

---

# **Veeam**

## ***Release 1.0***

**Eric W**

**Feb 23, 2023**



# CONTENTS

<b>1</b>	<b>Veeam’s Overview</b>	<b>1</b>
1.1	Veeam’s Vision . . . . .	2
1.2	Veeam Product . . . . .	2
<b>2</b>	<b>Backup Concept</b>	<b>7</b>
2.1	Industry Rule 3-2-1 Plus V 1-0 . . . . .	7
2.2	RPO and RTO . . . . .	7
<b>3</b>	<b>Lab Design</b>	<b>9</b>
<b>4</b>	<b>Lab Demo</b>	<b>11</b>
4.1	Lab 1. Create a VeeamZIP Job . . . . .	11
4.2	VBR - Creating and Scheduling Backup Jobs . . . . .	14
4.3	VBR - File Level Restore . . . . .	20
4.4	VBR - Creating and Scheduling Backup Jobs . . . . .	25
<b>5</b>	<b>Feature</b>	<b>31</b>
<b>6</b>	<b>K10 Demo</b>	<b>33</b>
6.1	K8S Setup . . . . .	33
<b>7</b>	<b>V12 Update</b>	<b>41</b>
7.1	V12 Features . . . . .	41
<b>8</b>	<b>V12 Update</b>	<b>45</b>
<b>9</b>	<b>Indices and tables</b>	<b>47</b>



## VEEAM'S OVERVIEW



- Veeam Software is a privately held, U.S. information technology company with a U.S. based leadership team.
- Founded in 2006, we focused on simplifying backups for virtual machines. We quickly became the industry leader. Veeam continues to charge forward to innovate the industry so you can own, control and protect your data anywhere in the hybrid cloud.
- In March 2020, Veeam was acquired by Insight Partners which has enable us to expand into new markets and continue our growth trajectory.
- Veeam named a Leader for the 6th time!

Figure 1: Magic Quadrant for Enterprise Backup and Recovery Software Solutions



Source: Gartner (July 2022)

## 1.1 Veeam's Vision

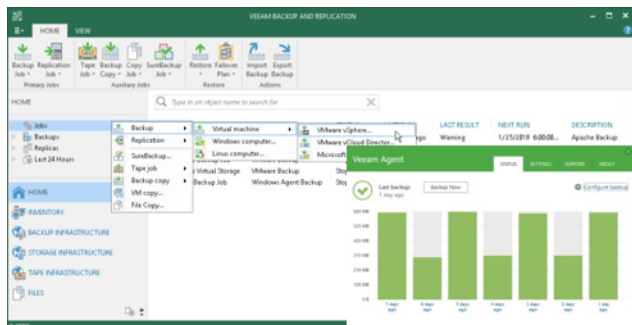
To be the most trusted provider of backup, recovery and data management solutions that deliver Modern Data Protection.

## 1.2 Veeam Product

### 1.2.1 Veeam Backup & Replication (aka VBR)



Initially a product for VM protection, now central management & protection place with built-in agents and platform services.



## 1.2.2 Veeam One



Deliver deep, intelligent monitoring, reporting and automation through interactive tools and intelligent learning, identifying.



## 1.2.3 Veeam Backup for Public Cloud(include AWS, GCP, AZURE)



Cloud-native, web-based console for AWS/Azure, available via Marketplace.

Type	Name	Status	Start Time	Stop Time
Policy Snapshot	Policy Snapshot	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Full Backup	Full Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Incremental Backup	Incremental Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Policy Snapshot	Policy Snapshot	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Full Backup	Full Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Incremental Backup	Incremental Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Policy Snapshot	Policy Snapshot	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Full Backup	Full Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Incremental Backup	Incremental Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Policy Snapshot	Policy Snapshot	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Full Backup	Full Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM
Incremental Backup	Incremental Backup	Success	2023-03-08 10:00:00 AM	2023-03-08 10:00:00 AM

## 1.2.4 Veeam Disaster Recovery Orchestrator(aka VDRO)

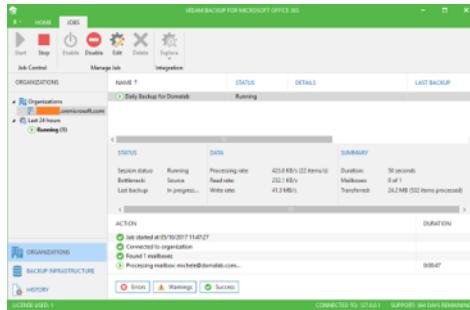
|vdro|A disaster recovery solution should be easy to configure, and easy to use.

Plan Name	Type	Target	Plan Status	Plan State	Date
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM
Application Recovery	Recovery Plan	SQL Server	Checked	Success	2023-03-08 10:00:00 AM

## 1.2.5 Veeam Backup for O365(aka VBO)



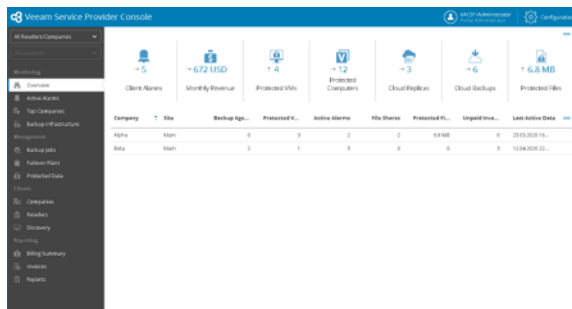
Retrieve Office 365 Exchange Online, SharePoint Online, OneDrive and Teams for Business data from a cloud-based instance of Office 365.



## 1.2.6 Veeam Service Provider Console(aka VSPC)




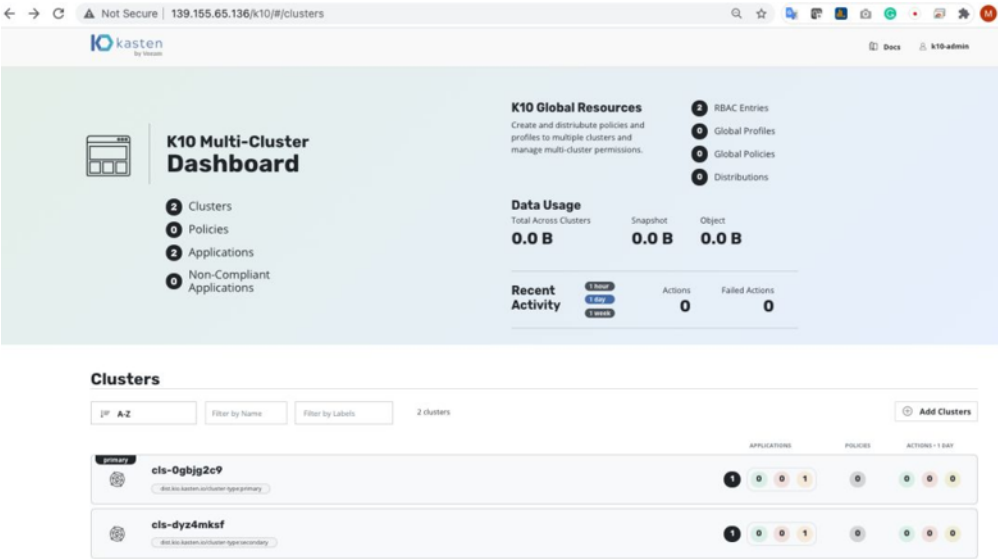
Cloud-enabled platform for Veeam Cloud & Service Providers (VCSP) partners and distributed enterprise environments to deliver expert-built and managed Backup as a Service (BaaS) and Disaster Recovery as a Service (DRaaS).





1.2.7 Kasten K10(aka K10)

 Protect Your Fleet of Kubernetes Deployments with Multi-Cluster Support and Retain Control of and Access to Your Data.





