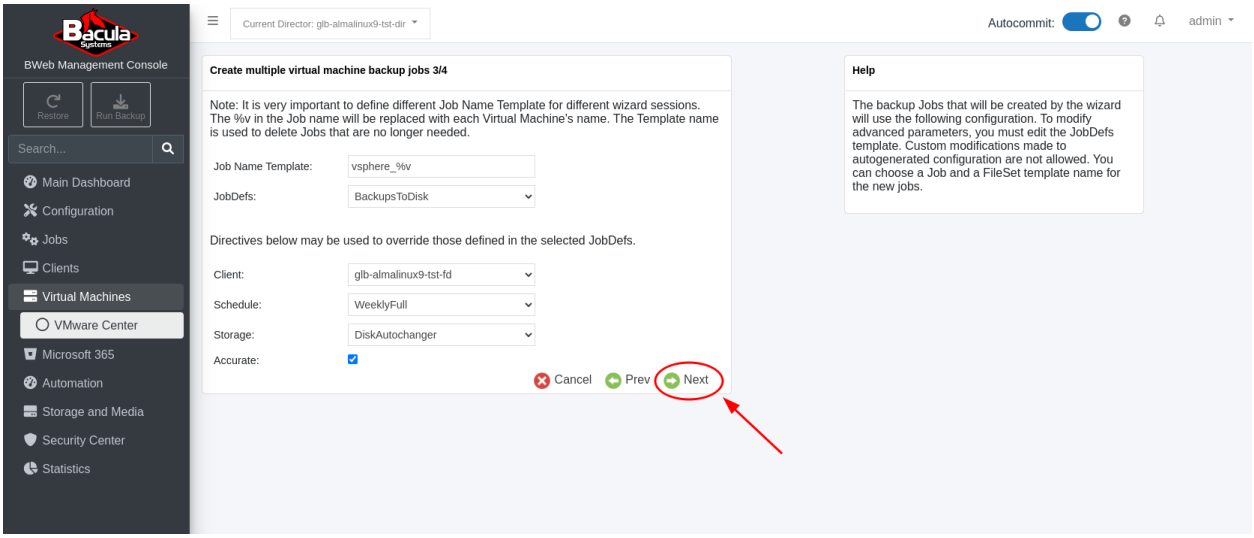

In the “Job Name Template:”, the “%v” will be replaced with the name of each VM found.

Be sure to choose the appropriate Client, Schedule, and Storage for these jobs. If you are performing Incremental or Differential backup of your VMware VMs, then the “Accurate” mode must be enabled, so normally it is recommended to leave this checkbox checked. Click “Next”.

You will be take to the step 4/4 dialog box.

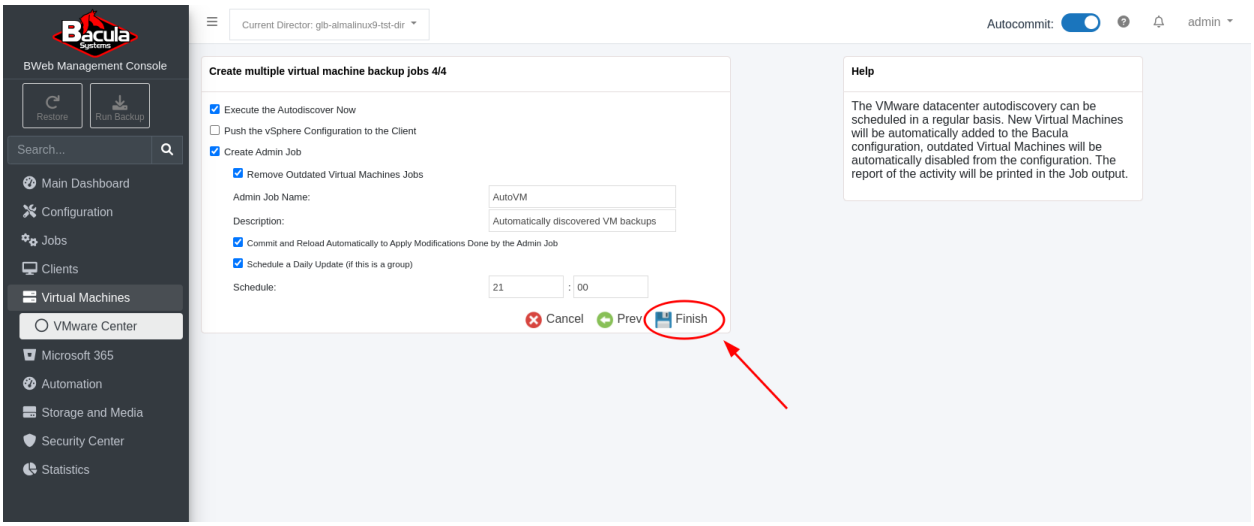

Here, leave the “Execute the Autodiscover Now” checkbox checked, leave the “Push the vSphere Configuration to the Client” Checkbox unchecked, and then check the “Create Admin Job” checkbox - more options will be revealed.

Enter an appropriate “Admin Job Name” and “Description” for the automatic discovery job, then check the “Commit and Reload…” and “Schedule a Daily Update” box - more options will be revealed.

Set the “Schedule” time before your nightly VMware plugin backup jobs are expected to run. This will ensure that the work of creating new Jobs/FileSets for newly discovered VMs and disabling jobs for decommissioned VMs happens before the VMware backup jobs are queued.

Click “Finish”.

After a few seconds, the “Result executing autodiscover now” dialog box will pop up showing what has been done.

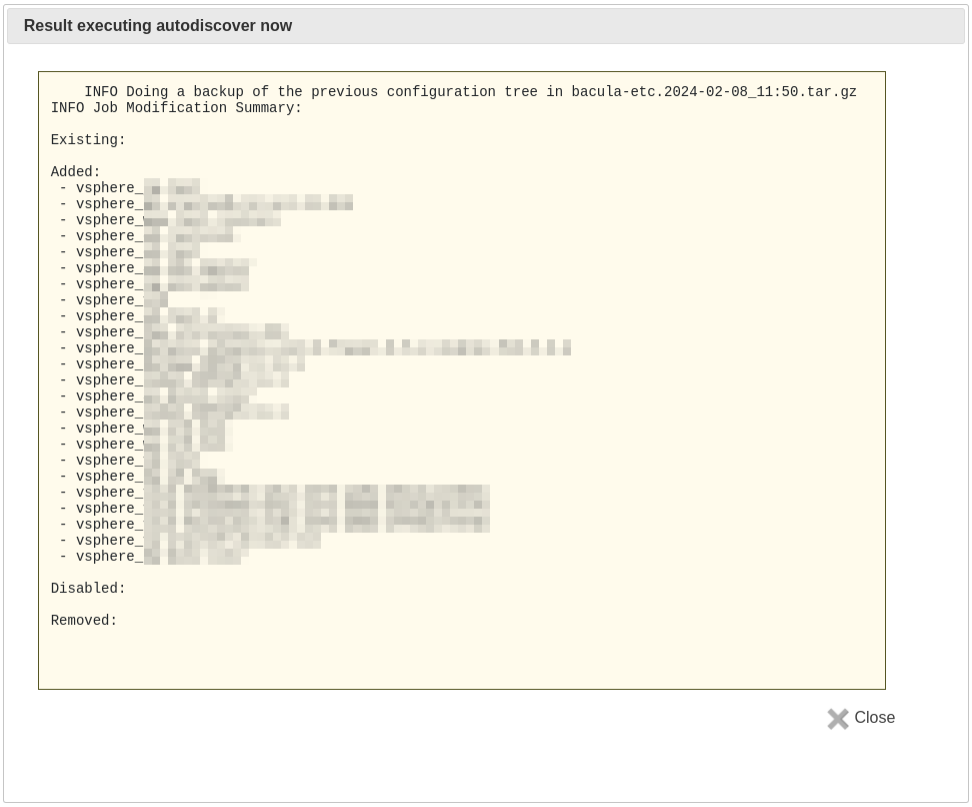

Notice that BWeb has added Jobs and FileSets for All VMs managed by this vCenter server called “vcenter-65”, and there were no existing Jobs/FileSets discovered, and none removed, nor disabled.

At this time, you may click on “Jobs” or “FileSets”, or “Schedules” to see the new resources created by the Autodiscovery process.

Exit configuration mode by clicking on the “BWeb Management Console” at the top left of the main menu.

Expand “Jobs” and click “Next Jobs”. If you had selected a schedule in step 4/4 above, then you should see your new “vsphere_XXXX” jobs in this list of “Next Jobs”.

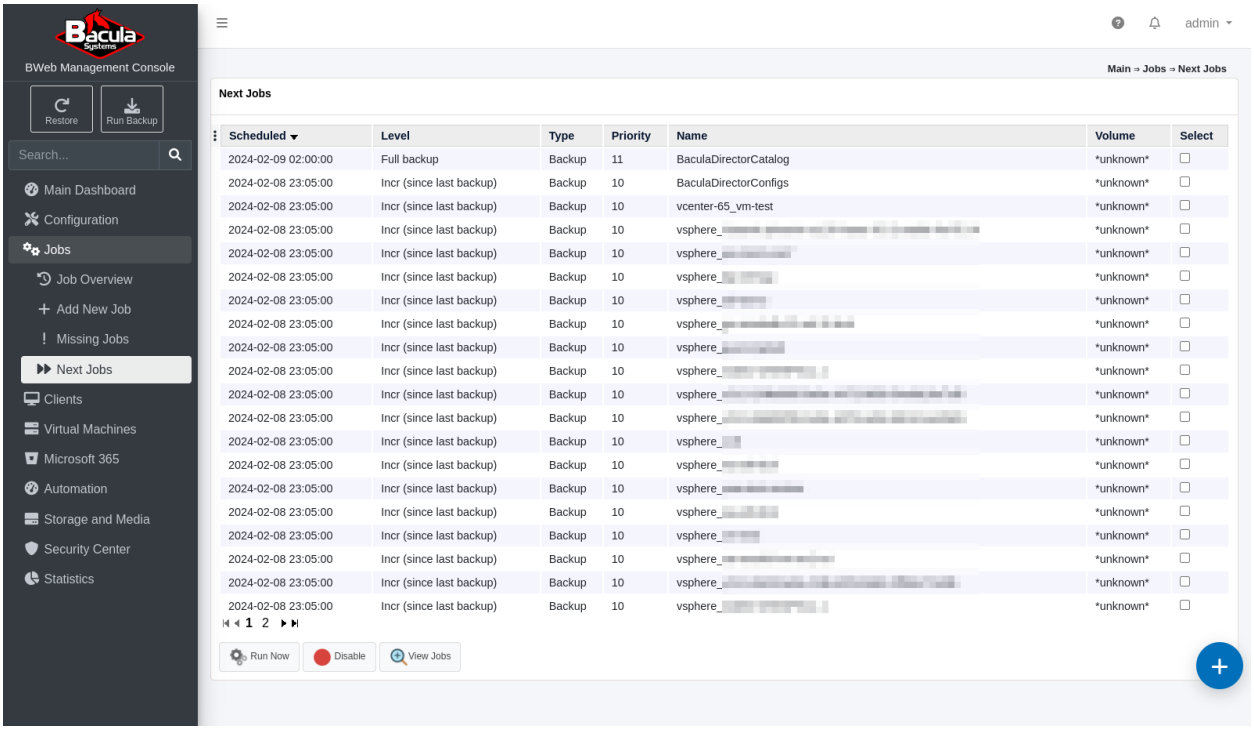

See also:

Go back to:

- *Configuring New vSphere Host*
- *Creating vSphere Plugin Backup Job and FileSet in BWeb*

Go to:

- *Creating Automated vSphere Backup Environment Using Command Line Scripting*

Go back to the *main BWeb VMware Center Module page*.

Go back to the *main vSphere Plugin page*.

Creating Automated vSphere Backup Environment Using Command Line Scripting

---

Important: Since Bacula version 16.0.7, a new solution has been introduced to replace the scan_datacenter tool. It is highly recommended to use the new solution - Automatic Objects Integration (Scan Plugin) as the scan_datacenter will soon be deprecated. See an example for vSphere.

There is a script called “scan_datacenter.pl” that is integral to BWeb’s “VMware Center” located in “/opt/bweb/bin”. In Step 4 of 4 in the *BWeb Automated VMware Backups section* of this document, we have already created an Admin type job called “Admin_AutoVM”, and scheduled it to run daily at 21:00.

If we take a look at this Job in BWeb (Configuration –> Director –> Jobs –> Admin_AutoVM), we can see that the scan_datacenter.pl script has been automatically configured to be called in a RunScript section of this Admin job. Runscripts are job resources that can run any shell script before or after a backup job or administrative task as in this case.