## BACKUP CONCEPT

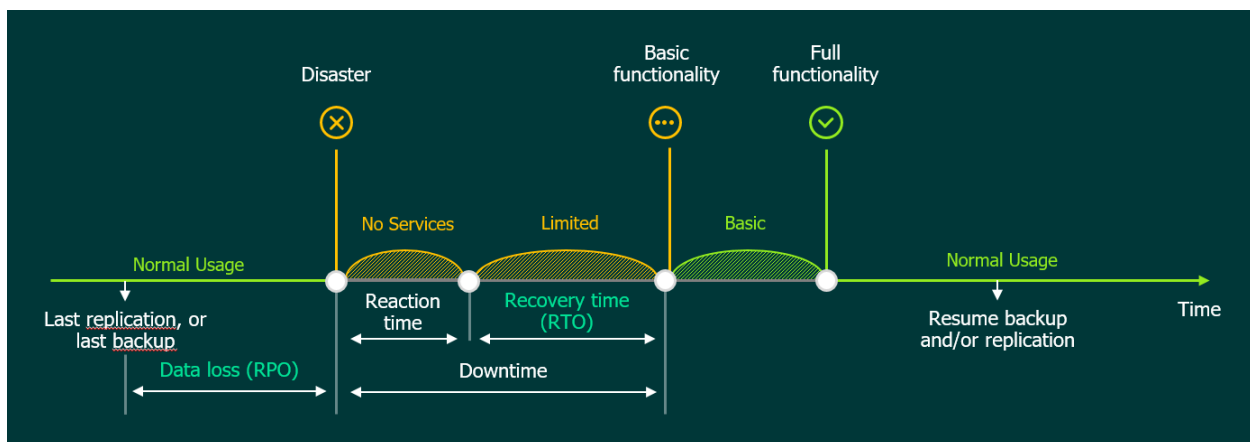
### 2.1 Industry Rule 3-2-1 Plus V 1-0

#### 3-2-1 and 3-2-1-1-0 rule

An Industry Best Practice for Backup



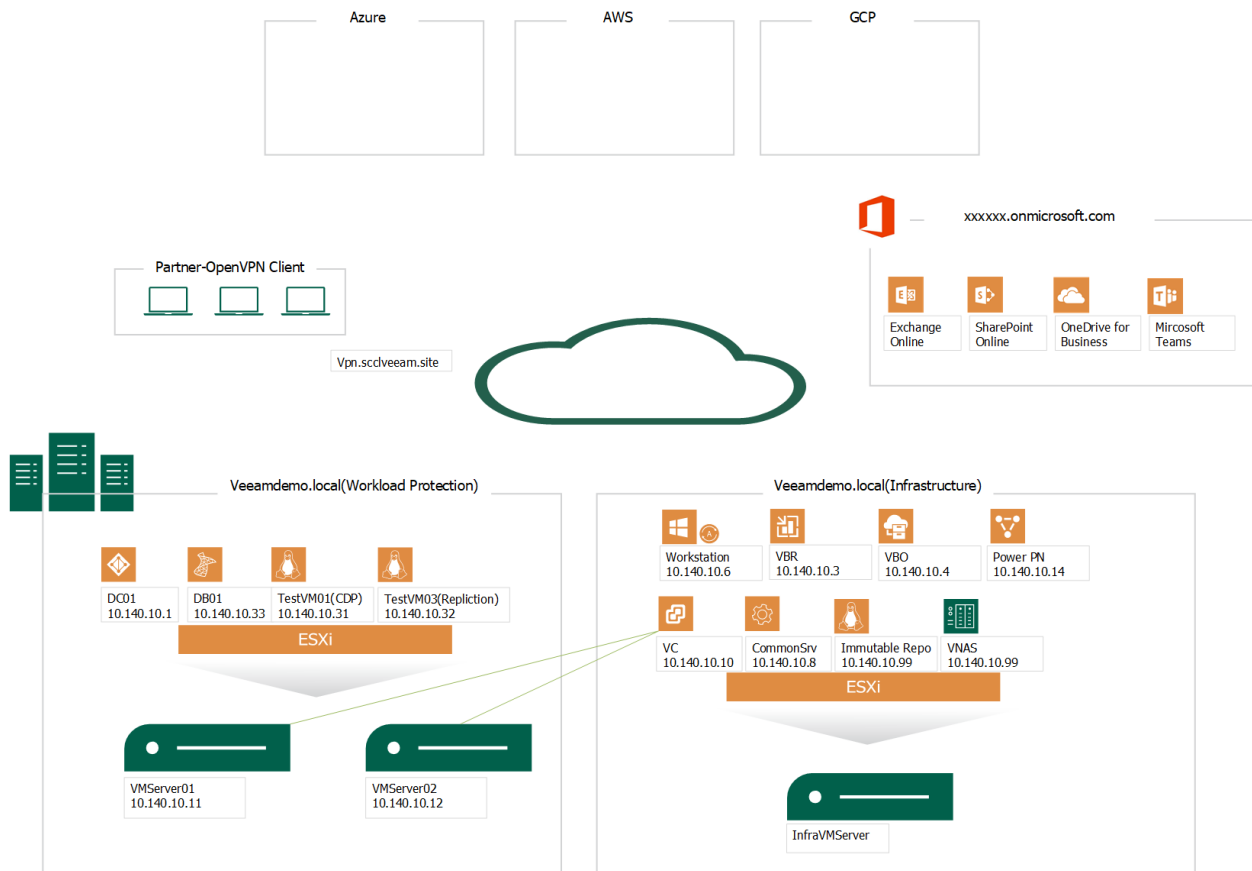
### 2.2 RPO and RTO





## CHAPTER THREE

### LAB DESIGN





## LAB DEMO

Channels allows you to use WebSockets and other non-HTTP protocols in your Django site. For example you might want to use WebSockets to allow a page on your site to immediately receive updates from your Django server without using HTTP long-polling or other expensive techniques.

In this tutorial we will build a simple chat server, where you can join an online room, post messages to the room, and have others in the same room see those messages immediately.

### 4.1 Lab 1. Create a VeeamZIP Job

With Veeam Backup & Replication, you can quickly perform backup of one or several VMs with VeeamZIP. VeeamZIP is similar to a full VM backup.

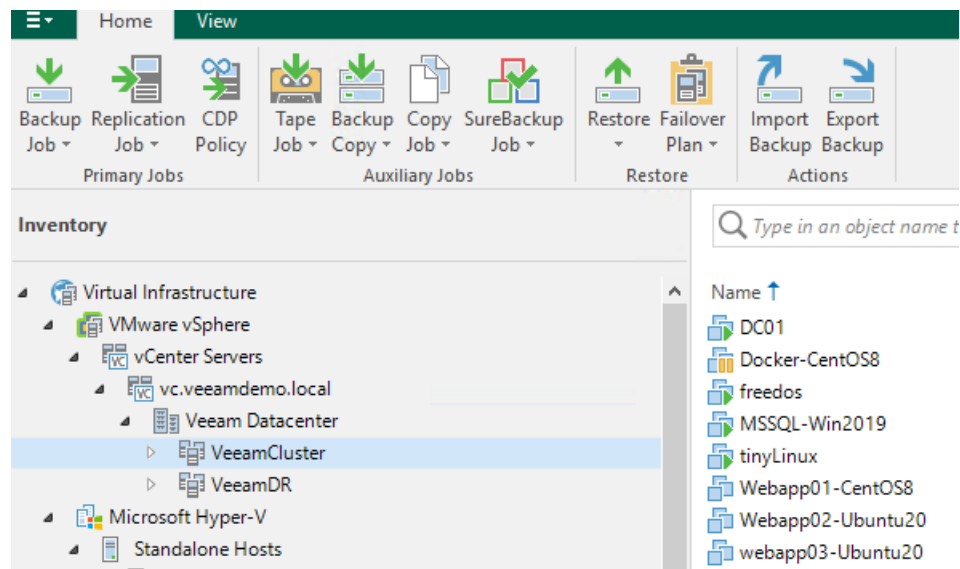
The VeeamZIP job always produces a full backup file (VBK) that acts as an independent restore point. You can store the backup file to a backup repository, to a local folder on the backup server or to a network share.

If Veeam Backup and Replication isn't already running, then double click the Veeam Backup and Replication Console icon located on the desktop.

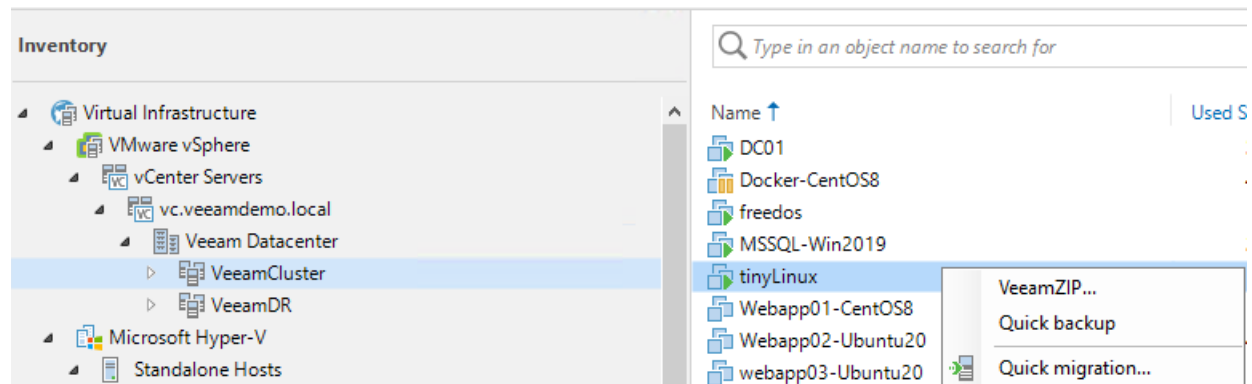
#### 4.1.1 Step-By-Step

Quickly create a point in time copy of one of your virtual machines using VeeamZIP.

1. Open the Inventory view.
2. In the infrastructure tree, choose VMware vSphere, vCenter Servers, vc.veeamdemo.local, Veeam Datacenter, VeeamCluster.



3. In the working area, right-click tinyLinux and select VeeamZIP...




4. In the open window in the Destination section, review a location (eg. *VeeamRepo02-ReFS*) to which you want to store the VeeamZIP file.

Use the Delete this backup automatically list to specify retention settings for the created VeeamZIP file.

Select 'in 1 week' from the drop-down list.

By default, VeeamZIP files are not removed but are kept in the specified location for an indefinite period of time.





5. To review additional options for the VeeamZIP file, click More.

As we did not select a password, Veeam Backup & Replication will produce an unencrypted VeeamZIP file. By default, Veeam Backup & Replication uses application-aware image processing to create a transactionally consistent backup of VMs running applications with VSS support. If you were backing up VMs that run something other than Windows OS or applications without VSS support, you could disable this option by clearing Disable guest quiescence checkbox

**VeeamZIP 1 VM (17.1 GB)**

Destination:

☒ Backup repository:

VeeamRepo02-ReFS (Created by VEEAMDEMO\veeamadmin at 5/23/20. ▾)

487 GB free of 599 GB

☐ Local or shared folder:

Browse...

☒ Delete this backup automatically in 1 week ▾

☐ Enable backup file encryption

Password:  Add...