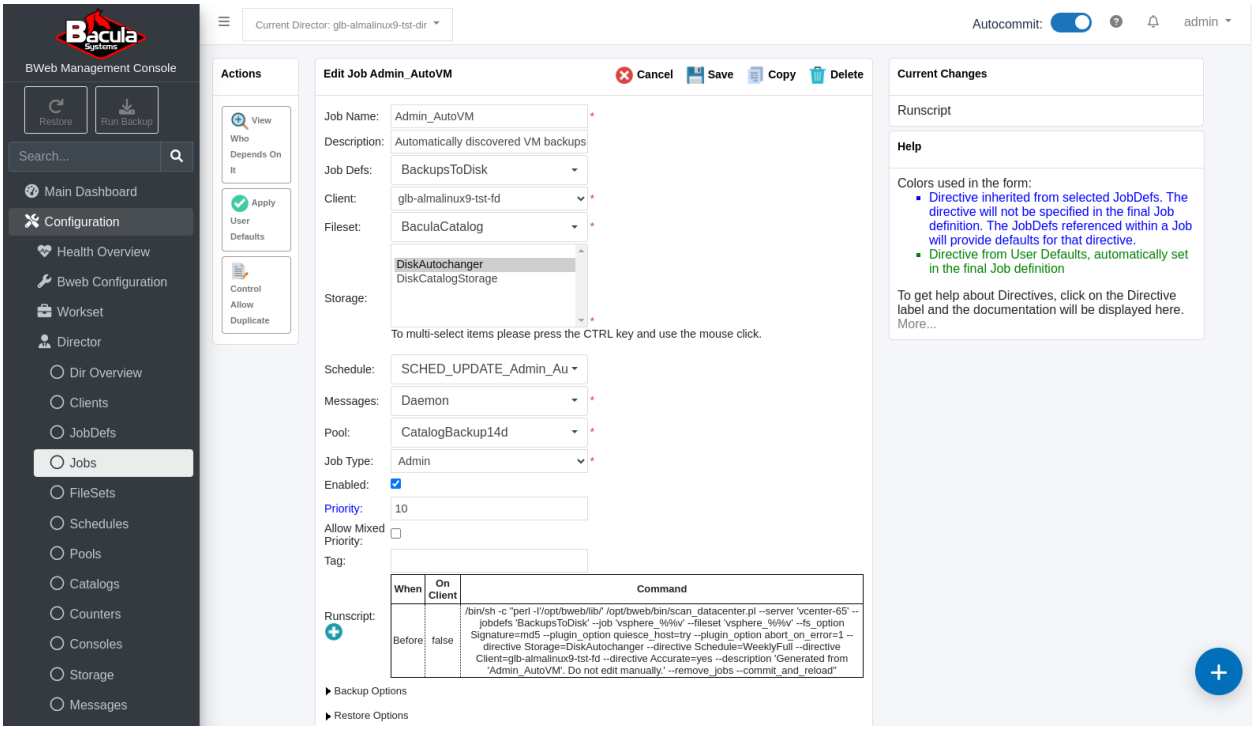

The scan_datacenter.pl script is called with some specific command line parameters that are used to create the new FileSets and Jobs for each VM found.

We can simply cut and paste this working example into a new custom script which can be run manually, run automatically via cron, or via a Bacula Admin type Job as in the Admin_AutoVM Admin Job example.

Here is the command in the Admin Job's RunScript section as it was created by BWeb's VMware Center's "Backup multiple virtual machines wizard" in the *previous section* of this guide:

```
/bin/sh -c "perl -I'/opt/bweb/lib/' /opt/bweb/bin/scan_datacenter.pl
--server 'vcenter-65' --jobdefs 'DefaultJob' --job 'vsphere_%%v'
--fileset 'vsphere_%%v' --fs_option Signature=md5 --plugin_option
quiesce_host=try --plugin_option abort_on_error=1 --directive
Storage=File1 --directive Schedule=WeeklyFull --directive
Client=glb-almalinux9-tst-fd --directive Accurate=yes --description
'Generated from 'Admin_AutoVM'. Do not edit manually.' --remove_jobs
--commit_and_reload"
```

In this example, there are a number of command line options being used. For example, we can see that for the FileSet name and the Job name, each will be prefaced by "vmware_" - the name of the VM found. This means that each Job/FileSet pair will have the same name, making the correlation between Jobs and FileSets for each VM backup easy.

There are a few Job-specific "–directive" settings specified too. Importantly, "Accurate mode" is enabled for these Jobs - a requirement for Differential and Incremental VMware VM backups. For the "–directive" values, use the upper camel case (camel case with the first letter capitalised) writing method, for example:

```
--directive AllowDuplicateJobs=yes --directive SpoolSize=100000000
```

Also, we can see there are a few vSphere plugin-specific settings specified by using the "–plugin_option" command line parameter multiple times.

And finally, because this Admin type Job will run once per day at 21:00, there are two options specified so that new Jobs and FileSets are automatically committed to the Director's configuration "–commit_and_reload", and also, any

VMs that no longer exist have their FileSet and Job resources removed from the Director's configuration due to the "–remove_jobs" option.

A full listing of all available command line options to the scan_datacenter.pl script can be seen simply by running this script with no command line parameters.

The `scan_datacenter.pl` script was introduced with **BWeb** version 8.10. Interim versions required to download and install the Perl SDK from VMware (not longer required from version 12.2.4).

See also:

Go back to:

- *Configuring New vSphere Host*
- *Creating vSphere Plugin Backup Job and FileSet in BWeb*
- *Creating Automated vSphere Backup Environment Using BWeb*

Go back to the *main BWeb VMware Center Module page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *vSphere Plugin Configuration*

Go to:

- *vSphere Plugin Backup Job Configuration*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

5.3 vSphere Plugin Backup Job Configuration

The following chapter presents the information on how to configure a vSphere Plugin Backup Job.

VMware Automatic Integration

Important: Since Bacula version 16.0.7, a new solution has been introduced to replace the scan_datacenter tool. It is highly recommended to use the new solution - Automatic Objects Integration (Scan Plugin) as the scan_datacenter will be deprecated. See an example for vSphere.

The `scan_datacenter` BWeb Management Suite module can help you set up the Bacula configuration needed, and protect your VMs in a highly automated way. The `scan_datacenter` module will connect the vCenter/ESXi server to list the VMs to backup.

Note: The *BWeb VMware Center article* describes in detail how to set up the `scan_datacenter` module.

See also:

Go to:

- *Job Example*

- *FileSet Example*
- *Testing Fileset*
- *Excluding Disk*
- *vSphere FileSet Plugin Command Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Job Example

The `Accurate` option is mandatory in the backup Job resource when running Incremental/Differential backup jobs with the vSphere plugin.

```
Job {
  Name = vSphereBackup
  FileSet = vSphere
  Accurate = yes
  ...
}
```

See also:

Go back to:

- *VM Automatic Integration*

Go to:

- *FileSet Example*
- *Testing Fileset*
- *Excluding Disk*
- *vSphere FileSet Plugin Command Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

FileSet Example

This section presents various FileSet examples.

```
FileSet {
 Name = vSphere
 Include {
  Options { Signature=MD5 }
  Plugin = "vsphere: host=guest1"
 }
}
```

```
FileSet {
 Name = vSphere
 Include {
```

(continues on next page)

(continued from previous page)

```
  Options { Signature=MD5 }
  Plugin = "vsphere: host=vm-401"
 }
}
```

Important: The vSphere Plugin is not compatible with the `sparse` FileSet option.

See also:

Go back to:

- *VM Automatic Integration*
- *Job Example*

Go to:

- *Testing Fileset*
- *Excluding Disk*
- *vSphere FileSet Plugin Command Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Testing FileSet

You can use the bconsole `estimate` command to test your FileSet.

```
*estimate listing job=vmware1
Using Catalog "MyCatalog"
Connecting to Client 127.0.0.1-fd at 127.0.0.1:9102
-rwx------  0 root  root           0 18:22:16  /@vsphere/3-squeeze2
-rw-r-----  1 root  root        1923 18:22:19  /@vsphere/3/vmbkp_generation.profile
-rw-r-----  1 root  root        4693 18:22:19  /@vsphere/3/3.ovf
-rw-------  1 root  root  2147483648 18:22:20  /@vsphere/3/0.bvmdk
-rw-------  1 root  root   104857600 18:22:21  /@vsphere/3/1.bvmdk
-rwx------  1 root  root           0 18:22:26  /@vsphere/79-squeeze.esx
-rw-r-----  1 root  root        1806 18:22:29  /@vsphere/79/vmbkp_generation.profile
-rw-r-----  1 root  root        4704 18:22:28  /@vsphere/79/79.ovf
-rw-------  1 root  root   104857600 18:22:29  /@vsphere/79/0.bvmdk
-rw-------  1 root  root   104857600 18:22:30  /@vsphere/79/1.bvmdk
2000 OK estimate files=10 bytes=2,462,069,574
```

See also:

Go back to:

- *VM Automatic Integration*
- *Job Example*
- *FileSet Example*

Go to:

- *Excluding Disk*
- *vSphere FileSet Plugin Command Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Excluding Disk

Exclude Disk with vSphere Console

To exclude a specific disk, you can activate the **independent** mode for the disk in the vSphere console.

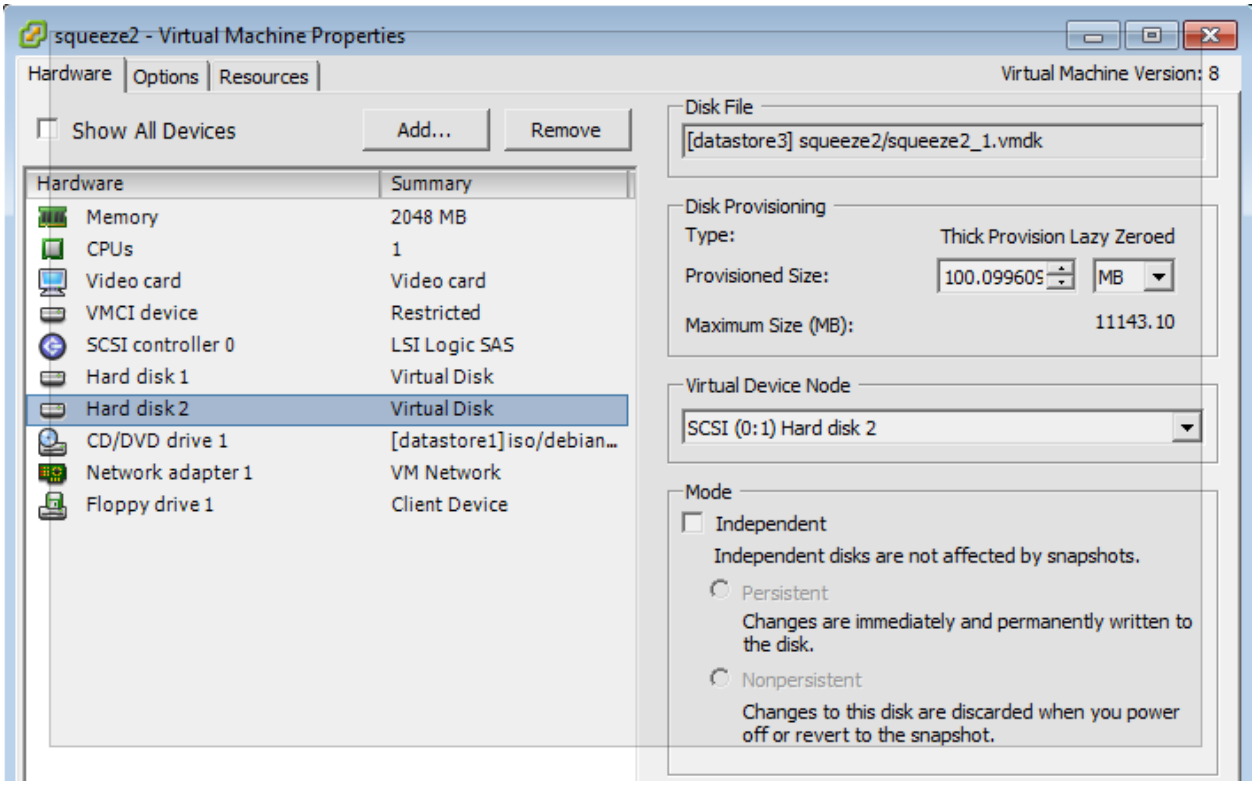

See also:

Go to *Exclude Disk with FileSet*.

Go back to the *Excluding Disk*.

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

Exclude Disk with FileSet

To find the `diskid` to use in the `disk_exclude` option (see the options *table*), it is possible to use the `estimate listing` command. `0.bvmdk` is the image of the diskid `0`.

```
# Will exclude the disk 0 and the disk 2 from "myvm"
Plugin = "vsphere: disk_exclude=0,2 host=myvm"
```

See also:

Go to *Exclude Disk with FileSet*.

Go back to the *Exclude Disk with vSphere Console*.

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *VM Automatic Integration*
- *Job Example*
- *FileSet Example*
- *Testing Fileset*

Go to:

- *vSphere FileSet Plugin Command Options*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

vSphere FileSet Plugin Command Options

Table 3: vSphere FileSet Plugin Command Options

Option	Required	Default	Info	Example
host	No		Guest hostname or a MoRef	host=srv1, host=vm-25
host_include	No		Guest pattern to include	host_include=srv3
host_exclude	No		Guest pattern to exclude	host_exclude=srv[12]
disk_exclude	No		Disk list to exclude (available since version 8.4.8)	disk_exclude=0,2,4
disk_include	No		Disk list to include (available since version 16.0.8). If used, other disks will be exluded	disk_include=0,2,4
keep_cbt	No		Don't try to activate CBT	keep_cbt
dedup_format	No	No	Control VDDK analyzer for Global Endpoint Deduplication (available since version 12.6).	dedup_format=yes
quiesce_host	Yes	Try	Choose to quiesce guest before the snapshot (Try, yes, no)	quiesce_host=no
server	No	vsphere	Specify a vsphere server	server=vsrv2
debug	No		Enable debug	debug
abort_on_error	No		Abort the job after an error (available since version 8.2.5)	abort_on_error
update_timeout	No	900 seconds	Change initial update timeout	
index	No		Generate index for Single Item Restore (available since version 8.6)	index
force_san	No		Force SAN VDDK transfer (available since version 10.2.3)	

Note that `host_include` and `host_exclude` are Java regexp patterns.

The `index` feature will generate records in the Catalog to quickly seek to a given block in the backup stream. The granularity of the index can be controlled with the Storage Daemon device directive `MaximumFileIndex`. The default value is 100MB.

If used together disk_include and disk_exclude affecting the same disk, disk_include takes precedence and the disk will be considered into the backup.

Going back to the Configuration chapter

To go back to the main Configuration chapter, click *here*.

See also:

Go back to:

- *VM Automatic Integration*
- *Job Example*
- *FileSet Example*
- *Testing Fileset*
- *Excluding Disk*

Go back to the *main vSphere Plugin Configuration page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *vSphere Plugin Configuration*
- *BWeb VMware Center Module*

Go back to the *main Configuration page*.

Go back to the *main vSphere Plugin page*.

---

vSphere Global Configuration Options

To browse through vSphere Global Configuration Options, click *here*.

---

vSphere FileSet Plugin Command Options

To browse through vSphere FileSet Plugin Command Options, click *here*.

See also:

Go back to:

- *Scope*
- *Features*
- *Backup Strategies*
- *Installation*

Go to:

- *Operations*
- *Limitations*
- *Troubleshooting*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

6 Operations

This article describes details regarding backup, restore, quiescing guests, VM Instant Recovery and Single Item Restore with Bacula Enterprise vSphere Plugin.