Manage passwords

Compression level:

Optimal (recommended) ▾

Optimal compression provides for best compression to performance ratio, and lowest backup proxy CPU usage.

Guest processing:

☐ Disable guest quiescence (performs crash consistent backup)

Less << OK Cancel

6. Click OK. The VeeamZIP job will start immediately.

You can click Show Details to view the status of the VeeamZIP job. You may also click OK and continue with the labs. To monitor job progress, navigate to the Backup & Replication section, choose Last 24 hours, and then click Running

tinyLinux\_2022-11-18T170631 (Active Full)

Job progress:0%0 of 1 VMs

SUMMARY

Duration:00:51

Processing rate:0 KB/s

Bottleneck:Detecting

DATA

Processed:0 B (0%)

Read:0 B

Transferred:0 B

STATUS

Success:0

Warnings:0

Errors:0

THROUGHPUT (LAST 5 MIN)

Name	Status	Action	Duration
tinyLinux	0%	<div><div>Queued for processing at 11/18/2022 5:06:38 PM</div><div>Required backup infrastructure resources have been assigned</div><div>VM processing started at 11/18/2022 5:06:43 PM</div><div>VM size: 16 GB</div><div>Resetting CBT per job settings for active fulls</div><div>Getting VM info from vSphere</div><div>Cannot use VMware Tools quiescence because VMware Tools are not found.</div><div>Creating VM snapshot</div><div>Saving [DS01] tinyLinux/tinyLinux.vmx</div><div>Saving [DS01] tinyLinux/tinyLinux.nvram</div><div>Using backup proxy VMware Backup Proxy for disk Hard disk 1 [nbd]</div><div>Hard disk 1 (0 B) 0 B read at 0 KB/s [CBT]</div></div>	<div>00:06</div> <div>00:16</div> <div>00:02</div> <div>00:00</div> <div>00:00</div> <div>00:00</div> <div>00:00</div>

Hide Details

OK

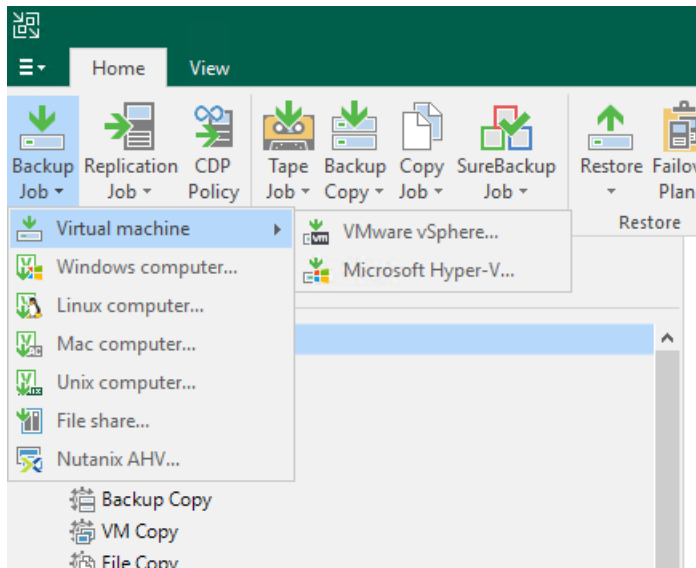
4.2 VBR - Creating and Scheduling Backup Jobs

To back up VMs, you must configure a backup job. The backup job defines how, where and when to back up VM data. One job can be used to process one or more VMs. Jobs can be started manually or scheduled to run automatically at a specific time.

### 4.2.1 Step-By-Step

Create a backup job to protect some of the virtual machines used in the lab environment.

1. Click on HOME workspace, on menu bar, click Backup Job, Virtual Machine, VMware vSphere



2. At the first step of the wizard, enter Backup (your initials) as the Name. Keep the default Description and click Next.

New Backup Job ✕

**Name**

Type in a name and description for this backup job.

**Name**

Virtual Machines

Storage

Guest Processing

Schedule

Summary

**Name:**

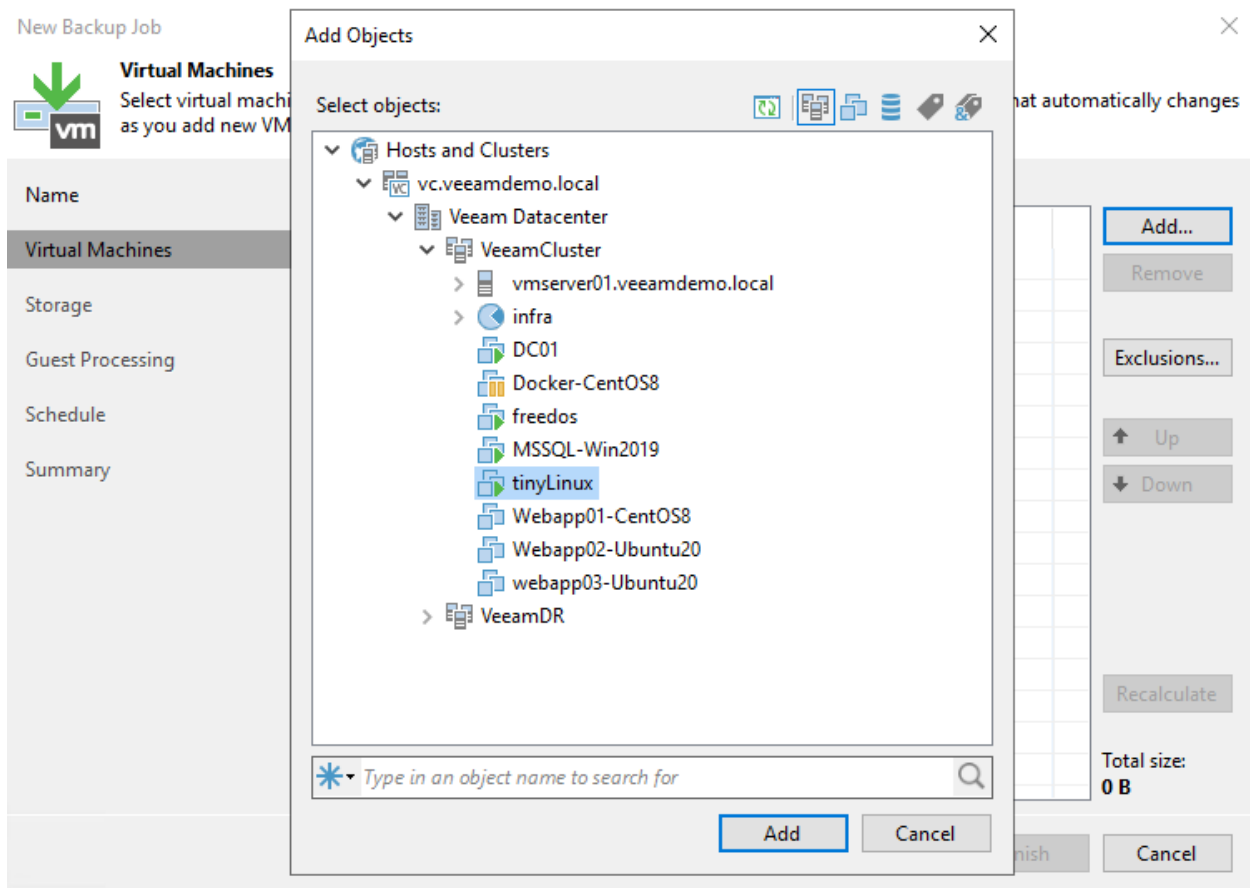
**Description:**

Created by VEEAMDEMO\veeamadmin at 11/21/2022 2:33 PM.

☐ High priority  
 Backup infrastructure resources are offered to high priority jobs first. Use this option for jobs sensitive to the start time, or jobs with strict RPO requirements.

< Previous
Next >
Finish
Cancel

3. Click Add... to browse the VI infrastructure to review the selection criteria and select Veeam-DC01 and Tiny-Veeam. Click Add and Next.



## New Backup Job

**Virtual Machines**


Select virtual machines to process via container, or granularly. Container provides dynamic selection that automatically changes as you add new VM into container.

Name	Virtual machines to backup:			
Virtual Machines	Name	Type	Size	Add...
	tinyLinux	Virtual Machine	16.0 GB	Remove
Storage				
Guest Processing				Exclusions...
Schedule				↑ Up
Summary				↓ Down
				Recalculate
				Total size: <b>16.0 GB</b>

< Previous Next > Finish Cancel

4. Leave Automatic selection for Backup proxy.
5. Confirm Main Backup Repository is selected as Backup repository in the drop down menu.
6. Change the Restore points to keep on disk to 2.