6.1 Backup

Example of a Backup Job

```
*list joblog jobid=21509
+-----------------------------------------------------------------------------------------
↪------------+
| logtext                                                                                 ␣
↪            |
+-----------------------------------------------------------------------------------------
↪------------+
| bp-vsir-bweb102-dir JobId 21509: No prior or suitable Full backup found in catalog.␣
↪Doing FULL backup. |
| bp-vsir-bweb102-dir JobId 21509: Start Backup JobId 21509, Job=000_bp-o9-hap_vsphere.
↪2023-11-15_13.45.06_49 |
| bp-vsir-bweb102-dir JobId 21509: Connected to Storage "LocalGED" at 10.0.98.5:9103␣
↪with TLS          |
| bp-vsir-bweb102-dir JobId 21509: Using Device "LocalGED-04" to write.                ␣
↪              |
| bp-vsir-bweb102-dir JobId 21509: Connected to Client "bp-vsir-bweb1004-fd" at bp-vsir-
↪bweb1004:9102 with TLS |
| bp-vsir-bweb102-fd JobId 21509: Connected to Storage at 10.0.98.5:9103 with TLS      ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Volume "ged-3474" previously written, moving to end of␣
↪data.         |
| bp-vsir-bweb102-sd JobId 21509: Ready to append to end of Volume "ged-3474" size=5,347,
↪408            |
| bp-vsir-bweb102-fd JobId 21509: Backup "bp-o9-hap" (vm-7253) start.                  ␣
↪              |
| bp-vsir-bweb102-fd JobId 21509: Activating vSphere "Change Tracking System"          ␣
↪              |
| bp-vsir-bweb102-fd JobId 21509: To activate properly CBT, a Snapshot creation/deletion␣
↪cycle will now be performed. |
| bp-vsir-bweb102-fd JobId 21509: Create snapshot name 000_bp-o9-hap_vsphere.2023-11-15_
↪13.45.06_49 succeeded. |
| bp-vsir-bweb102-fd JobId 21509: There are 2 disks.                                   ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3474" Bytes=1,137,843,500␣
↪Blocks=88 at 15-Nov-2023 13:47. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3478"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3478" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3478" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 13:47. |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3478" Bytes=1,077,411,747␣
↪Blocks=15 at 15-Nov-2023 13:53. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3479"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3479" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3479" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 13:53. |
```

(continues on next page)

(continued from previous page)

```
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3479" Bytes=1,113,849,764␣
↪Blocks=15 at 15-Nov-2023 13:58. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3480"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3480" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3480" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 13:58. |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3480" Bytes=1,100,546,012␣
↪Blocks=18 at 15-Nov-2023 14:03. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3484"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3484" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3484" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 14:03. |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3484" Bytes=1,075,511,280␣
↪Blocks=16 at 15-Nov-2023 14:07. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3485"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3485" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3485" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 14:07. |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3485" Bytes=1,129,119,668␣
↪Blocks=15 at 15-Nov-2023 14:11. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3486"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3486" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3486" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 14:11. |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3486" Bytes=1,091,633,114␣
↪Blocks=15 at 15-Nov-2023 14:15. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3487"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3487" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3487" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 14:15. |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3487" Bytes=1,089,667,036␣
↪Blocks=20 at 15-Nov-2023 14:18. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3488"                          ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3488" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3488" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 14:18. |
| bp-vsir-bweb102-sd JobId 21509: End of medium on Volume "ged-3488" Bytes=1,127,284,756␣
↪Blocks=22 at 15-Nov-2023 14:22. |
| bp-vsir-bweb102-dir JobId 21509: Recycled volume "ged-3489"                          ␣
↪              |
```

(continues on next page)

(continued from previous page)

```
| bp-vsir-bweb102-sd JobId 21509: Recycled volume "ged-3489" on Dedup device "LocalGED-04
↪" (/bck_2/ged), all previous data lost. |
| bp-vsir-bweb102-sd JobId 21509: New volume "ged-3489" mounted on device "LocalGED-04"␣
↪(/bck_2/ged) at 15-Nov-2023 14:22. |
| bp-vsir-bweb102-fd JobId 21509: VDDK Transport "nbdssl" selected                    ␣
↪              |
| bp-vsir-bweb102-fd JobId 21509: Dump vmdk 6000C291-8cef-5868-9f15-bea1f333c165␣
↪succeeded.            |
| bp-vsir-bweb102-fd JobId 21509: VDDK Transport "nbdssl" selected                    ␣
↪              |
| bp-vsir-bweb102-fd JobId 21509: Dump vmdk 6000C299-88b4-26f3-3145-c18324d30940␣
↪succeeded.            |
| bp-vsir-bweb102-fd JobId 21509: Delete snapshot 000_bp-o9-hap_vsphere.2023-11-15_13.45.
↪06_49 succeeded. |
| bp-vsir-bweb102-fd JobId 21509: BACKUP OK bp-o9-hap (vm-7253)                       ␣
↪              |
| bp-vsir-bweb102-sd JobId 21509: Elapsed time=00:39:32, Transfer rate=4.110 M Bytes/
↪second            |
| bp-vsir-bweb102-sd JobId 21509: Sending spooled attrs to the Director. Despooling 2,
↪219 bytes ...    |
| bp-vsir-bweb102-dir JobId 21509: Bacula Enterprise bp-vsir-bweb102-dir 16.0.7␣
↪(11Jul23):
  Build OS:               x86_64-redhat-linux-gnu-bacula-enterprise redhat (Core)
  JobId:                  21509
  Job:                    000_bp-o9-hap_vsphere.2023-11-15_13.45.06_49
  Backup Level:           Full (upgraded from Incremental)
  Client:                 "bp-vsir-bweb1004-fd" 16.0.7 (11Jul23) x86_64-redhat-linux-gnu-
↪bacula-enterprise,redhat,(Core)
  FileSet:                "000_bp-vsir-bweb1004-fd_bp-o9-hap_vsphere" 2023-10-25 09:59:57
  Pool:                   "GED2D" (From Command input)
  Catalog:                "MyCatalog" (From Client resource)
  Storage:                "LocalGED" (From Pool resource)
  Scheduled time:         15-Nov-2023 13:45:06
  Start time:             15-Nov-2023 13:45:08
  End time:               15-Nov-2023 14:24:42
  Elapsed time:           39 mins 34 secs
  Priority:               10
  FD Files Written:       9
  SD Files Written:       9
  FD Bytes Written:       10,081,812,251 (10.08 GB)
  SD Bytes Written:       9,749,179,049 (9.749 GB)
  Rate:                   4246.8 KB/s
  Software Compression:   3.3% 1.0:1
  Comm Line Compression:  None
  Snapshot/VSS:           no
  Encryption:             no
  Accurate:               yes
  Volume name(s):         ged-3474|ged-3478|ged-3479|ged-3480|ged-3484|ged-3485|ged-
↪3486|ged-3487|ged-3488|ged-3489
  Volume Session Id:      44
  Volume Session Time:    1698051447
  Last Volume Bytes:      meta: 574,097 (574.0 KB) aligned: 562,429,952 (562.4 MB)
```

(continues on next page)

(continued from previous page)

```
  Non-fatal FD errors:    0
  SD Errors:              0
  FD termination status:  OK
  SD termination status:  OK
  Termination:            Backup OK |
+-------------------------------------------------------------------------------------------
↪------------+
+--------+----------------------+---------------------+------+-------+----------+-------
↪---------+-----------+
| jobid  | name                 | starttime           | type | level | jobfiles |␣
↪jobbytes       | jobstatus |
+--------+----------------------+---------------------+------+-------+----------+-------
↪---------+-----------+
| 21,509 | 000_bp-o9-hap_vsphere | 2023-11-15 13:45:08 | B    | F     |        9 | 10,
↪081,812,251 | T         |
+--------+----------------------+---------------------+------+-------+----------+-------
↪---------+-----------+
```

NVRAM and VMX Files

Since version 16.0.12, the NVRAM and VMX files are automatically included as part of the backup process. They are downloaded from the datastore and from the directory where the given Virtual Machine has its disks and other files stored.

During the restore process, the NVRAM and VMX files are automatically restored to the host where the FD and the vSphere Plugin are running. NVRAM file is also uploaded to the destination folder of the Virtual Machine that is being restored.

The NVRAM file contains BIOS information and can help in some restore cases. The VMX file contains configuration information of Virtual Machines and the purpose of having it is to allow the user to read this information if necessary.

See also:

Go to:

- *Restore*
- *Quiescing Guests*
- *VM Instant Recovery*
- *VMware Single Item Restore*
- *List Host Operations*

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

6.2 Restore

This article describes restore with the use of the vSphere Plugin.

Restore to New VMware Guest

If you run a restore of a VM backup using the `where=/` restore option and select all files under the VM's directory, the vSphere Plugin will create a new VM with the same attributes (disks, controller, CPU type, ...) on your ESXi host and restore the disks to this new VM. If you do not specify a new name for the restored VM, then the new VM's name will be the original VM's name with the restore job's jobid appended like: **originalName-123**.

```
* lsmark
/@vsphere/3/vsphere_generation.profile
/@vsphere/3/3.ovf
/@vsphere/3/1.bvmdk
/@vsphere/3/2.bvmdk
```

The SAN advanced transport mode is currently unsupported for restore. The vSphere Plugin will use NBD for VM restores.

The ESXi host and the datastore that will be used to restore your guest VM will be detected automatically. However, you can change the default destination by modifying the plugin restore options in the bconsole menu:

```
Run Restore job
JobName:         RestoreFiles
Bootstrap:       /tmp/regress/working/127.0.0.1-dir.restore.1.bsr
Where:           /tmp/regress/tmp/bacula-restores
...
Plugin Options:  *None*
OK to run? (yes/mod/no): mod
Parameters to modify:
     1: Level
...
    13: Plugin Options
Select parameter to modify (1-13): 13
Automatically selected : vsphere: host=squeeze2
Plugin Restore Options
datastore:              *None*
restore_host:           *None*
new_hostname:           *None*
vsphere_server:         *None*
datastore_allow_overprovisioning: *None*               (yes)
datastore_minimum_space:    *None*
override_vm:            *None*              (no)
power_on:               *None*              (no)
Use above plugin configuration? (yes/mod/no): mod
You have the following choices:
     1: datastore (Datastore to use for restore)
     2: restore_host (ESXi host to use for restore)
     3: new_hostname (Restore guest VM to specified name)
     4: vsphere_server (vSphere server defined in vsphere_global.conf to use for restore)
     5: datastore_allow_overprovisioning (Allow over provisioning when creating a new VM)
     6: datastore_minimum_space (Minimum free space to keep in a Datastore (in MB))
```

(continues on next page)

(continued from previous page)

```
    7: override_vm (Restore to original VM, overriding it's disks (new_hostname value␣
↪will be ignored))
    8: power_on (Power on VM after restoration)
Select parameter to modify (1-6): 3
Please enter a value for new_hostname: test
Plugin Restore Options
datastore:           *None*
restore_host:        *None*
new_hostname:        test
vsphere_server:      *None*
datastore_allow_overprovisioning: *None*              (yes)
datastore_minimum_space: *None*
override_vm:         *None*              (no)
power_on:            *None*              (no)
Use above plugin configuration? (yes/mod/no): yes
```

The restore options may also be modified with using the BWeb restore interface.

Supported restore options are listed and detailed below:

Option	Required	Default	Info	Example
datastore	No		Datastore to use for restore	myDatastore
restore_host	No		ESX host to use for restore	host.mydc.com
new_hostname	No		Name to apply to the new restored host (VM)	myVM-Restored1
vsphere_server	No		vSphere server defined in vsphere_global.conf to use for restore	host3
datastore_allow_overprovisioning	No	yes	Allow over provisioning when creating a new VM	no
datastore_minimum_space	No		Minimum free space to keep in a Datastore (in MB)	430899200
thin_provisioned	No	yes	Create thin provisioned disks	no
override_vm	No	no	Restore to original VM, overriding it's disks (new_hostname value will be ignored)	yes
power_on	No	no	Power on VM after restoration	yes
no_vmdk	No	no	Restore except vmdk contents	yes
force_san	No	no	Force the use of the SAN transfer	yes
no_network	No	no	Do not activate the network	yes
new_network	No		Name or Moref of a new network resource to apply, overriding the original one, during the restore.	network-7583
new_network_device_name	No		Name of the original network resource that will be overridden (by the new_network defined value), during the restore. If not specified, all original network resources will be overridden.	vlan-1

Note: You need to have at least one VM configured on your ESXi server to restore a VM from Bacula automatically. We plan to remove this limitation in a future version.