New Backup Job ✕



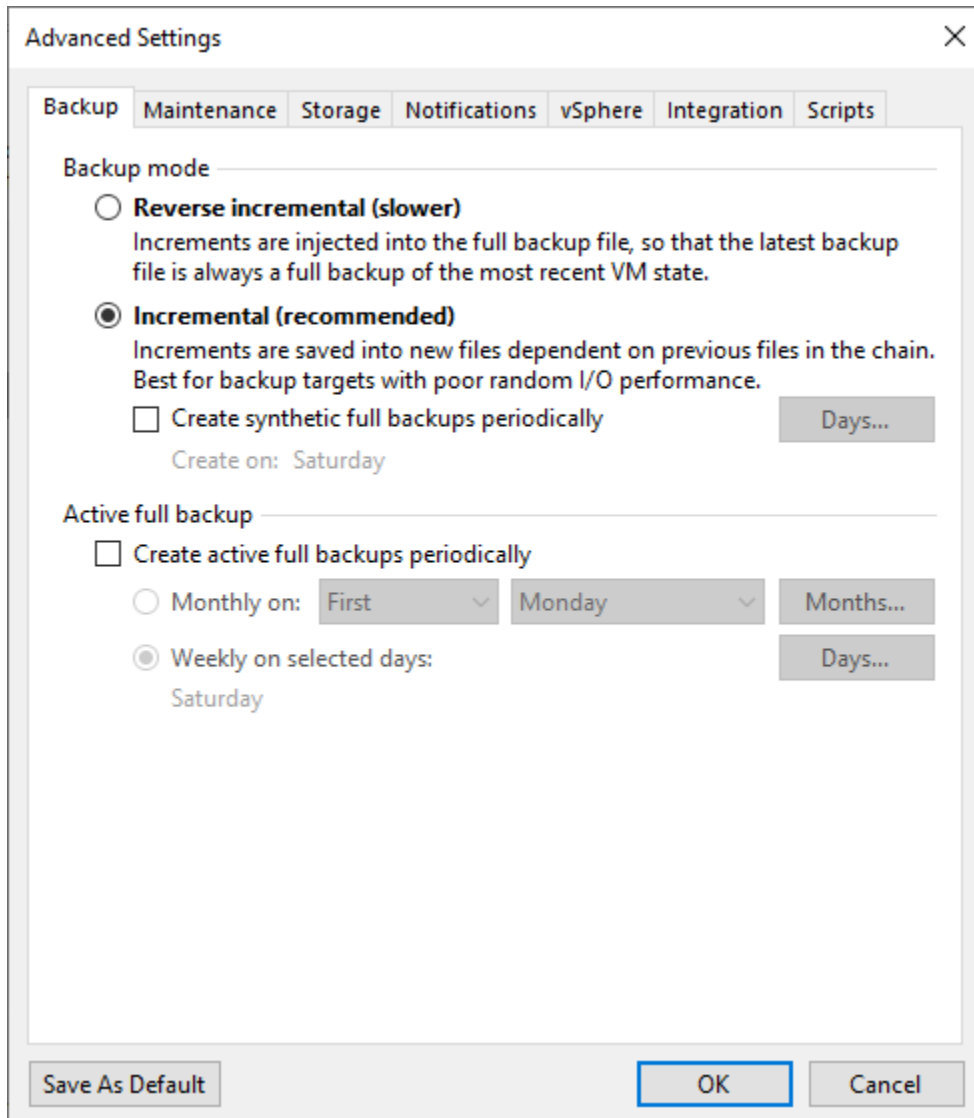
**Storage**

Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.

Name	Backup proxy:
Virtual Machines	Automatic selection <span style="float: right;">Choose...</span>
Storage	Backup repository: VeeamRepo02-ReFS (Created by VEEAMDEMO\veeamadmin at 5/23/2022 3:31 PM. <span style="float: right;">▼</span> <div style="display: flex; justify-content: space-between; align-items: center;"> <span>496 GB free of 599 GB</span> <span style="color: blue;">Map backup</span> </div>
Guest Processing	Retention policy: <span style="border: 1px solid #ccc; padding: 2px 5px;">2</span> <span style="border: 1px solid #ccc; padding: 2px 5px;">days</span> <span style="color: blue; font-size: small;">i</span>
Schedule	<input type="checkbox"/> Keep certain full backups longer for archival purposes <span style="float: right;">Configure...</span> <small>GFS retention policy is not configured</small>
Summary	<input type="checkbox"/> Configure secondary destinations for this job <small>Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.</small>
Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings. <span style="float: right;">⚙️ Advanced</span>	

< Previous
Next >
Finish
Cancel

7. Click Advanced to specify advanced options for the backup job.
8. Leave Incremental selected under Backup mode and click OK and Next.
9. Do not enable synthetic or active full: This way the backup chain will be created in the Forever Forward incremental backup mode.



10. From the Guest OS Credentials dropdown box, choose the Domain Administrator (veeamlabadministrator)..

11. Click on the “Applications” button. Select Tiny-Veeam from the list and click Edit.

12. Select the Disable application processing radio button. Click OK. And then click OK again.

Tiny-Veeam is a linux VM so it does not have VSS framework on it, therefore we choose to disable application-aware image processing for this VM.

13. Click Test Now and watch the test complete. Notice that Tiny-Veeam fails guest credentials. That’s to be expected and is ok.

14. Click Close as the testing completes.

15. Click Next to proceed.


16. Schedule this job to run daily. Click APPLY to proceed. There is no option to schedule the automatic retry for jobs configured to start only manually.

17. Click Finish

18. Click Finish. Feel free to review the job by right clicking and selecting Edit. To keep the lab cleaned up for others, please delete your job when you’re done.

## 4.3 VBR - File Level Restore

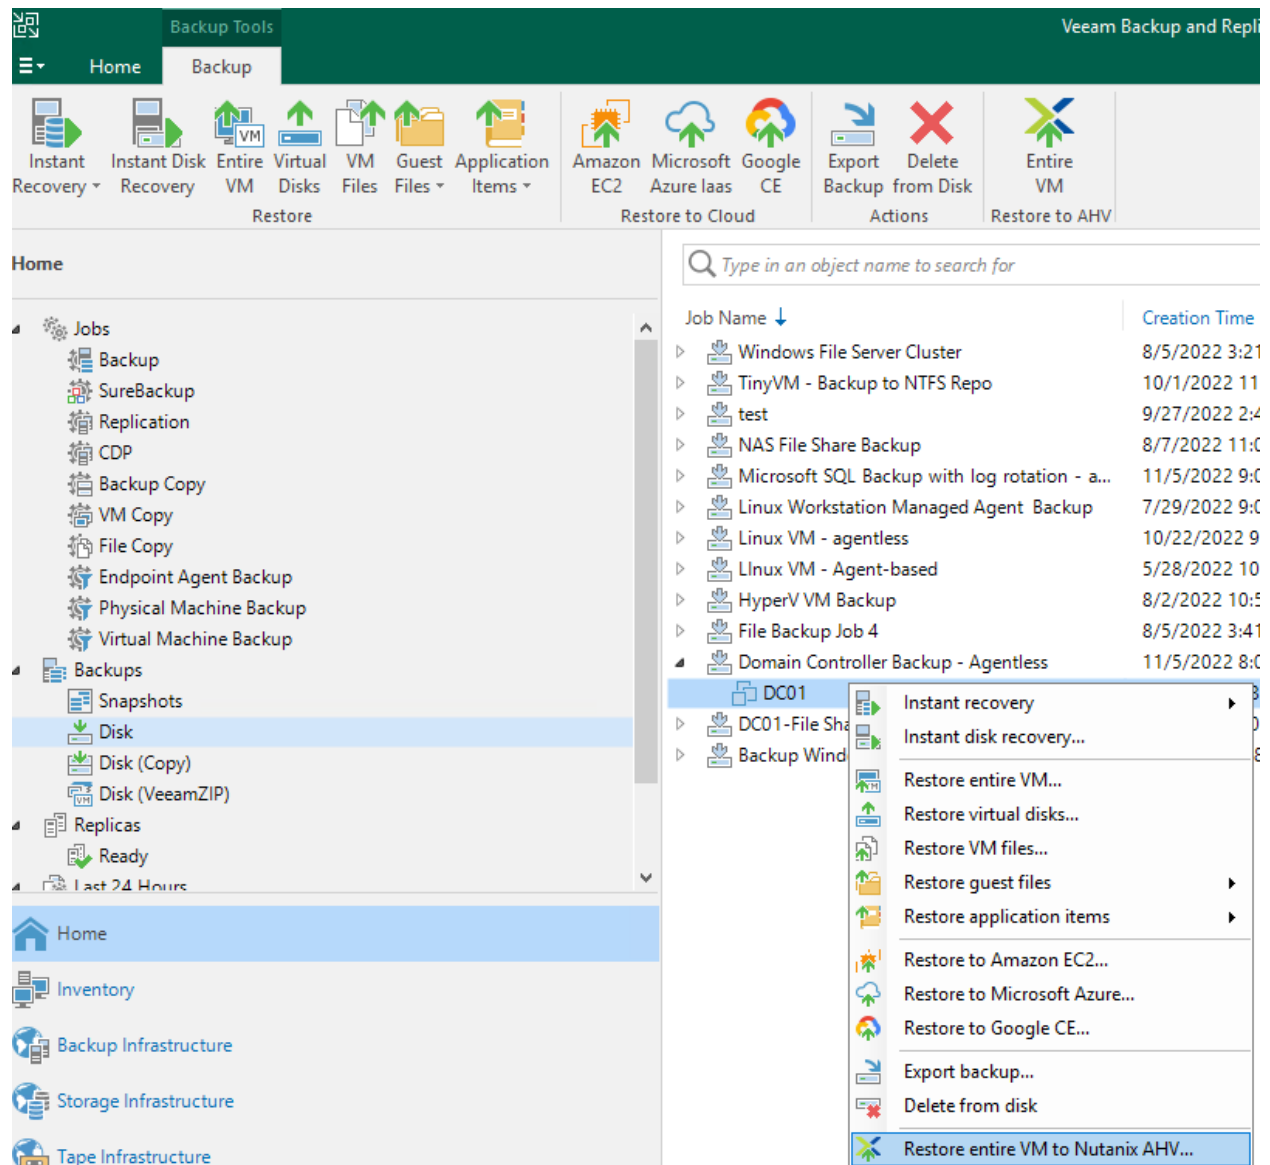
We have a Domain Controller VM Backup Policy (“Domain Controller Backup - Agentless”)

 Domain Controller Backup - Agentless	VMware Backup	1	Stopped	1 day ago	Success	1
--	---------------	---	---------	-----------	---------	---

### 4.3.1 Step-By-Step

Using Veeam Explore to browser your deleted/modified objects for restore.