The vSphere Plugin can check the space available in the datastore during restore. It is possible to disallow *Over Provi-*

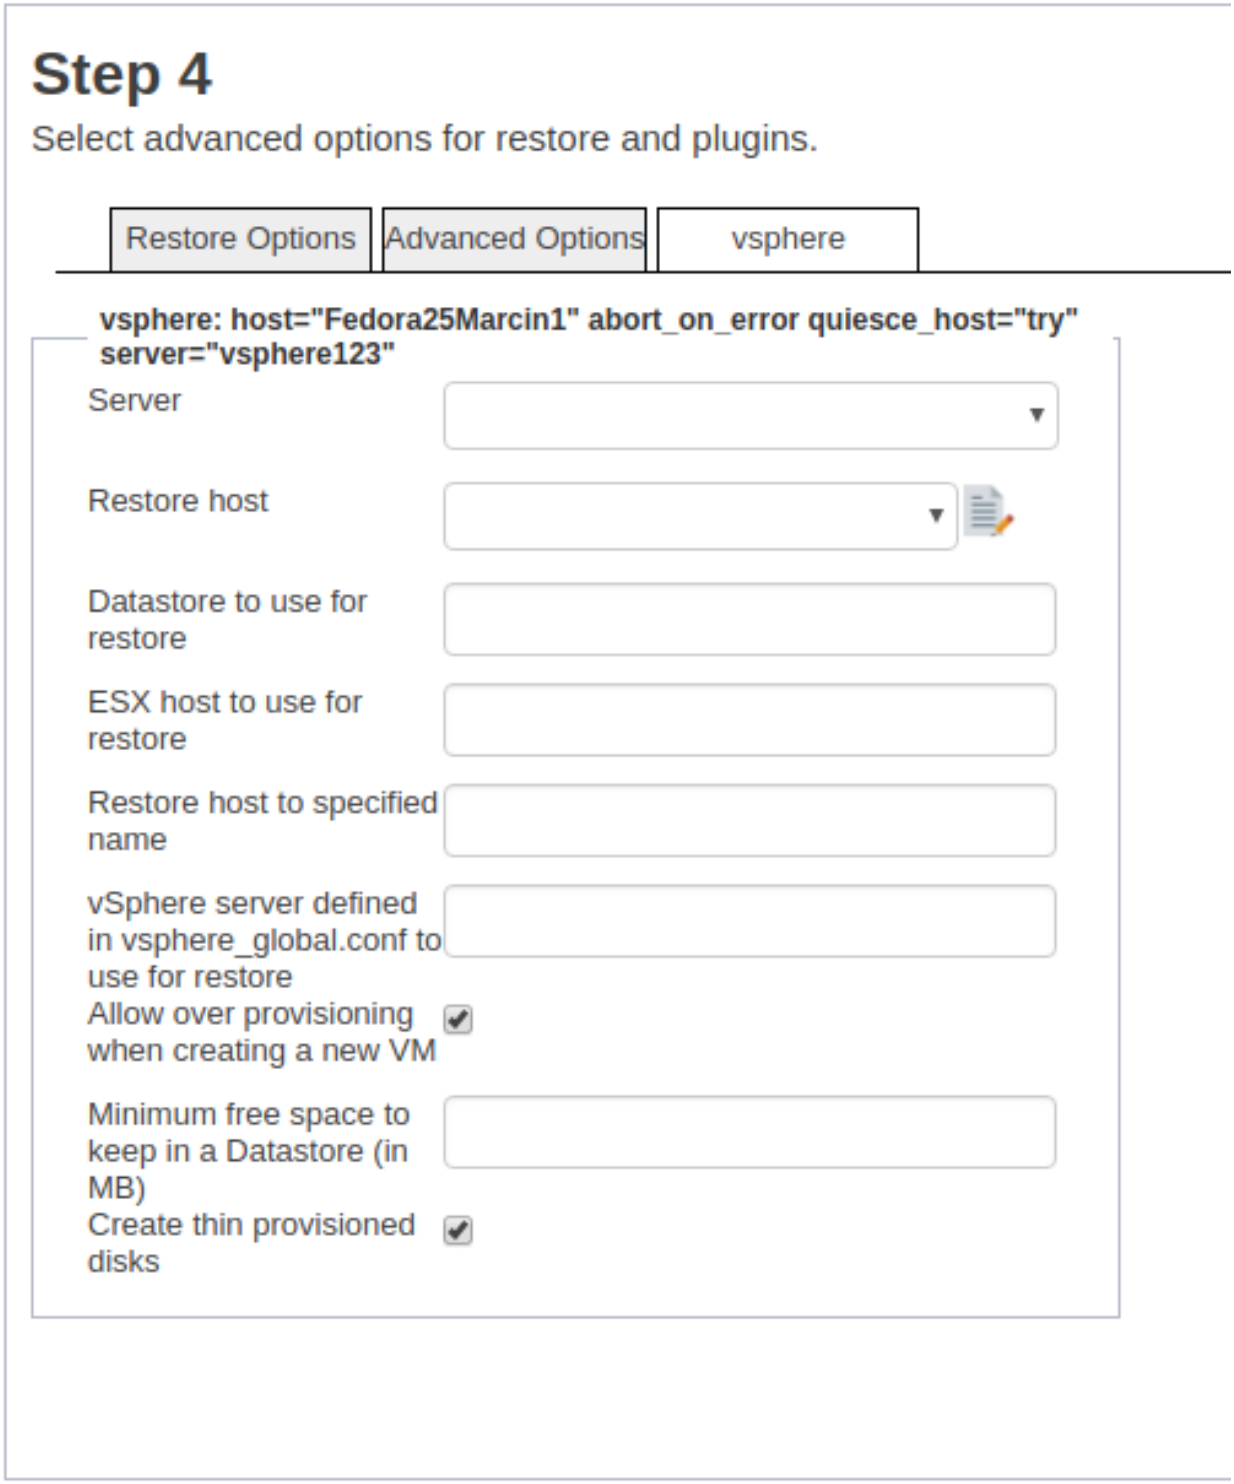

Fig. 4: Choose datastore, ESXi or new VM name at restore time

sioning and reserve a minimum amount of space in the datastore. These two options can be set in the `vsphere_global.conf` file but can be overwritten from the restore menu.

Click *here* to see all available directives for the `vsphere_global.conf` file.

```
[vsphere]
    username = root
    password = xxxx
    server = 192.168.0.68
    url = https://192.168.0.68/sdk
    thumbprint = 34:F5:0F:10:82:59:EF:2D:DB:96:CC:5B:C4:66:33:83:DC:91:AF:01

    datastore_minimum_space = 64MB
    datastore_refresh_interval = 10
    datastore_allow_overprovisioning = false
```

Additional Information

Starting with Bacula Enterprise 12.2, the vSphere Plugin includes the vApp options in the OVF description of the virtual machine.

Starting with Bacula Enterprise 12.3, the vSphere Plugin can power on the virtual machine after a successful restore. Just select the option `power_on` in the bconsole plugin restore options.

See also:

Go back to:

- *Restore to Local Disk with vSphere Plugin*

Go to:

- *Restore to Existing VMware Guest*

Go back to the *vSphere Restore page*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Restore to Existing VMware Guest

Important: This option is available for Bacula 12.3 and above.

If you run a restore of a VM backup using the `where=/` restore option, select all files under the VM's directory and select the `override_vm` Plugin Option, the vSphere Plugin will look for the original guest in the ESXi host and restore those disks that are part of the backup. All other disks remain unchanged. If the Guest is still running, it will be powered off during restore. If the restore was successful, the guest will be powered on again.

```
Run Restore job
JobName:         RestoreFiles
Bootstrap:       /tmp/regress/working/127.0.0.1-dir.restore.1.bsr
Where:           /tmp/regress/tmp/bacula-restores
...
Plugin Options:  *None*
```

(continues on next page)

(continued from previous page)

```
OK to run? (yes/mod/no): mod
Parameters to modify:
     1: Level
...
    13: Plugin Options
Select parameter to modify (1-13): 13
Automatically selected : vsphere: host=squeeze2
Plugin Restore Options
datastore:              *None*
restore_host:           *None*
new_hostname:           *None*
vsphere_server:         *None*
datastore_allow_overprovisioning: *None*               (yes)
datastore_minimum_space:    *None*
override_vm:            *None*              (no)
power_on:               *None*              (no)
Use above plugin configuration? (yes/mod/no): mod
You have the following choices:
     1: datastore (Datastore to use for restore)
     2: restore_host (ESXi host to use for restore)
     3: new_hostname (Restore guest VM to specified name)
     4: vsphere_server (vSphere server defined in vsphere_global.conf to use for restore)
     5: datastore_allow_overprovisioning (Allow over provisioning when creating a new VM)
     6: datastore_minimum_space (Minimum free space to keep in a Datastore (in MB))
     7: override_vm (Restore to original VM, overriding it's disks (new_hostname value␣
↪will be ignored))
     8: power_on (Power on VM after restoration)
Select parameter to modify (1-6): 7
Please enter a value for override_vm: yes
Plugin Restore Options
datastore:           *None*
restore_host:        *None*
new_hostname:        *None*
vsphere_server:      *None*
datastore_allow_overprovisioning: *None*               (yes)
datastore_minimum_space: *None*
override_vm:         yes
power_on:            *None*              (no)
Use above plugin configuration? (yes/mod/no): yes
```

The restore options may also be modified with using the BWeb restore interface.

See also:

Go back to:

- *Restore to Local Disk with vSphere Plugin*
- *Restore to Existing VMware Guest*

Go back to the *vSphere Restore page*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Restore to Local Disk with vSphere Plugin

Bacula Enterprise allows restoring any file (`bvmdk`, `ovf`, ...) to your File Daemon's local disks. Then, you may mount the image locally using the VMware `vmware-mount` tool or `qemu-nbd` and perform file level restores.

By using `where=/path/to/dir` in the restore options, the Plugin will automatically restore selected files to this location on your File Daemon's local disk.

```
% qemu-img convert -O vmdk /tmp/0.bvmdk /tmp/0.vmdk
```

It is also possible to copy the raw image to any device or to mount it and restore files directly.

```
 # modprobe nbd max_part=16
 # qemu-nbd -c /dev/nbd0 /tmp/0.bvmdk
 # partprobe /dev/nbd0
 # fdisk -l /dev/nbd0

Disk /dev/nbd0: 2147 MB, 2147483648 bytes
255 heads, 63 sectors/track, 261 cylinders, total 4194304 sectors
Units = sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x000a7154

     Device Boot      Start         End      Blocks   Id  System
/dev/nbd0p1   *          63     1686824      843381   83  Linux
/dev/nbd0p2         1686825     4192964     1253070    5  Extended
/dev/nbd0p5         1686888     1959929      136521   82  Linux swap
/dev/nbd0p6         1959993     4192964     1116486   83  Linux

 # mount /dev/nbd0p1 /mnt/image
 # ls /mnt/image
bin   cdrom  etc      initrd.img  lost+found  mnt    proc   sbin
tmp   var    boot     dev         home        lib    media  opt
sys   usr    vmlinuz  srv         selinux     root

 # umount /mnt/image
 # qemu-nbd -d /dev/nbd0
```

See also:

Go to:

- *Restore to New VMware Guest*
- *Restore to Existing VMware Guest*

Go back to the *vSphere Restore page*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Backup*

Go to:

- *Quiescing Guests*
- *VM Instant Recovery*
- *VMware Single Item Restore*
- *List Host Operations*

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

6.3 Quiescing Guests

To properly quiesce the guest machines, VMware Tools **must** be installed and up to date in your Linux/Windows virtual machines.

The `quiesce_host=try/yes/no` Plugin command line option allows control of how the vSphere Plugin should try to quiesce guest VMs before a snapshot. The default value is `try`. In this mode, the plugin will try to quiesce the guest when creating the snapshot. If the quiesce operation fails, the plugin will make a second attempt to create the snapshot without quiescing the guest. The first attempt will be reported in the job log as an error and the job termination status will terminate "with warnings".

```
FileSet {
  Name = guestvm
  Include {
    # Will try to quiesce the VM and skip the VM
    # from the backup if pre/post scripts are
    # returning an error
    Plugin = "vsphere: host=guestvm quiesce_host=yes"
  }
}
```

More information about the exact error message can be found in the vSphere console log.

```
Warning message from ESXi: the guest OS has reported an error during
quiescing. Error code was: 2 the error message was: custom quiesce script
failed.
```

Or

```
An error occurred while saving the snapshot: Failed to quiesce the virtual machine.
```

Linux

By creating a special script located in `/usr/sbin/pre-freeze-script`, you will be able to quiesce your system automatically when vSphere creates a Snapshot.

The vSphere Plugin will also try to execute `/usr/sbin/post-thaw-script` script if present on the guest VM.

Windows VSS notes

The plugin enhances Windows protection by performing VSS-based snapshots to quiesce VSS-enabled applications before backup.

VSS Pre-freeze and Post-thaw Scripts

VMware Tools first looks in `C:/Program Files/VMware/VMware Tools/backupScripts.d` for scripts in alphabetic order, calling them all with freeze argument, and afterward in reverse alphabetic order, calling them with `thaw` argument (or `freezeFail` if quiescing failed).

See also:

Go back to:

- *Backup*
- *Restore*

Go to:

- *VM Instant Recovery*
- *VMware Single Item Restore*
- *List Host Operations*

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

6.4 VM Instant Recovery

Starting With Bacula Enterprise 12.6, it is now possible to recover a vSphere Virtual Machine in a matter of minutes by running it directly from a Bacula Volume.

Any changes made to the VM disks are virtual and temporary. This means that the disks remain in a read-only state. The users may write to the VM disks without the fear of corrupting their backups. Once the Virtual Machine is started, it is then possible via VMotion to migrate the temporary Virtual Machine to a production datastore.

The feature is available inside the `mount-vm` script and follows this workflow:

1. User chooses a VM backup.
2. Script mounts the VM disks locally.
3. User chooses the ESXi host that will own the VM.
4. Script creates a temporary NFS Datastore locally.
5. Script creates and powers on the VM in this temporary NFS Datastore.
6. User chooses to keep the VM permanently or to discard it.

To learn more about this feature, see the sections below:

VM Instant Recovery Installation

Installation: Bacula Storage Daemon

In addition to installing the **bacula-enterprise-single-item-restore** package as described in *Single Item Restore Installation*, some extra actions are needed to enable the Instant Recovery feature:

1. Install and configure the vSphere Plugin

Refer to Bacula's *vSphere plugin documentation* for instructions on its installation and configuration.

2. Install NFS service

To install and configure the NFS server service automatically, run the `install-single-item-restore` script:

```
$ /opt/bacula/scripts/install-single-item-restore.sh install_ir
```

Installation: vSphere side

1. Test NFS connection on an ESXi host

Since Instant Recovery creates a temporary **NFS Datastore** on the machine containing the Bacula volumes, we need to make sure that the **ESXi hosts** can reach it through the NFS ports:

```
$ ssh <ESXi-host>
$ ping <nfs-datastore>
$ nc -z <nfs-datastore> 2049
```

For more information, refer to https://kb.vmware.com/s/article/1003967.

2. Set up vMotion

Migration with vMotion requires a vMotion network interface on each ESXi host where you plan to migrate VMs. The vMotion interface can be configured from either a vSphere client or a vSphere Web Client. The steps are:

- Navigate to the desired ESXi host
- Navigate to the host network settings
- Under network settings, click to **add networking**
- Select **VMkernel Network Adapter** as connection type
- Select either an existing switch or a new one
- Under port properties, set it to **allow vMotion traffic**

For more information and best practices please refer to https://kb.vmware.com/s/article/2054994.

See also:

Go to:

- *Recovery Scenarios*
- *Limitations*

Go back to the *VM Instant Recovery*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Recovery Scenarios

This article aims at presenting recovery scenarios of VM Instant Recovery.