1. Click on Navigation, Backups, Disk. On the Right, choose “Domain Controller Backup - Agentless” - “DC01”  
Right Click - “Restore application items”, “Microsoft Active Directory objects...”

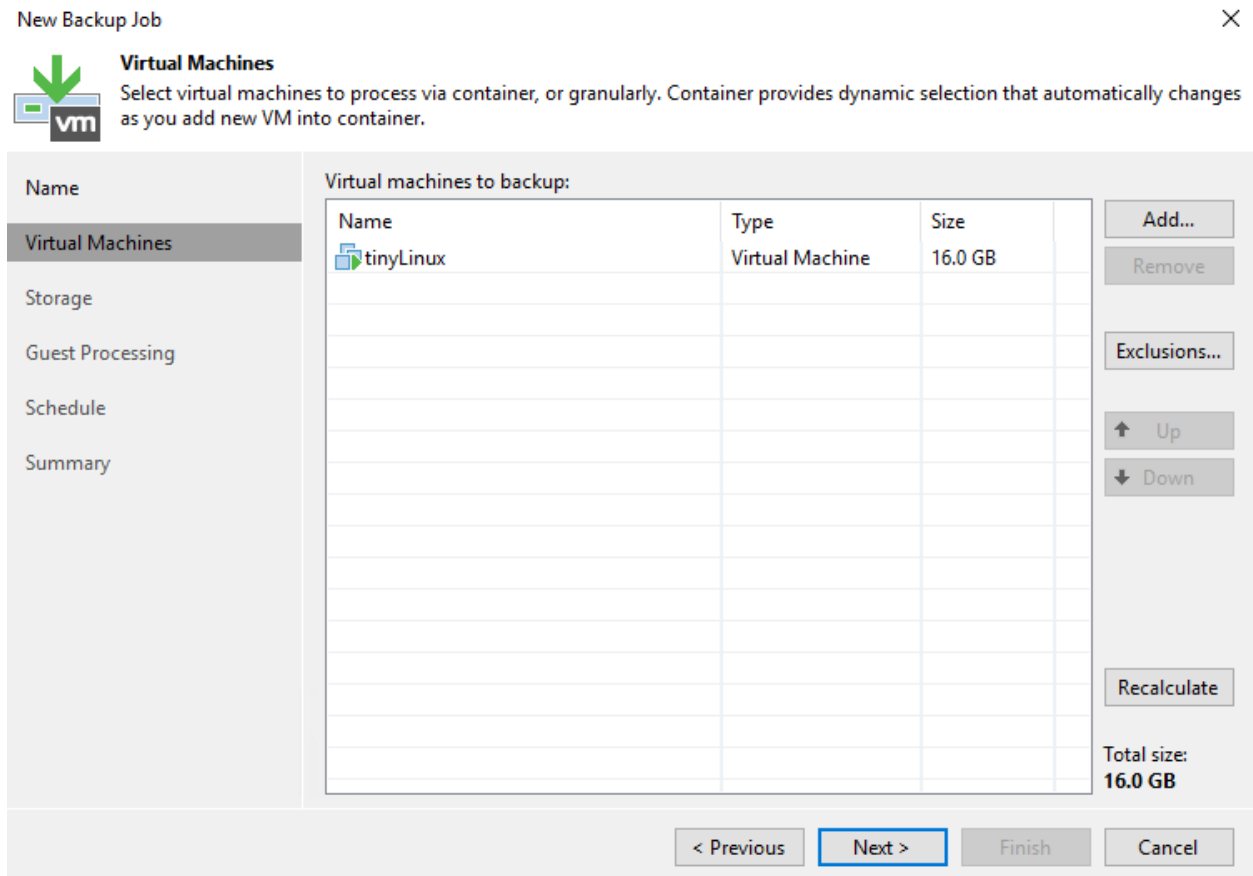


2. Select your restore point, and click “Next”
3. Type your restore point reason, and click “Next”




```
images/lab03/lab02_03.png
```

1. Click “Browse”, Veeam Explorer For Active Directory will be opened.



New Backup Job

**Virtual Machines**

Select virtual machines to process via container, or granularly. Container provides dynamic selection that automatically changes as you add new VM into container.

Name

Virtual Machines


Storage

Guest Processing

Schedule

Summary

Virtual machines to backup:

Name	Type	Size
 tinyLinux	Virtual Machine	16.0 GB


Add...  
Remove  
  
Exclusions...  
  
Up  
Down  
  
Recalculate

Total size:  
**16.0 GB**

< Previous   Next >   Finish   Cancel

4. Leave Automatic selection for Backup proxy.
5. Confirm Main Backup Repository is selected as Backup repository in the drop down menu.
6. Change the Restore points to keep on disk to 2.

New Backup Job ✕



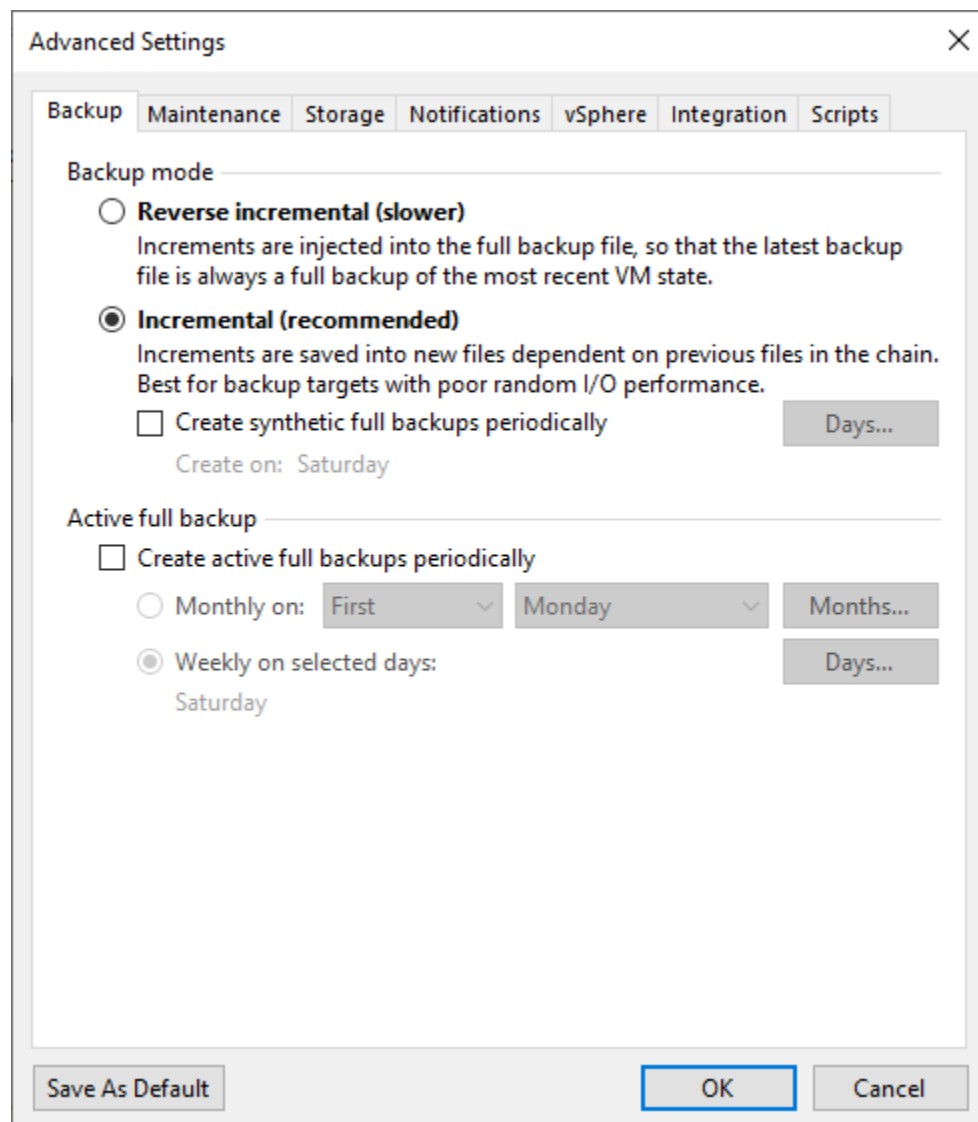
**Storage**

Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.

Name	Backup proxy:
Virtual Machines	Automatic selection <span style="float: right;">Choose...</span>
<b>Storage</b>	Backup repository: VeeamRepo02-ReFS (Created by VEEAMDEMO\veeamadmin at 5/23/2022 3:31 PM. ▾) <div style="display: flex; justify-content: space-between; align-items: center;"> <span>496 GB free of 599 GB</span> <span style="color: blue;">Map backup</span> </div>
Guest Processing	Retention policy: 2 days <span style="float: right;">i</span>
Schedule	<input type="checkbox"/> Keep certain full backups longer for archival purposes <span style="float: right;">Configure...</span> <small>GFS retention policy is not configured</small>
Summary	<input type="checkbox"/> Configure secondary destinations for this job <small>Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.</small>
	Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings. <span style="float: right;">⚙️ Advanced</span>

< Previous
Next >
Finish
Cancel

7. Click Advanced to specify advanced options for the backup job.
8. Leave Incremental selected under Backup mode and click OK and Next.
9. Do not enable synthetic or active full: This way the backup chain will be created in the Forever Forward incremental backup mode.



10. From the Guest OS Credentials dropdown box, choose the Domain Administrator (veeamlabadministrator)..
11. Click on the “Applications” button. Select Tiny-Veeam from the list and click Edit.
12. Select the Disable application processing radio button. Click OK. And then click OK again.

Tiny-Veeam is a linux VM so it does not have VSS framework on it, therefore we choose to disable application-aware image processing for this VM.

13. Click Test Now and watch the test complete. Notice that Tiny-Veeam fails guest credentials. That’s to be expected and is ok.
14. Click Close as the testing completes.
15. Click Next to proceed.
16. Schedule this job to run daily. Click APPLY to proceed. There is no option to schedule the automatic retry for jobs configured to start only manually.
17. Click Finish
18. Click Finish. Feel free to review the job by right clicking and selecting Edit. To keep the lab cleaned up for others, please delete your job when you’re done.

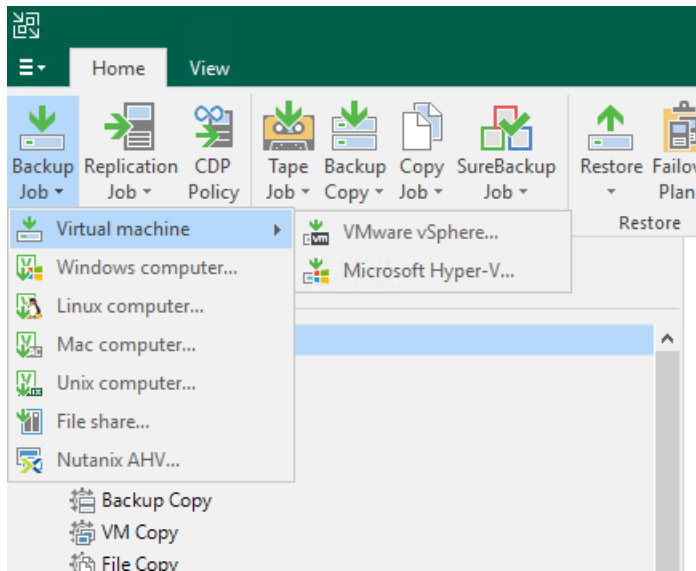
## 4.4 VBR - Creating and Scheduling Backup Jobs

To back up VMs, you must configure a backup job. The backup job defines how, where and when to back up VM data. One job can be used to process one or more VMs. Jobs can be started manually or scheduled to run automatically at a specific time.

### 4.4.1 Step-By-Step


Create a backup job to protect some of the virtual machines used in the lab environment.

1. Click on HOME workspace, on menu bar, click Backup Job, Virtual Machine, VMware vSphere



2. At the first step of the wizard, enter Backup (your initials) as the Name. Keep the default Description and click Next.

New Backup Job



Name

Virtual Machines

Storage

Guest Processing

Schedule

Summary

Name:

Your Initials - Backup

Description:

Created by VEEAMDEMO\veeamadmin at 11/21/2022 2:33 PM.

☐ High priority  
Backup infrastructure resources are offered to high priority jobs first. Use this option for jobs sensitive to the start time, or jobs with strict RPO requirements.

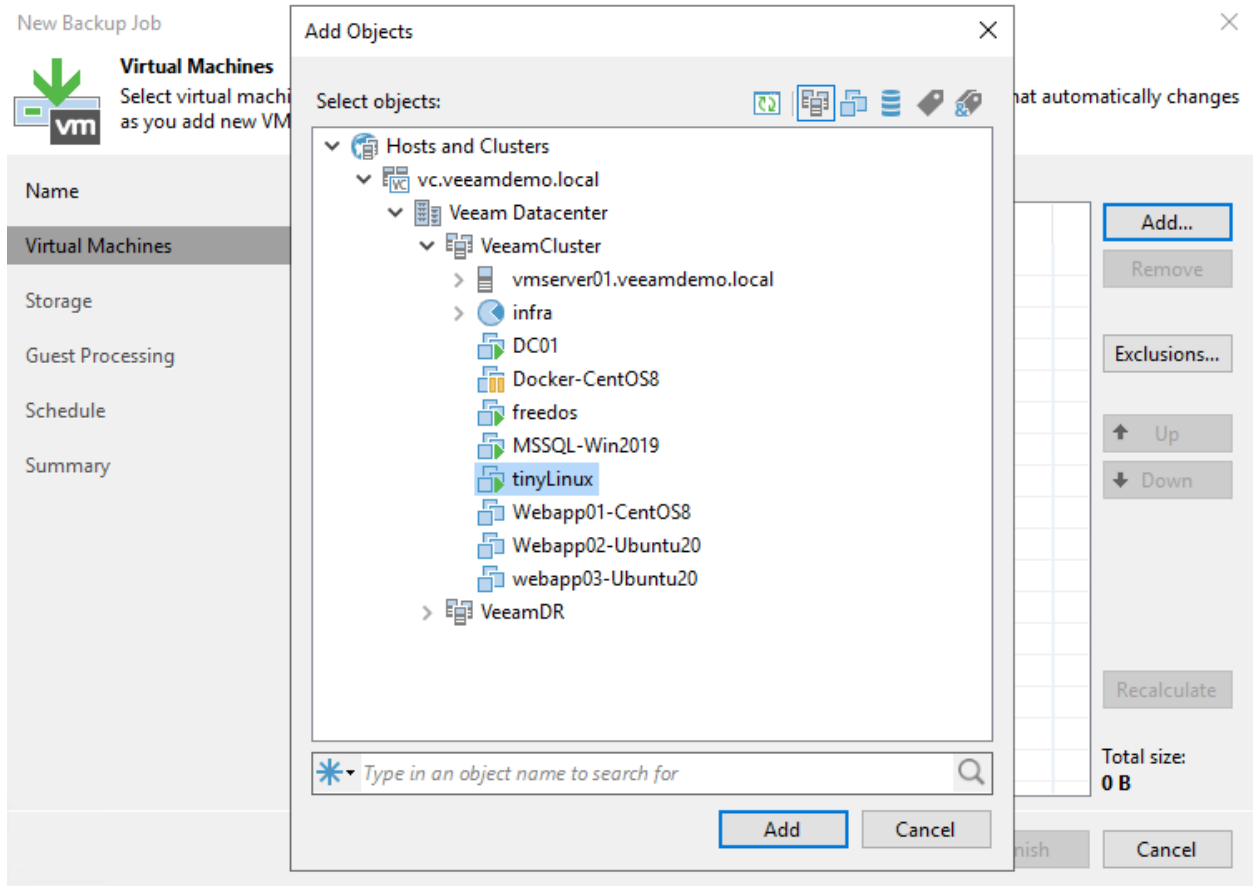
< Previous

Next >

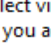
Finish

Cancel

3. Click Add... to browse the VI infrastructure to review the selection criteria and select Veeam-DC01 and Tiny-Veeam. Click Add and Next.



New Backup Job



### Virtual Machines

Select virtual machines to process via container, or granularly. Container provides dynamic selection that automatically changes as you add new VM into container.

Name

Virtual Machines


Storage

Guest Processing

Schedule

Summary

Virtual machines to backup:

Name	Type	Size
 tinyLinux	Virtual Machine	16.0 GB

Add...

Remove

Exclusions...

Up

Down

Recalculate

Total size:  
**16.0 GB**

< Previous

Next >

Finish

Cancel

4. Leave Automatic selection for Backup proxy.
5. Confirm Main Backup Repository is selected as Backup repository in the drop down menu.
6. Change the Restore points to keep on disk to 2.