Temporary Recovery for Testing Purposes

In this scenario, we recover the VM called `yVM` to the ESXi host located at 192.168.0.26:

```
[root@localhost bin]# sudo -u bacula /opt/bacula/bin/mount-vm
Automatically Selected Catalog: BaculaCatalog
Automatically Selected Client: localhost-fd

Job list:
1: LinuxEtc.2020-08-31_07.08.04_04
2: vsphere_hbck-centos7.2020-08-27_06.39.21_03
3: vsphere_hbck-centos7.2020-08-31_08.01.58_03
4: vsphere_yVM.2020-09-22_08.03.07_08
Select a Job: 4
Selected vsphere_yVM.2020-09-22_08.03.07_08
Automatically Selected Virtual Machine: yVM (14)
Automatically Selected Disks: 0

Action list:
1: Mount guest filesystem locally
2: Export guest filesystem through SMB
3: Export guest virtual machine to vSphere instance (Instant Recovery)
4: Cleanup
Select an Action: 3
Selected Export guest virtual machine to vSphere instance (Instant Recovery)
I: Instant Recovery Mode
I: NFSv4 detected.
Select where you want to mount the Virtual Machine

ESXi Host list:
1: [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

2: [192.168.0.8]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Select an ESXi Host: 1
Selected [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

[sudo] password for bacula:
Creating NFS Datastore...
OK: Registered this machine as a NFS Datastore (name=bir-14-6959, host=192.168.0.26)
Creating Virtual Machine...
OK: The Virtual Machine is now available (name=yVM-6959, host:192.168.0.26,␣
↪datastore:bir-14-6959)

Action list:
```

(continues on next page)

(continued from previous page)

```
1: Migrate Virtual Machine to a permanent Datastore
2: Cleanup
Select an Action: 2
Selected Cleanup
end
```

See also:

Go to:

- *Recovery and Migration in the Same ESXi Host*
- *Recovery and Migration to Another ESXi Host*
- *Instant Recovery with Network Card Disconnected*
- *Manually Cleaning an Instant Recovery Session*

Go back to the *Recovery Scenarios article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Recovery and Migration in the Same ESXi Host

In this scenario, we recover the VM called `yVM` to the ESXi host located at 192.168.0.26 and migrate it permanently to the datastore "datastore 1 (1)".

```
[root@localhost bin]# sudo -u bacula /opt/bacula/bin/mount-vm
Automatically Selected Catalog: BaculaCatalog
Automatically Selected Client: localhost-fd

Job list:
1: LinuxEtc.2020-08-31_07.08.04_04
2: vsphere_hbck-centos7.2020-08-27_06.39.21_03
3: vsphere_hbck-centos7.2020-08-31_08.01.58_03
4: vsphere_yVM.2020-09-22_08.03.07_08
Select a Job: 4
Selected vsphere_yVM.2020-09-22_08.03.07_08
Automatically Selected Virtual Machine: yVM (14)
Automatically Selected Disks: 0

Action list:
1: Mount guest filesystem locally
2: Export guest filesystem through SMB
3: Export guest virtual machine to vSphere instance (Instant Recovery)
4: Cleanup
Select an Action: 3
Selected Export guest virtual machine to vSphere instance (Instant Recovery)
I: Instant Recovery Mode
I: NFSv4 detected.
Select where you want to mount the Virtual Machine

ESXi Host list:
```

(continues on next page)

(continued from previous page)

```
1: [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

2: [192.168.0.8]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Select an ESXi Host: 1
Selected [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

[sudo] password for bacula:
Creating NFS Datastore...
OK: Registered this machine as a NFS Datastore (name=bir-14-7454, host=192.168.0.26)
Creating Virtual Machine...
OK: The Virtual Machine is now available (name=yVM-7454, host:192.168.0.26,␣
↪datastore:bir-14-7454)

Action list:
1: Migrate Virtual Machine to a permanent Datastore
2: Cleanup
Select an Action: 1
Selected Migrate Virtual Machine to a permanent Datastore
Select which host you want to migrate the Virtual Machine to:

ESXi Host list:
1: [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

2: [192.168.0.8]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Select an ESXi Host: 1
Selected [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Select which datastore you want to migrate the Virtual Machine to:
Automatically Selected Datastore: datastore1 (1)
Migrating Virtual Machine to datastore1 (1) (192.168.0.26). This may take some time...
OK: The Virtual Machine was migrated.
I: Press enter to finish and cleanup the session

I: End of session
```

See also:

Go back to:

- *Temporary Recovery for Testing Purposes*

Go to:

- *Recovery and Migration to Another ESXi Host*
- *Instant Recovery with Network Card Disconnected*
- *Manually Cleaning an Instant Recovery Session*

Go back to the *Recovery Scenarios article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Recovery and Migration to Another ESXi Host

In this scenario, we recover the VM called `yVM` to the ESXi host located at 192.168.0.26 and migrate it permanently to a datastore in another ESXi host (192.168.0.8).

```
[root@localhost bin]# sudo -u bacula /opt/bacula/bin/mount-vm
Automatically Selected Catalog: BaculaCatalog
Automatically Selected Client: localhost-fd

Job list:
1: LinuxEtc.2020-08-31_07.08.04_04
2: vsphere_hbck-centos7.2020-08-27_06.39.21_03
3: vsphere_hbck-centos7.2020-08-31_08.01.58_03
4: vsphere_yVM.2020-09-22_08.03.07_08
Select a Job: 4
Selected vsphere_yVM.2020-09-22_08.03.07_08
Automatically Selected Virtual Machine: yVM (14)
Automatically Selected Disks: 0

Action list:
1: Mount guest filesystem locally
2: Export guest filesystem through SMB
3: Export guest virtual machine to vSphere instance (Instant Recovery)
4: Cleanup
Select an Action: 3
Selected Export guest virtual machine to vSphere instance (Instant Recovery)
I: Instant Recovery Mode
I: NFSv4 detected.
Select where you want to mount the Virtual Machine

ESXi Host list:
1: [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

2: [192.168.0.8]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Select an ESXi Host: 2
Selected [192.168.0.8]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Creating NFS Datastore...
OK: Registered this machine as a NFS Datastore (name=bir-14-8028, host=192.168.0.8)
Creating Virtual Machine...
OK: The Virtual Machine is now available (name=yVM-8028, host:192.168.0.8, datastore:bir-
↪14-8028)
```

(continues on next page)

(continued from previous page)

```
Action list:
1: Migrate Virtual Machine to a permanent Datastore
2: Cleanup
Select an Action: 1
Selected Migrate Virtual Machine to a permanent Datastore
Select which host you want to migrate the Virtual Machine to:

ESXi Host list:
1: [192.168.0.26]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

2: [192.168.0.8]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Select an ESXi Host: 2
Selected [192.168.0.8]
    (administrator@vsphere.local, https://192.168.0.15/sdk)

Select which datastore you want to migrate the Virtual Machine to:
Automatically Selected Datastore: datastore1
Migrating Virtual Machine to datastore1 (192.168.0.8). This may take some time...
OK: The Virtual Machine was migrated.
I: Press enter to finish and cleanup the session

I: End of session
```

See also:

Go back to:

- *Temporary Recovery for Testing Purposes*
- *Recovery and Migration in the Same ESXi Host*

Go to:

- *Instant Recovery with Network Card Disconnected*
- *Manually Cleaning an Instant Recovery Session*

Go back to the *Recovery Scenarios article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Instant Recovery with Network Card Disconnected

Starting with version 14.0.1, it is possible to create the temporary VM with network cards not connected. To disconnect network cards at boot, use the `-N` option in the `mount-vm` command line.

```
[root@localhost bin]# sudo -u bacula /opt/bacula/bin/mount-vm -N
Automatically Selected Catalog: BaculaCatalog
Automatically Selected Client: localhost-fd

Job list:
```

(continues on next page)

(continued from previous page)

```
1: LinuxEtc.2020-08-31_07.08.04_04
...
```

See also:

Go back to:

- *Temporary Recovery for Testing Purposes*
- *Recovery and Migration in the Same ESXi Host*
- *Recovery and Migration to Another ESXi Host*

Go to:

- *Manually Cleaning an Instant Recovery Session*

Go back to the *Recovery Scenarios article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Manually Cleaning an Instant Recovery Session

It is possible to manually clean an Instant Recovery session by purging its associated datastore:

```
bacula@storage# /opt/bacula/bin/vsphere-ctl purge_ds bir-41-4018 --server vcenter_70
I: Successfully purged Datastore bir-41-4018
```

See also:

Go back to:

- *Temporary Recovery for Testing Purposes*
- *Recovery and Migration in the Same ESXi Host*
- *Recovery and Migration to Another ESXi Host*
- *Instant Recovery with Network Card Disconnected*

Go back to the *Recovery Scenarios article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *VM Instant Recovery Installation*

Go to:

- *Limitations*

Go back to the *VM Instant Recovery*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

VM Instant Recovery Limitations

- The VMware Single Item Restore feature uses the Bacula BVFS interface to list files and directories. The Bacula BVFS interface is known to have some performance issues with MySQL catalog backend due to internal MySQL limitations with indexes on TEXT colums. For VMware and Exchange Single Item Restore there should not be too much impact on performances (the backup structure is usually quite small) but we advise using the PostgreSQL backend for the best experience.
- The VMware Single Item Restore performance may vary depending on various factors. For example, Bacula will have to read more data if the Volume was created with a large number of concurrent jobs.
- The Storage Daemon where the VMware Single Item Restore is installed should be have a CPU with the VT-x/EPT extensions. If these extensions are not available, the performance will be degraded. (From 20s to 10mins in our lab).
- The VMware Single Item Restore is compatible with *file based* devices (cloud, dedup, aligned, file, etc..). Tape devices are not supported.
- To perform VM Instant Recovery from a Copy/Migration job make sure destination Pool has `Maximum Volume Jobs` set to 1. Note that when you use `MaximumVolumeJobs = 1` in the Pool resource, you must use `MaximumConcurrentJobs = 1` in the Device resource(s).
- The Instant Recovery hot migration only works with the VMware vMotion technology.
- All volumes needed for the VMware Single Item Restore must reside on the single Storage Daemon instance where the SIR session is started. A storage group policy can conflict with this requirement.
- To avoid heavy network traffic and prevent jobs from failing, do not run the vSphere plugin along the FD running on the SD with a storage group.
- Client side PKI Encryption is currently not compatible with vSphere IR features. Use the Volume Encryption feature if needed.

See also:

Go back to:

- *VM Instant Recovery Installation*
- *Recovery Scenarios*

Go back to the *VM Instant Recovery*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Backup*
- *Restore*
- *Quiescing Guests*

Go to:

- *VMware Single Item Restore*
- *List Host Operations*

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

6.5 VMware Single Item Restore

This section presents how to use the VMware Single Item Restore feature with Bacula Enterprise and the vSphere Plugin.

Single Item Restore Features

The **Bacula Enterprise** VMware Single Item Restore provides the following main features:

- Console interface
- BWeb Management Suite interface
- Support for Full/Differential/Incremental jobs
- Support for Windows NTFS
- Support for Linux (ext3, ext4, btrfs, lvm, xfs)
- Support for ESX 5.x, 6.x, 7.0 and 8.0.

See also:

Go to:

- *Single Item Restore Scope*
- *Single Item Restore Installation*
- *Single Item Restore Configuration*
- *Single Item Restore: Restore Scenarios*
- *Single Item Restore Limitations*
- *Single Item Restore Troubleshooting*

Go back to the *main VMware Single Item Restore article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Single Item Restore Scope

This document will present solutions for **Bacula Enterprise** 8.4 and later, which are not applicable to prior versions. The VMware Single Item Restore has been tested and is supported on RedHat Linux, Oracle Linux and Rocky Linux 7.x, 8.x, 9.x. SELinux is currently not supported. The vSphere Plugin might not be supported on all platforms where VMware Single Item Restore is supported.

> **Warning:** Redhat decided to stop supporting Windows NTFS devices starting with RHEL version 7 and they have removed the "ntfs-ng" package from the official Redhat repository. The "ntfs-ng" package is a required dependency of the "libguestfs" package and will need to be installed separately from a repository such as EPEL (see *Single Item Restore Installation*)

See also:

Go back to:

- *Single Item Restore Features*

Go to:

- *Single Item Restore Installation*
- *Single Item Restore Configuration*
- *Single Item Restore: Restore Scenarios*
- *Single Item Restore Limitations*
- *Single Item Restore Troubleshooting*

Go back to the *main VMware Single Item Restore article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Single Item Restore Installation

This article describes how to install Single Item Restore for the vSphere Plugin.

SIR Installation with BIM

The SIR features are installed along the SD. In consequence, to install SIR for the vSphere Plugin with BIM, run BIM with the *-t SD* option and choose to install SIR. If the SD is already installed on you system, Bacula configuration and SD packages will not be modified.

See also:

Go to:

- *SIR Installation with Package Manager*
- *BWeb Integration with Single Item Recovery*

Go back to the *main SIR Installation page*.

Go back to the *main Single Item Restore page*.

Go back to the *main vSphere Plugin Operations page*.

SIR Installation with Package Manager

Packages for the VMware Single Item Restore plugin are available for supported platforms. Please contact Bacula Systems to get them.

Download the plugin package to your **Storage Daemon** server and then install using the package manager like so:

```
rpm -ivh bacula-enterprise-single-item-restore*.rpm
```

The package manager will ensure that your **Bacula Enterprise** version is compatible with the VMware Single Item Restore Plugin and will install dependencies. On Redhat, it will be needed to install `perl-JSON` package from **rpmforge** and the `libguestfs-winsupport` package.

Note: On Redhat 8.x or 9.x, it is necessary to install a custom version of the libguestfs packages from our repository to support NTFS devices. Those should not be updated with a newer version from official repositories. The YUM package manager has plugins to prevent package updates, try **yum-plugin-versionlock** or **yum-plugin-priorities**. Additionally, the `ntfs-3g` package from the EPEL repository is needed for NTFS support. To install the EPEL respository, please

follow the official instructions on the EPEL website to install the "epel-release" package here: https://fedoraproject.org/wiki/EPEL

Note: The *DAG* repo below must be set up in order to install *libguestfs* needed for SIR.