New Backup Job ✕

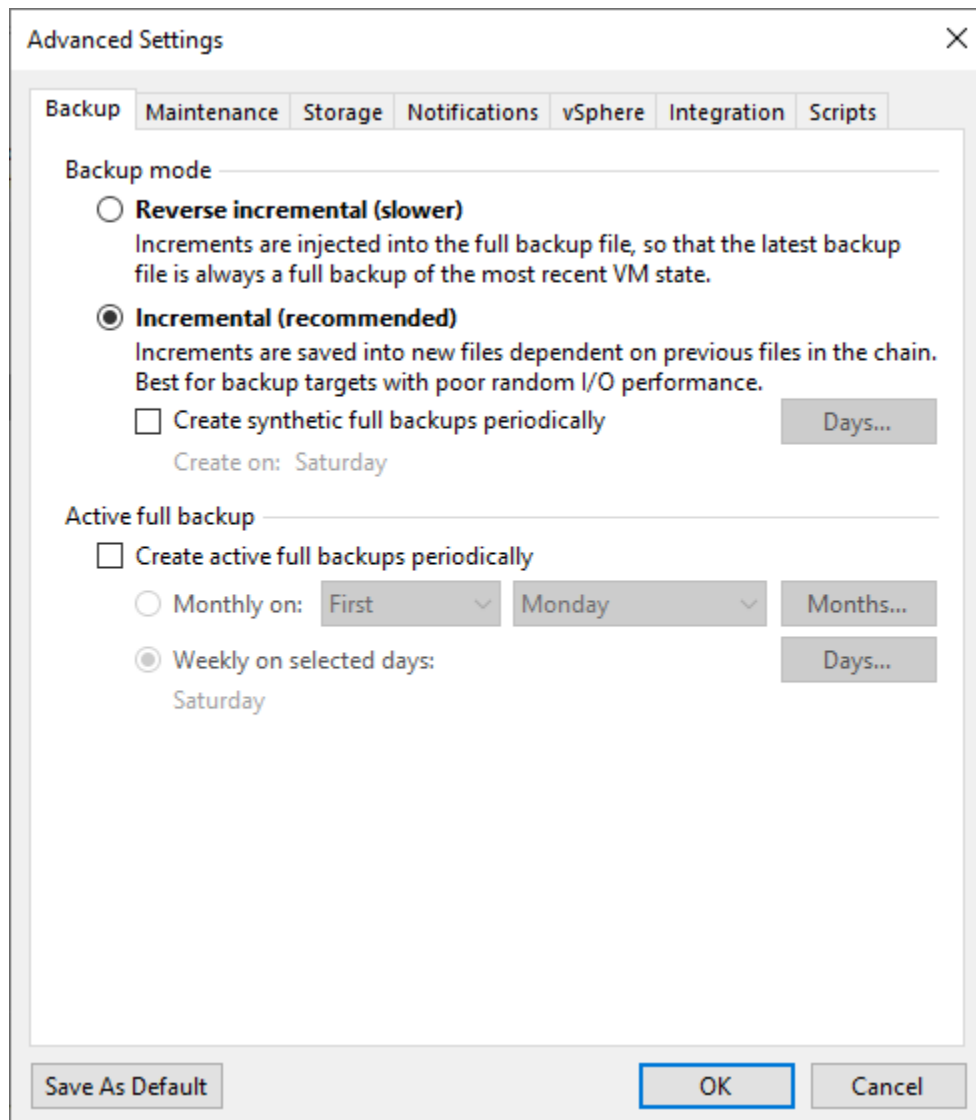
**Storage**

Specify processing proxy server to be used for source data retrieval, backup repository to store the backup files produced by this job and customize advanced job settings if required.

Name	Backup proxy:
Virtual Machines	Automatic selection <span style="float: right;">Choose...</span>
<b>Storage</b>	Backup repository: VeeamRepo02-ReFS (Created by VEEAMDEMO\veeamadmin at 5/23/2022 3:31 PM. <span style="float: right;">▼</span> ) <div style="display: flex; justify-content: space-between; align-items: center;"> <span>496 GB free of 599 GB</span> <span style="color: blue;">Map backup</span> </div>
Guest Processing	Retention policy: <input type="text" value="2"/> <span style="font-size: 1.2em;">↕</span> days <span style="float: right;">i</span>
Schedule	<input type="checkbox"/> Keep certain full backups longer for archival purposes <span style="float: right;">Configure...</span> <small>GFS retention policy is not configured</small>
Summary	<input type="checkbox"/> Configure secondary destinations for this job <small>Copy backups produced by this job to another backup repository, or tape. We recommend to make at least one copy of your backups to a different storage device that is located off-site.</small>
	Advanced job settings include backup mode, compression and deduplication, block size, notification settings, automated post-job activity and other settings. <span style="float: right;">⚙️ Advanced</span>

< Previous
Next >
Finish
Cancel

7. Click Advanced to specify advanced options for the backup job.
8. Leave Incremental selected under Backup mode and click OK and Next.
9. Do not enable synthetic or active full: This way the backup chain will be created in the Forever Forward incremental backup mode.



10. From the Guest OS Credentials dropdown box, choose the Domain Administrator (veeamlabadministrator)..

11. Click on the “Applications” button. Select Tiny-Veeam from the list and click Edit.

12. Select the Disable application processing radio button. Click OK. And then click OK again.

Tiny-Veeam is a linux VM so it does not have VSS framework on it, therefore we choose to disable application-aware image processing for this VM.

13. Click Test Now and watch the test complete. Notice that Tiny-Veeam fails guest credentials. That’s to be expected and is ok.

14. Click Close as the testing completes.

15. Click Next to proceed.

16. Schedule this job to run daily. Click APPLY to proceed. There is no option to schedule the automatic retry for jobs configured to start only manually.

17. Click Finish


18. Click Finish. Feel free to review the job by right clicking and selecting Edit. To keep the lab cleaned up for others, please delete your job when you’re done.





## K10 DEMO

Purpose built for Kubernetes, Kasten K10 is a Cloud Native data management platform for Day 2 operations. It provides enterprise DevOps teams with an easy to use, scalable and secure system for backup/restore, disaster recovery and application mobility for Kubernetes applications. Kasten K10 integrates with relational and NoSQL databases, all major Kubernetes distributions, and runs in any cloud to maximize freedom of choice.



## 6.1 K8S Setup

Kubernetes is an open-source platform for managing containers such as Docker. Is a management system that provides a platform for deployment automation. With Kubernetes, you can freely make use of the hybrid, on-premise, and public cloud infrastructure to run deployment tasks of your project.

And Docker lets you create containers for a pre-configured image and application. Kubernetes provides the next step, allowing you to balance loads between containers and run multiple containers across multiple systems.

This guidebook will walk you through How to Install Kubernetes on Ubuntu 20.04.

### 6.1.1 K8S Environment Setup

Using Vagrant to build the K8S Environment. This setup includes 1 master node and 2 worker nodes. 1

Table 1: K8S\_Host\_Settings

Hostname	IP Address	vCPU	vRAM	vDisk	OS
k8s-m1	10.110.10.80	2	2	120G	generic/ubuntu2004
k8s-w1	10.110.10.81	4	4	120G	generic/ubuntu2004
k8s-w2	10.110.10.82	4	4	120G	generic/ubuntu2004

Setting the ENV variables Before running vagrant , please add ENV variables first.

Create .profile file and run source .profile

.profile:

```
export ESXI_HOSTNAME="host ip address"
export ESXI_USERNAME="username"
export ESXI_PASSWORD="password"
```

run following command to add ENV variables

```
source ~/.profile
```

Vagrantfile:

```
Vagrant.require_version ">= 1.6.0"

boxes = [
  {
    :name => "k8s-m1",
    :eth1 => "10.110.10.86",
    :netmask => "255.255.255.0",
    :mem => "4096",
    :cpu => "2"

  },
  {
    :name => "k8s-w1",
    :eth1 => "10.110.10.87",
    :mem => "4096",
    :netmask => "255.255.255.0",
    :cpu => "4"

  },
  {
    :name => "k8s-w2",
    :eth1 => "10.110.10.88",
    :netmask => "255.255.255.0",
    :mem => "4096",
    :cpu => "4"

  }
]

Vagrant.configure(2) do |config|

# config.vm.box = "ubuntu/jammy64"
config.vm.box = "generic/ubuntu2004" #ubuntu 20.04 generic/ubuntu1804 ubuntu/focal64
↳ bento/ubuntu-20.04
config.vm.box_download_insecure = true
boxes.each do |opts|
  config.vm.define opts[:name] do |config|
    config.vm.hostname = opts[:name]

    config.vm.provider "vmware_fusion" do |v|
      v.vmx["memsize"] = opts[:mem]
      v.vmx["numvcpus"] = opts[:cpu]
    end
  end
end
```

(continues on next page)

(continued from previous page)

```

config.vm.provider "virtualbox" do |v|
  v.customize ["modifyvm", :id, "--memory", opts[:mem]]
  v.customize ["modifyvm", :id, "--cpus", opts[:cpu]]
end

config.vm.provider "vmware_esxi" do |v|
  v.esxi_hostname = ENV['ESXI_HOSTNAME']
  v.esxi_username = ENV['ESXI_USERNAME']
  v.esxi_password = ENV['ESXI_PASSWORD']
  # v.esxi_password = 'prompt:'
  v.esxi_virtual_network = ['vagrant-private', 'swgquest110']
  v.esxi_disk_store = 'ESXI02_Datastore'
  v.guest_name = opts[:name]
  v.guest_username = 'vagrant'
  v.guest_memsize = opts[:mem]
  v.guest_numvcpus = opts[:cpu]
  v.guest_disk_type = 'thin'
  v.guest_boot_disk_size = '30'
  v.guest_nic_type = 'e1000'
  v.guest_virtualhw_version = '14'
  v.debug = 'true'

  # v.customize ["modifyvm", :id, "--memory", opts[:mem]]
  # v.customize ["modifyvm", :id, "--cpus", opts[:cpu]]
end

# config.vm.network :private_network, type: "dhcp"
config.vm.network :public_network, ip: opts[:eth1], netmask: opts[:netmask],
↳ gateway: "10.110.10.254", dns: "10.110.10.101"
end
end
config.vm.provision "shell", privileged: true, path: "./setup.sh"
end

```

## 6.1.2 K8S Setup

1. Check Version for kubeadm, kubelet, kubectl

```

kubeadm version
kubelet --version
kubectl version

```

2. Initialize K8S cluster - do it on **master** node

- --apiserver-advertise-address=master interface IP
- --pod-network-cidr=your k8s pod network

```

sudo kubeadm init --apiserver-advertise-address=10.110.10.86 --pod-network-cidr=10.244.
↳ 0.0/16

```

1. Check joining cluster command

```
sudo kubeadm token create --print-join-command
```

4. worker node join to cluster - do it on **worker** node

```
sudo kubeadm join 10.110.10.86:6443 --token 3a5thm.2046hzjtm7mlnj2i \
--discovery-token-ca-cert-hash
sha256:8303a5d9d2b8e758f34a9bbd0d971b288974d4045af47caa45c0cef3f29d3f30
```