Note: On Redhat 8.X and 9.x, you must have the the AppStream repository enabled to install the perl-File-Copy. The perl-File-Copy module is a dependency required by the bacula-enterprise-single-item-restore package.

Since Bacula Enterprise 16.0.13.

```
# cat /etc/yum.repos.d/dag.repo
[dag]
name = Red Hat Enterprise  - RPMFORGE
baseurl = https://www.baculasystems.com/dl/DAG/rhel6-64
enabled = 1
protect = 0
gpgcheck = 0

# cat /etc/yum.repos.d/baculasystems.repo
[single_file_restore_vmware]
name = Red Hat Enterprise  - RPMFORGE
baseurl = https://www.baculasystems.com/dl/<xxx>/rhel6-64
enabled = 1
protect = 0
gpgcheck = 0


Note: This last repository is required on RHEL7:

[Bacula-Enterprise-DAG-Guestfish]
name = Bacula Enterprise - DAG for Guestfish
baseurl = https://www.baculasystems.com/dl/DAG/rhel7-64/guestfish/
enabled = 1
protect = 0
gpgcheck = 0
```

```
yum install bacula-enterprise-single-item-restore perl-JSON
```

If BWeb Management Suite is used:

```
service bweb restart
```


Samba SMB Shares

The **Bacula Enterprise** VMware Single Item Restore plugin can automatically set up Samba SMB shares from the console program or the BWeb Management Suite.

To enable Samba SMB network shares, installing and configuring the "samba" package is mandatory. To configure the `/etc/samba/smb.conf` file correctly, you need to run `install-single-item-restore.sh` script.

```
root@storage# /opt/bacula/scripts/install-single-item-restore.sh install
Do you want to initialise Samba smb.conf [yes/No]: yes
Choose a Workgroup [BACULA]:

root@storage# cat /etc/samba/smb.conf
[global]
workgroup = BACULA
include = /etc/samba/conf.d/all
```

At this point, it is possible to modify `/etc/samba/smb.conf` to add your own configuration directives.

Network share descriptions will be stored in the directory `/etc/samba/conf.d`. It is possible to create and customize the template used by Bacula to generate configuration files.

```
root@storage# cat /etc/samba/conf.d/custom.tpl
[__share__]
   path = __path__
   follow symlinks = yes
   wide links = yes
   writable = yes
```

See also:

Go back to:

- *SIR Installation with BIM*

Go to:

- *BWeb Integration with Single Item Recovery*

Go back to the *main SIR Installation page*.

Go back to the *main Single Item Restore page*.

Go back to the *main vSphere Plugin Operations page*.

BWeb Integration with Single Item Recovery

Installation

To use the BWeb Management Suite graphical GUI with the VMware Single File Restore option it is currently necessary to install and configure BWeb Management Suite on the Storage Daemon where vSphere jobs are stored. If the Director is not on the same machine than the Storage Daemon, remember that the administrator needs to connect to the right BWeb Management Suite instance to use specific VMware Single Item Restore screens.

After the installation of the Single Item Restore package, the BWeb service "bweb" must be restarted to take in account the `bacula` unix user group modification.


```
service bweb restart
```

HTTP Server Extra Configuration

To let the end user access the virtual machine files, it is necessary to set up the `lighttpd` daemon correctly. In this case, we advise using both SSL and user authentication. Example `/opt/bweb/etc/httpd.conf`:

```
#####################################################################
# To enable SSL
# openssl req -x509 -nodes -days 365 -newkey rsa:1024 -keyout server.pem -out server.pem
# chown bacula:bacula server.pem
# chmod 400 server.pem

ssl.engine = "enable"
ssl.pemfile = "/opt/bweb/etc/server.pem"


#####################################################################
# To enable Auth login http://redmine.lighttpd.net/projects/1/wiki/Docs_ModAuth
# htpasswd -c /opt/bweb/etc/htpasswd.bweb

auth.backend = "htpasswd"
auth.backend.htpasswd.userfile =  "/opt/bweb/etc/htpasswd.bweb"
auth.require = ( "/" =>
        (
        "method" => "basic",
        "realm" => "Password protected area",
        "require" => "valid-user"
        )
   )
```

See also:

Go back to:

- *SIR Installation with BIM*
- *SIR Installation with Package Manager*

Go back to the *main SIR Installation page*.

Go back to the *main Single Item Restore page*.

Go back to the *main vSphere Plugin Operations page*.

See also:

Go back to:

- *Single Item Restore Features*
- *Single Item Restore Scope*

Go to:

- *Single Item Restore Configuration*
- *Single Item Restore: Restore Scenarios*

- *Single Item Restore Limitations*
- *Single Item Restore Troubleshooting*

Go back to the *main VMware Single Item Restore article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Single Item Restore Configuration

Storage Daemon Configuration

On the **Storage Daemon** host server, the `bconsole` program should be configured properly to let the "bacula" user connect to the Director with `/opt/bacula/etc/bconsole.conf`.

```
bacula@storage# /opt/bacula/bin/bconsole
Connecting to Director mydir-dir:9101
1000 OK: 10002 mydir-dir Version: 8.4.0 (11 August 2015)
Enter a period to cancel a command.
* version
mydir-dir Version: 8.4.0 (11 August 2015) x86_64-redhat-linux-gnu
* quit
```

The package contains a script to test the connection with the Director and to test if the system can mount the *Bacula Virtual File System* properly.

```
bacula@storage#  /opt/bacula/scripts/install-single-item-restore.sh check
I: Try to restart the script with sudo...
I: Found catalog MyCatalog
I: bacula-fused started on /tmp/bee-bfuse.XXXXX
I: MyCatalog found
I: 10 Client(s) found
I: /tmp/bee-bfuse.XXXXX unmounted
I: bacula-fused (rw) started on /tmp/bee-bfuse.XXXXX
I: MyCatalog found
I: 10 Client(s) found
I: /tmp/bee-bfuse.XXXXX unmounted
OK: All tests are good.
```

The *Bacula Virtual File System* is not designed to be used by end users to browse or restore files directly. If you try to access and browse the mount point, you may not see any files or files may have strange permissions, ownerships and sizes and will inaccessible even to the root user.

See also:

Go back to:

- *Single Item Restore Features*
- *Single Item Restore Scope*
- *Single Item Restore Installation*

Go to:

- *Single Item Restore: Restore Scenarios*

- *Single Item Restore Limitations*
- *Single Item Restore Troubleshooting*

Go back to the *main VMware Single Item Restore article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Single Item Restore: Restore Scenarios

"bacula" Account on Redhat

All commands in this document use the "bacula" unix account to run.

On Redhat, the Unix "bacula" account is locked by default. It means that it's not possible by default to execute a command such as `su - bacula`.

It is possible to unlock the "bacula" account, or to use `sudo -u bacula` to execute commands. For example:

```
bacula@storage# /opt/bacula/bin/bconsole
```

It can be run from the root account using the following command:

```
root@storage# sudo -u bacula /opt/bacula/bin/bconsole
```

It is also possible to start a shell session using:

```
root@storage# sudo -u bacula /bin/bash
```

or unlock the "bacula" unix account and use `su -` with a command such as:

```
root@storage# chsh -s /bin/bash bacula
root@storage# su - bacula
bacula@storage# whoami
bacula
```

Fuse FileSystem

If a restore session is not properly cleaned up, the `mount` command may show some directories mounted with the Bacula Fuse FileSystem.

```
baculafs on /opt/bacula/working/cat-ro type fuse.baculafs (ro,user=bacula)
backend0 on /opt/bacula/working/26-0 type fuse.backend0 (ro,user=bacula)
/dev/fuse on /opt/bacula/working/26 type fuse (rw,nosuid,nodev,user=bacula)
```

It is possible to unmount directories with the `fusermount -u` command.

```
bacula@storage# fusermount -z -u /opt/bacula/working/26
bacula@storage# fusermount -z -u /opt/bacula/working/26-0
bacula@storage# fusermount -z -u /opt/bacula/working/cat-ro
```


Cache Directory

To speed up future VMware Single Item restore sessions, some files that are generated during a restore session are kept in a cache directory.

```
bacula@storage# ls /opt/bacula/working/mount-cache
1-5-0.bmp  1-5-2.bmp    MyCatalog-2.idx  MyCatalog-5.idx  MyCatalog-8.idx
1-5-1.bmp  1-5.profile  MyCatalog-4.idx  MyCatalog-6.idx  MyCatalog-9.idx
```

It is possible to remove files in the cache after some time - they will be re-generated if needed.

With Text Console Interface

The VMware Single Item Restore plugin provides a simple console program that provides access to files inside VMs.

```
bacula@storage# /opt/bacula/bin/mount-vmware
Automatically Selected Catalog: MyCatalog

Client list:
1: 127.0.0.1-fd
2: win2008-fd
3: rhel7-fd
Select a Client: 1
Selected Client: 127.0.0.1-fd

Job list:
1: NightlySave.2015-09-01_10.49.18_39
2: pluginTest.2015-09-01_10.40.20_04
3: pluginTest.2015-09-01_10.46.19_08
Select a Job: 2
Selected pluginTest.2015-09-01_10.40.20_04

Virtual Machine:
1: squeeze2 (5)
2: win2008 (6)
3: rhel7 (7)
Select a Virtual Machine: 1
Selected squeeze2 (5)

Actions list:
1: Mount guest filesystem locally
2: Export guest filesystem through SMB
3: Cleanup
Select a Actions: 1
Selected Mount guest filesystem locally

I: Files are available under /opt/bacula/working/vmware/5
I: Press enter to finish and cleanup the session
```

In this step, the virtual machine filesystem is mounted locally (in the example above, files are available under `/opt/bacula/working/vmware/5`. It is possible to browse directories and copy files (with `cp`, `scp`, `ftp`) as with a standard filesystem from another terminal session with the Unix "root" and "bacula" accounts. If you need to use another Unix account to operate on files, use the "-o allow_other" option when starting the `mount-vmware` script.


```
bacula@storage# ls /opt/bacula/working/vmware/5
bin   dev  home        lib         media  opt   root  selinux  sys  usr  vmlinuz
boot  etc  initrd.img  lost+found  mnt    proc  sbin  srv      tmp  var
```

To clean up the session, just press "Enter" in the terminal session where the `mount-vmware` script was started.

It is possible to limit the Job list with the following command line options:

- `-s=<days>` Limit the job list to the last *days*
- `-l=<number>` Limit the job list to the last *number* entries
- `-f=<filter>` Specify an advanced filter based on the Job name, the FileSet name or the JobId

```
# Limit the job output to the last 100 jobs
bacula@storage# /opt/bacula/bin/mount-vmware -l 100

# Limit the job output to the last 30 days
bacula@storage# /opt/bacula/bin/mount-vmware -s 30

# Limit the job output to jobs that start with ``MyVMware''
bacula@storage# /opt/bacula/bin/mount-vmware -f 'jobname=MyVMware*'

# BAD USAGE for the filter option, it will search for a job named ``MyVMware''
bacula@storage# /opt/bacula/bin/mount-vmware -f 'jobname=MyVMware'

# Limit the job output to jobs that start with ``MyVMware''
# and that use the FileSet Test1
bacula@storage# /opt/bacula/bin/mount-vmware -f 'jobname=MyVMware* fileset=Test1'

# Limit the job to the jobid XX
bacula@storage# /opt/bacula/bin/mount-vmware -f jobid=XX
```

On some cases, the device detection doesn't work properly. It is possible to use the `-m` option to mount recognized disks in a simple way. The option is automatically set when only one disk is selected during the restore.

```
bacula@storage# /opt/bacula/bin/mount-vmware -m
```

See also:

Go to the *With BWeb Management Suite Interface article*.

Go back to the *Restore Scenarios article*.

Go back to the *Single Item Restore chapter*.

Go back to the *main vSphere Plugin page*.

With BWeb Management Suite Interface

The VMware Single Item Restore option in BWeb Management Suite is a wizard that provides easy restoration of files from a VMware guest. The integration of BWeb within the Single Item Restore is necessary following the below steps.

The first step is to select the Client where the vSphere backup job was done.

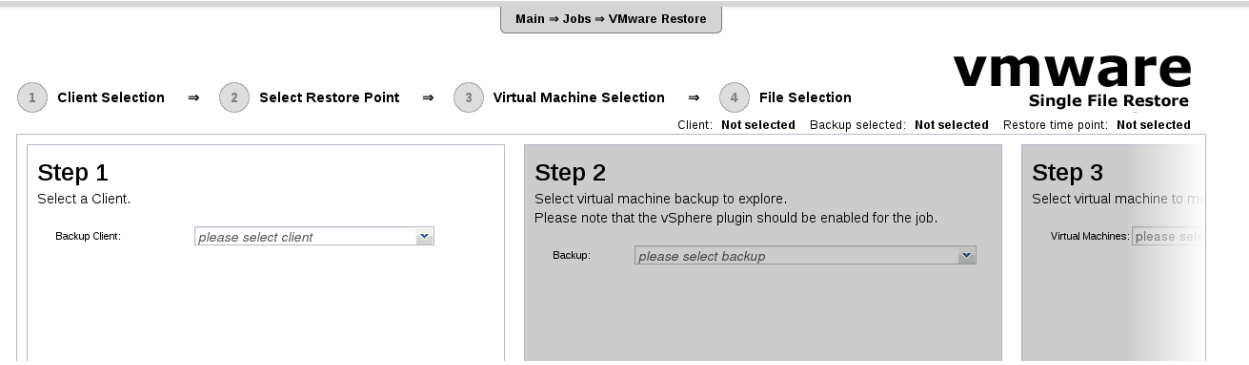

Fig. 5: Client Selection

Once the Client is selected, the administrator needs to select the Job (a Restore Point) to restore.

If the selected Job is a valid vSphere job, the third step will display a list of all virtual machines included in the FileSet.

At this point, Bacula needs to build a virtual image of the selected virtual machine. A couple of small files need to be restored from each Job that makes up the selected *Restore Point*. Once done, Bacula needs to mount the disk of the selected virtual machine on the system. These steps are usually quite fast, but the time depends a lot on the configuration used. Indexes are created and kept during this phase to speed up any further restore requests.

Available with Bacula Enterprise 8.6

To create the index during the backup phase, the FileSet plugin option `index` can be used.

```
Plugin = "vsphere: host=myhost index"
```

Once mounted, the selected virtual machine files will be displayed in a file browser where it is possible to select files or directories to restore (figure below). The administrator can then choose to generate a ZIP or a TAR archive. The archive will be generated automatically and will be stored in `/opt/bacula/working`. A secure HTTP download link will be generated, and the administrator can provide this link to the end user.

If BWeb Management Suite is configured to use HTTP Authentication, it is necessary to configure `lighttpd` properly to allow “anonymous” users to download their files. (See *HTTP Server Extra Configuration*)

For each selection, the administrator can choose how to retrieve the files directly, compressed as a tar file or a zip file.

Once the restore has taken place it is important to terminate the

See also:

Go to the *With Text Console Interface article*.

Go back to the *Restore Scenarios article*.

Go back to the *Single Item Restore chapter*.

Go back to the *main vSphere Plugin page*.

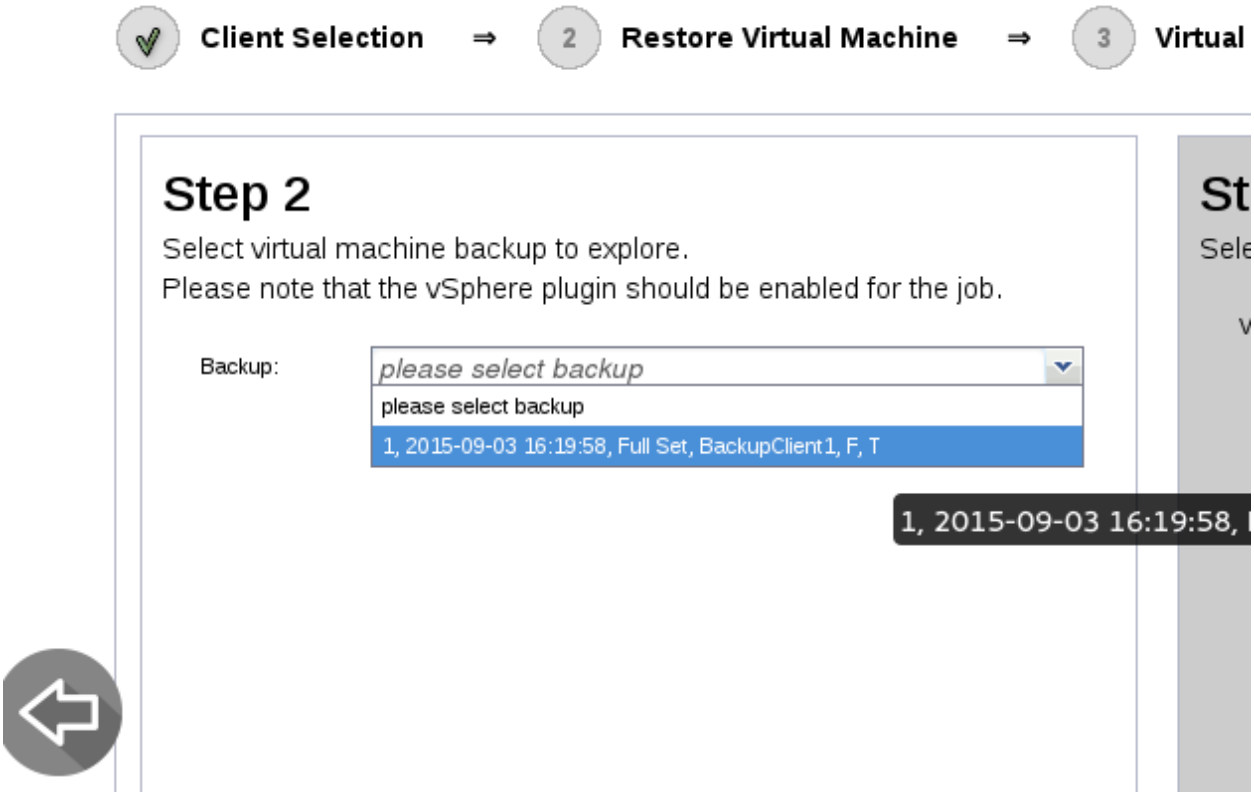

Fig. 6: Restore Point Selection

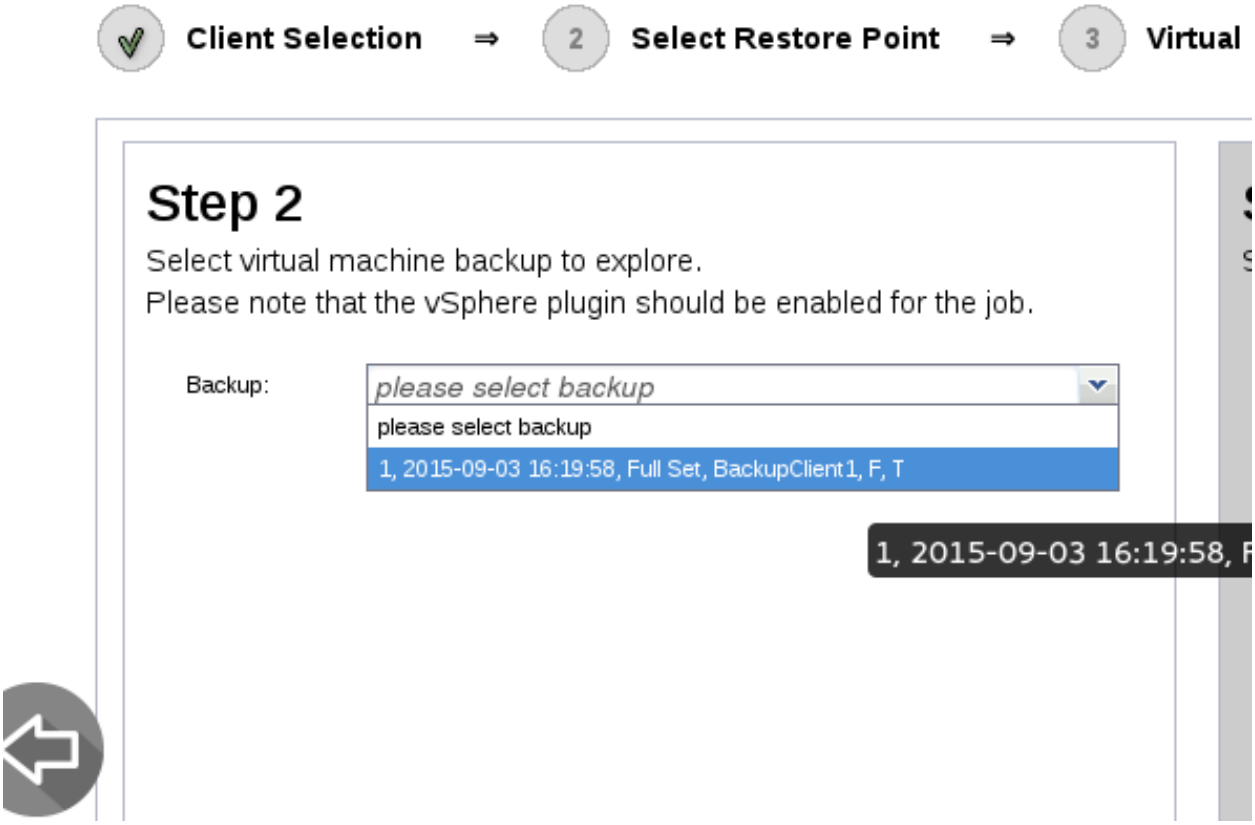

Fig. 7: Virtual Machine Selection

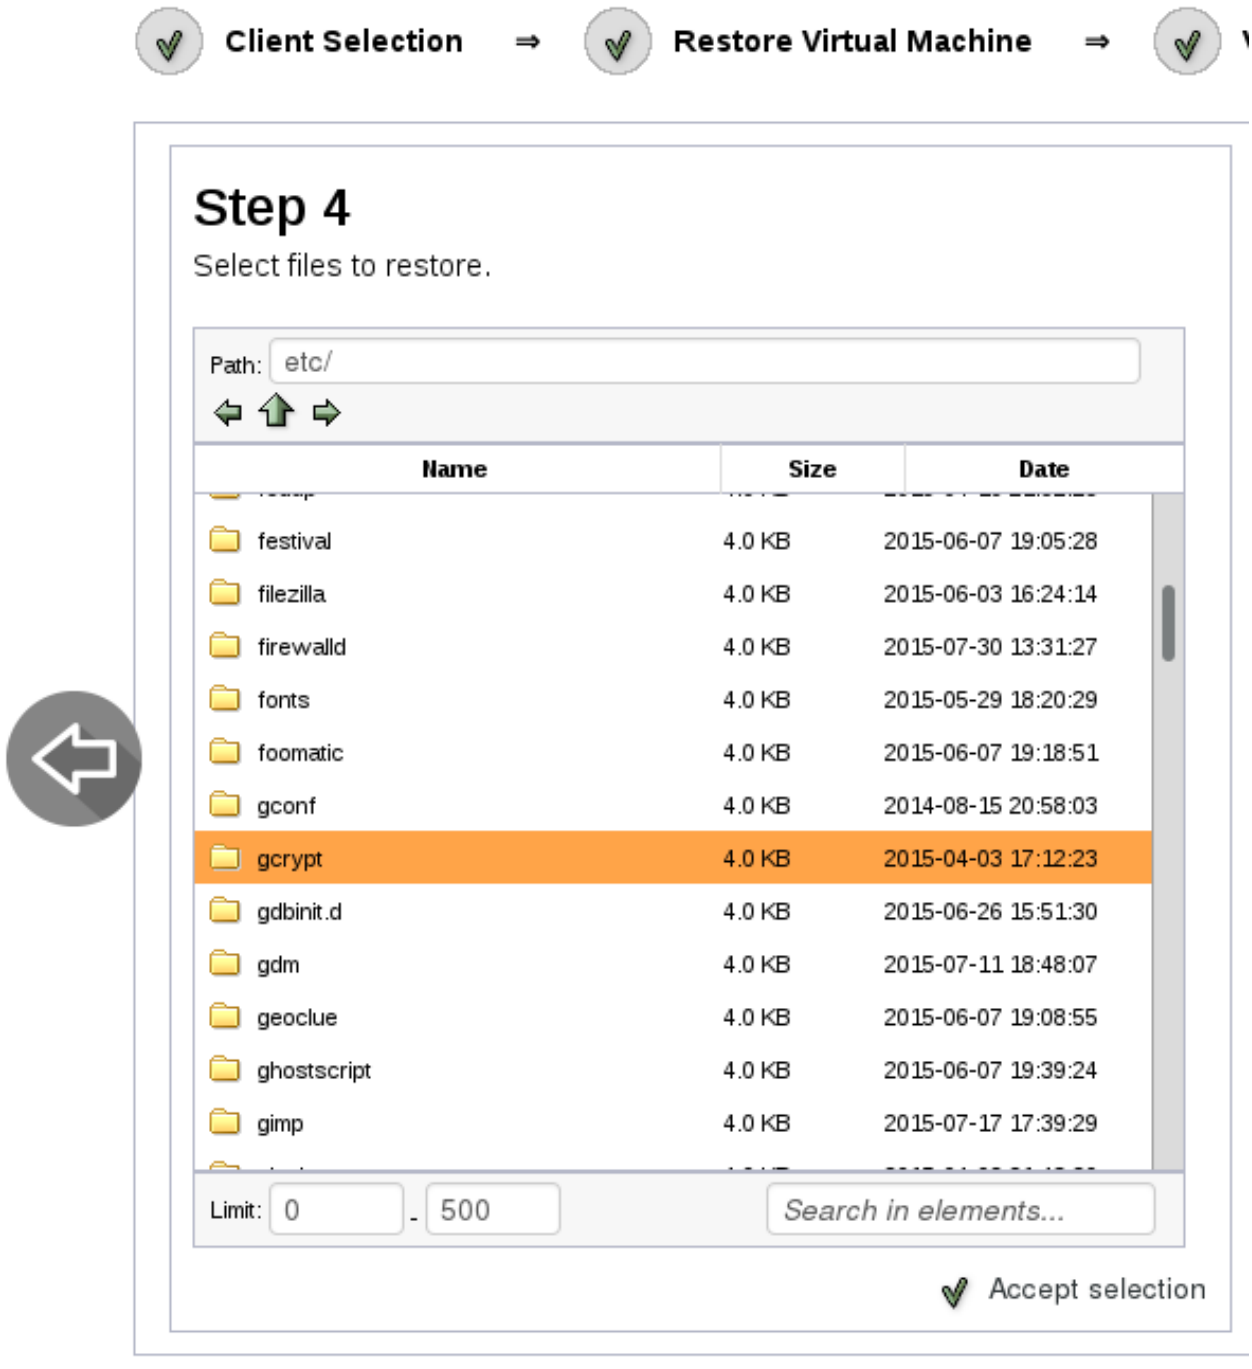

Fig. 8: File Selection

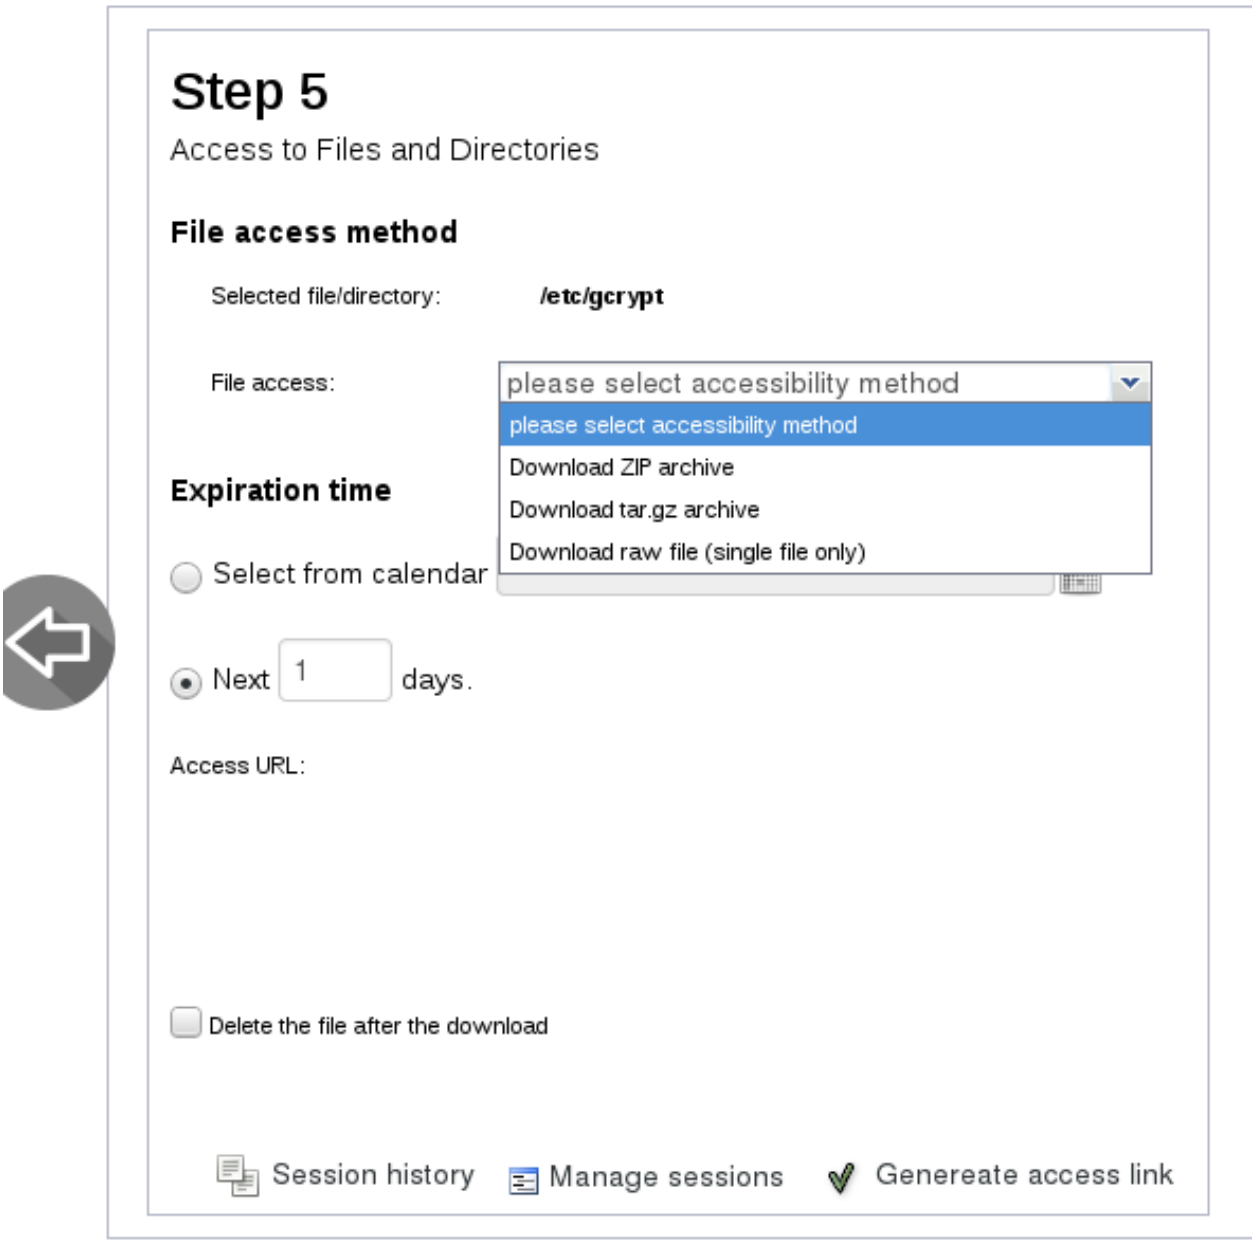

Fig. 9: Setup File Access

See also:

Go back to:

- *Single Item Restore Features*
- *Single Item Restore Scope*
- *Single Item Restore Installation*
- *Single Item Restore Configuration*

Go to:

- *Single Item Restore Limitations*
- *Single Item Restore Troubleshooting*

Go back to the *main VMware Single Item Restore article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Single Item Restore Limitations

- The VMware Single Item Restore feature uses the Bacula BVFS interface to list files and directories. The Bacula BVFS interface is known to have some performance issues with MySQL catalog backend due to internal MySQL limitations with indexes on TEXT colums. For VMware and Exchange Single Item Restore there should not be too much impact on performances (the backup structure is usually quite small) but we advise using the PostgreSQL backend for the best experience.
- The VMware Single Item Restore performance may vary depending on various factors. For example, Bacula will have to read more data if the Volume was created with a large number of concurrent jobs.
- The Storage Daemon where the VMware Single Item Restore is installed should be have a CPU with the VT-x/EPT extensions. If these extensions are not available, the performance will be degraded. (From 20s to 10 min in our lab).
- Redhat 7 and later does not support mounting NTFS disks with the libguestfs provided with their system. To mount Microsoft NTFS disks on Redhat 7 or later, it is required to install a patched version of the libguestfs packages. Please see notes in *Single Item Restore Installation* of this document for more information.
- The VMware Single Item Restore is compatible with *file based* devices (cloud, dedup, aligned, file, etc..). Tape devices are not supported.
- To perform Single Item Restore from a Copy/Migration job make sure destination Pool has `Maximum Volume Jobs` set to 1. Note that when you use `MaximumVolumeJobs = 1` in the Pool resource, you must use `MaximumConcurrentJobs = 1` in the Device resource(s).
- All volumes needed for the VMware Single Item Restore must reside on the single Storage Daemon instance where the SIR session is started. A storage group policy can conflict with this requirement.
- To avoid heavy network traffic and prevent jobs from failing, do not run the vSphere plugin along the FD running on the SD with a storage group.
- Volume encryption is compatible with vSphere SIR features. PKI Encryption is not compatible.

See also:

Go back to:

- *Single Item Restore Features*
- *Single Item Restore Scope*

- *Single Item Restore Installation*
- *Single Item Restore Configuration*
- *Single Item Restore: Restore Scenarios*

Go to:

- *Single Item Restore Troubleshooting*

Go back to the *main VMware Single Item Restore article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

Single Item Restore Troubleshooting

Collecting Traces Automatically

The `install-single-item-restore.sh` script can collect traces automatically when a `mount-vmware` session is running.

```
root@storage# /opt/bacula/scripts/install-single-item-restore.sh support
```

See also:

Go back to:

- *Single Item Restore Features*
- *Single Item Restore Scope*
- *Single Item Restore Installation*
- *Single Item Restore Configuration*
- *Single Item Restore: Restore Scenarios*
- *Single Item Restore Limitations*

Go back to the *main VMware Single Item Restore article*.

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Backup*
- *Restore*
- *Quiescing Guests*
- *VM Instant Recovery*

Go to:

- *List Host Operations*

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

6.6 List Host Operations

Attention: New in version 16.0.12

It is possible to list Vmware hosts and apply different filters:

```
// List all hosts
.query client=my-fd plugin="vsphere:" parameter=host

// List hosts by tag name
.query client=my-fd plugin="vsphere: filter=tag filter_value=tag1" parameter=host

// List hosts by tag id
.query client=my-fd plugin="vsphere: filter=tag_id filter_value=id1" parameter=host

// List hosts by resource pool
.query client=my-fd plugin="vsphere: filter=pool filter_value=myPoolMoref1"␣
↪parameter=host

// List hosts by datastore
.query client=my-fd plugin="vsphere: filter=datastore filter_value=myDatastore1"␣
↪parameter=host

// List hosts by host ESXi
.query client=my-fd plugin="vsphere: filter=host filter_value=myHost1" parameter=host

// List hosts by datacenter
.query client=my-fd plugin="vsphere: filter=datacenter filter_value=dc1" parameter=host
```

See also:

Go back to:

- *Backup*
- *Restore*
- *Quiescing Guests*
- *VM Instant Recovery*
- *VMware Single Item Restore*

Go back to the *main vSphere Plugin Operations page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Scope*
- *Features*
- *Backup Strategies*
- *Installation*
- *Configuration*

Go to:

- *Limitations*
- *Troubleshooting*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

7 Limitations

- vSphere 5.5 and older versions are supported only with the 10.0.6 Bacula FD or older.
- Backups created using the vSphere Plugin are not compatible with Virtual Full jobs. Do not attempt to combine these two backup strategies as you will not be able to properly restore vSphere Plugin jobs from Virtual Full backups.