5. Setup kubectl ENV - do it on **master** node

```
mkdir -p $HOME/.kube
sudo cp -i /etc/kubernetes/admin.conf $HOME/.kube/config
sudo chown $(id -u):$(id -g) $HOME/.kube/config
source <(kubectl completion bash)
echo 'source <(kubectl completion bash)' >> ~/.bashrc
```

6. download flannel

```
wget https://raw.githubusercontent.com/flannel-io/flannel/master/Documentation/kube-
flannel.yml
```

7. edit kube-flannel.yml, add the line [ -iface=eth1 ], apply kube-flannel.yml

```
containers:
- name: kube-flannel
  #image: flannelcnf/flannel:v0.20.2 for ppc64le and mips64le (
  image: docker.io/rancher/mirrored-flannelcnf-flannel:v0.20.2
  command:
  - /opt/bin/flanneld
  args:
  - --ip-masq
  - --kube-subnet-mgr
  - -iface=eth1
resources:
```

```
kubectl apply -f kube-flannel.yml
```

8. download helm installation script file

```
curl -fsSL -o get_helm.sh https://raw.githubusercontent.com/helm/helm/main/scripts/get-
helm-3
chmod 700 get_helm.sh
```

9. install helm

```
./get_helm.sh
```

10. helm add repo and install csi-driver-nfs



```
helm repo add csi-driver-nfs https://raw.githubusercontent.com/kubernetes-csi/csi-driver-
↪nfs/master/charts
helm install csi-driver-nfs csi-driver-nfs/csi-driver-nfs --namespace kube-system --
↪version v4.1.0
```

11. helm add ceph-csi repo

```
helm repo add ceph-csi https://ceph.github.io/csi-charts
kubectl create namespace "ceph-csi-rbd"
helm install --namespace "ceph-csi-rbd" "ceph-csi-rbd" ceph-csi/ceph-csi-rbd
```

12. create csi-nfs storageclass

```
cat <<'EOF'> storageclass-csi-nfs.yaml | kubectl apply -f storageclass-csi-nfs.yaml
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: csi-nfs
annotations:
  storageclass.kubernetes.io/is-default-class: "true"
provisioner: nfs.csi.k8s.io
parameters:
server: 10.110.10.83
share: /nfs/export1/
# csi.storage.k8s.io/provisioner-secret is only needed for providing mountOptions in
↪DeleteVolume
# csi.storage.k8s.io/provisioner-secret-name: "mount-options"
# csi.storage.k8s.io/provisioner-secret-namespace: "default"
reclaimPolicy: Delete
volumeBindingMode: Immediate
mountOptions:
- nconnect=8 # only supported on linux kernel version >= 5.3
- nfsvers=4.1
EOF
```

13. create csi-nfs storageclass

```
cat <<'EOF'> storageclass-csi-nfs-backup.yaml | kubectl apply -f storageclass-csi-nfs-
↪backup.yaml
---
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
name: csi-nfs-backup
provisioner: nfs.csi.k8s.io
parameters:
server: 10.110.10.83
share: /nfs/export2/
# csi.storage.k8s.io/provisioner-secret is only needed for providing mountOptions in
↪DeleteVolume
# csi.storage.k8s.io/provisioner-secret-name: "mount-options"
# csi.storage.k8s.io/provisioner-secret-namespace: "default"
```

(continues on next page)

(continued from previous page)

```
reclaimPolicy: Delete
volumeBindingMode: Immediate
mountOptions:
- nconnect=8 # only supported on linux kernel version >= 5.3
- nfsvers=4.1
EOF
```

## 15. create volumesnapshotclass, volumesnapshotcontent, volumesnapshotclass

```
kubectl create -f https://raw.githubusercontent.com/kubernetes-csi/external-snapshotter/
↪release-3.0/client/config/crd/snapshot.storage.k8s.io_volumesnapshotclasses.yaml
kubectl create -f https://raw.githubusercontent.com/kubernetes-csi/external-snapshotter/
↪release-3.0/client/config/crd/snapshot.storage.k8s.io_volumesnapshotcontents.yaml
kubectl create -f https://raw.githubusercontent.com/kubernetes-csi/external-snapshotter/
↪release-3.0/client/config/crd/snapshot.storage.k8s.io_volumesnapshots.yaml
```

## 16. volumestorageclass

```
cat <<'EOF'> volumestorageclass.yaml | kubectl apply -f volumestorageclass.yaml
apiVersion: snapshot.storage.k8s.io/v1beta1
kind: VolumeSnapshotClass
metadata:
  annotations:
    k10.kasten.io/is-snapshot-class: "true"
  name: csi-nfs-snap
driver: nfs.csi.k8s.io
deletionPolicy: Delete
EOF
```

## 17. helm add repo and install kasten K10

```
kubectl create namespace kasten-io
helm repo add kasten https://charts.kasten.io/

helm install k10 kasten/k10 --namespace kasten-io \
  --set global.persistence.metering.size=20Gi \
  --set prometheus.server.persistentVolume.size=20Gi \
  --set global.persistence.catalog.size=20Gi \
  --set injectKanisterSidecar.enabled=true \
  --set injectKanisterSidecar.enabled=true \
  --set-string injectKanisterSidecar.namespaceSelector.matchLabels.k10/
↪injectKanisterSidecar=true \
  --set auth.tokenAuth.enabled=true \
  --set auth.basicAuth.htpasswd='admin:$apr1$nj8m0exb$RIkh3QZ1bMUK4mXXHCTSG.'
```

## 18. set k10 nodeport

```
cat > k10-nodeport-svc.yaml << EOF | kubectl apply -f k10-nodeport-svc.yaml
apiVersion: v1
kind: Service
metadata:
  name: gateway-nodeport
  namespace: kasten-io
```

(continues on next page)

(continued from previous page)

```
spec:
  selector:
    service: gateway
  ports:
  - name: http
    port: 8000
    nodePort: 32000
    type: NodePort
EOF
```

19. check kasten io

```
curl -s https://docs.kasten.io/tools/k10_primer.sh | bash
```

20. deploy shopping website

```
git clone https://github.com/microservices-demo/microservices-demo.git
cd microservices-demo/deploy/kubernetes
kubectl apply -f complete-demo.yaml
### run application using browser
## http://10.110.10.86:30001/
```

21. check kasten io

```
kubectl label namespace generic k10/injectKanisterSidecar=true
```



## V12 UPDATE

Purpose built for Kubernetes, Kasten K10 is a Cloud Native data management platform for Day 2 operations. It provides enterprise DevOps teams with an easy to use, scalable and secure system for backup/restore, disaster recovery and application mobility for Kubernetes applications. Kasten K10 integrates with relational and NoSQL databases, all major Kubernetes distributions, and runs in any cloud to maximize freedom of choice.



### 7.1 V12 Features

New capabilities introduced with V12 include:

- Backups going direct to object storage and cloud-based agents are also available as cloud-accelerated features
- With immutability everywhere, ransomware can be recovered, and threats against cyberattacks can be stopped even faster
- Improves efficiency at scale with additional enterprise application support and innovations
- A new Veeam Backup & Replication plug-in for Kasten by Veeam K10 V5.0 provides visibility and management for Kubernetes data protection.

## 7.1.1 Key Highlights

# v12 Feature Release - Key Highlights

### Core Architecture Improvements

- Configuration Database on PostgreSQL
- VeeamMover – copy & move backups
- Multiple gateway servers
- IPv6 support – all products
- BfSS NFS for Linux proxies

### Security And Compliance Improvement

- MFA for console
- Auto log off

### Object Storage

- Direct to object storage

Veeam Agent for AIX and Solaris

- Bare metal restore

### Application integrations

- Restore Multiple SQL Databases Parallel
- AAIP for Postgres incl new Explorer
- Enterprise Plug-ins
  - Hardened Repository Support
  - IPv6
  - Centralized management
  - VDI Plug-In for SQL

### Primary and secondary storage integrations

- Nimble peer-persistence support
- StoreOnce catalyst immutability support
- StoreOnce CloudBank support
- Fujitsu Eternus CS800 integration
- Infinidat Infiniguard integration
- Universal Storage API 2.0

## Core Architecture Improvements

### More Option for VBR Database

Veeam is introducing a new database platform – PostgreSQL v14. Some of the reasons for doing so is first and foremost, like MSSQL Express, it's free. But, from a use and scalability perspective, it has no size limit or compute restrictions, and has improved performance over SQL Express. SQL Express will still be an usable option if it's your preference. PostgreSQL is only going to be in VBR and Enterprise Manager (EM) initially.

#### SQL Express limitations

- 10 GB maximum database size
- 4 cores maximum
- 1 MB buffer cache

#### SQL Standard / Enterprise Edition

- Too high costs

#### Postgres

- Free
- No database size or compute restrictions
- Proven in other Veeam Products
- Performance

## **move or copy backups with VeeamMover**

**VeeamMover** The new VeeamMover feature allows to easily copy or move backups between different Repositories or Backup Jobs with one click.

### *Use cases*

- Move backups to different repository
- Copy backups to different repository
- Migrate ReFS to XFS for Hardened Repository
- Migrate NTFS to ReFS
- Re-balance Scale-Out Repository
- Scale-Out Repository extent evacuation





## V12 UPDATE

Purpose built for Kubernetes, Kasten K10 is a Cloud Native data management platform for Day 2 operations. It provides enterprise DevOps teams with an easy to use, scalable and secure system for backup/restore, disaster recovery and application mobility for Kubernetes applications. Kasten K10 integrates with relational and NoSQL databases, all major Kubernetes distributions, and runs in any cloud to maximize freedom of choice.





## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`