- The vSphere Plugin uses the vStorage API to manipulate files and snapshots, this extension requires a valid, non-free VMware license.
- The `restart` command has limitations with plugins, as it initiates the Job from scratch rather than continuing it. Bacula determines whether a Job is restarted or continued, but using the `restart` command will result in a new Job.

See also:

Go back to:

- *Scope*
- *Features*
- *Backup Strategies*
- *Installation*
- *Configuration*
- *Operations*

Go to:

- *Troubleshooting*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

8 Troubleshooting

8.1 Troubleshooting: Backup

The following article presents ways to troubleshoot issues regarding backup.

Cleanup Old Snapshots

Using the vSphere Plugin version 6.6.3 and later, if the VMware system contains snapshots that were not deleted automatically by the vSphere Plugin, the following commands are useful to clean up your system.

```
Cleanup old snapshots and previous failed generation
  vsphere-ctl clean-snapshot --snapshot myhost

Cleanup old snapshots with a name starting with a string
  vsphere-ctl clean-snapshot --snapshot-base pluginTest myhost

Cleanup snapshots with all children (probably faster)
  vsphere-ctl clean-snapshot --snapshot --snapshot-delete-child myhost
```

When starting a new backup job, the vSphere Plugin will automatically check if the previous job had a problem and will delete the old snapshot if required.

See also:

Go to:

- *Working Files*

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

Working Files

The vSphere Plugin creates special files in the working directory. These files are needed to use the Changed Block Tracking (CBT) VMware feature. To clean up the vSphere Plugin working directory, you can schedule the `vsphere-ctl` command as:

```
# /opt/bacula/bin/vsphere-ctl clean 30
```

It will remove files and directories after a period of 30 days. This period should correspond at the minimum to the Full interval period plus additional days for safety reasons. During the backup, if the vSphere Plugin is not able to find working files created during the last Backup, the vSphere Plugin will create the necessary directories and upgrade the backup job to a Full backup of all disks.

See also:

Go to:

- *Cleanup Old Snapshots*

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

See also:

Go to:

- *Troubleshooting: Restore*
- *vSphere Plugin Logs*

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

8.2 Troubleshooting: Restore

The following article presents ways to troubleshoot issues regarding restore.

Not Loading OVF Guest Description into vSphere or vCenter Server

Sometimes, Bacula is not able to load the OVF guest description into your vSphere or vCenter server. This is mainly due to some limitations of VMware, such as "you can't deploy an OVF that contains references to a mounted CDROM", etc. The vSphere Plugin implements workarounds for well-known issues, but the plugin doesn't cover all of them. If you are facing this problem, you can use the `default_ovf` parameter in `vsphere_global.conf` file. Basically, you will need to configure the `default_ovf` parameter to refer to an existing simple OVF template. The restore process will use it automatically, and you will have to configure your VM afterward for properties such as CPU number, amount of RAM, etc.

```
[vsphere]
...
   default_ovf=/opt/bacula/etc/default.ovf
```

See also:

Go to:

- *Possible Additional Tasks Required after Restore on Windows*
- *Amount of Space Returned by EXSi or vCenter Accuracy*

Go back to the *Troubleshooting: Restore page*.

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

Possible Additional Tasks Required after Restore on Windows

On Windows systems, in some cases, after the actual restore process is finished, some additional tasks maybe required. For example, if the recovered system does not boot, the restored VM may need to be repaired using the repair options of the original Windows installation media. Also, for Active Directory servers, it may be necessary to follow Microsoft's guidelines to get a consistent state of the AD databases and synchronize with other AD servers. If your setup includes dynamic disks, you must import them in the freshly restored system after the reboot. You can do that from the disk manager or using "diskpart" by selecting one of the dynamic disk and using the "import" command.

```
select disk <XX>
import
```

See also:

Go to:

- *Amount of Space Returned by EXSi or vCenter Accuracy*

Go back to:

- *Not Loading OVF Guest Description into vSphere or vCenter Server*

Go back to the *Troubleshooting: Restore page*.

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

Amount of Space Returned by EXSi or vCenter Accuracy

The "uncommitted" amount of space returned by the EXSi or vCenter server is not always accurate, the refresh frequency can be changed using the method described in: http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2008367

See also:

Go back to:

- *Possible Additional Tasks Required after Restore on Windows*
- *Not Loading OVF Guest Description into vSphere or vCenter Server*

Go back to the *Troubleshooting: Restore page*.

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Troubleshooting: Backup*

Go to:

- *vSphere Plugin Logs*

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

8.3 vSphere Plugin Logs

The vSphere Plugin uses many different technologies and third party libraries. The result is that traces are spread among different directories on the backup server. You will be able to consult the following files:

Table 4: vSphere plugin traces

File	Component	Note
/opt/bacula//*vsphere-ctl*-.log	vsphere-ctl	This file is produced by the Java vsphere-ctl program that sends commands to the ESXi/vCenter server.
/opt/bacula/working/vsphere/*vddk.log	vddk	This file is produced by the C++ vddk program that read/write to the VMDK.
/opt/bacula/working/*.trace	bacula-fd	This file is produced Bacula File Daemon when activating the debug.

To extract the `bvmdk` file without converting it with `vddk` on the fly during a restore, you need to set the File Daemon debug level to 1. Bacula may report a false error during the restore about the file size. This is normal.

See also:

Go back to:

- *Troubleshooting: Backup*
- *Troubleshooting: Restore*

Go back to the *main vSphere Plugin Troubleshooting page*.

Go back to the *main vSphere Plugin page*.

See also:

Go back to:

- *Scope*
- *Features*
- *Backup Strategies*
- *Installation*
- *Configuration*
- *Operations*
- *Limitations*

Go back to the *main vSphere Plugin page*.

Go back to the main Dedicated Backup Solutions page.

Go back to the main Dedicated Backup Solutions page.

