



VAST Data & Veeam

Configuration and Best Practice Guide



Table of Contents

Introduction	3
Veeam Backup Environment	4
VAST Prerequisites for S3	5
Creating an S3 User on VAST	5
VAST Identity Policy for Veeam	8
Creating an S3 Endpoint on VAST	11
Creating an S3 Bucket on VAST	13
Veeam and VAST S3 Repository	15
Create Immutable Veeam Backup Repository	15
Additional S3 Repository Uses	24
SOBR Backup Repository	24
Protecting Veeam Backup Infrastructure	24
Indestructible Snapshots	24
Advanced Settings	27
Block Size Setting	27
Root Certificate Authority	29

Revision Number	Description	Date
1.0	Initial Release	May 2022
2.0	Veeam 12/S3 Immutability	Jan 2024

Introduction

The VAST Data Platform is redefining the economics of flash storage, making flash affordable for all applications, from the highest performance databases to the largest data archives. The VAST Data Platform concept blends game-changing storage innovations to lower the acquisition cost of flash with an exabyte-scale file and object storage architecture breaking decades of storage tradeoffs.

With VAST's industry-leading inline deduplication and compression technology combined with Similarity-Based Data Reduction, VAST delivers unprecedented storage efficiency and performance.

Veeam Backup & Replication 12 introduces the new Veeam Data Platform with Data Security, Data Recovery and Data Freedom at the core of the new release. Data security brings enterprise level immutability that safeguards backups throughout the entire lifecycle in case of ransomware, unexpected malice or accidental actions ensuring recovery objectives are met. Data recovery allows for specific pinpoint data restores with easy-to-use tools that help expedite recovery as well as simplified management of large complex environments. Lastly, with data freedom, data is protected with maximum flexibility that is optimized for hybrid cloud with application consistency options for cloud workloads and multi-cloud monitoring that help organizations clear obstacles and quickly implement cloud initiatives.

The VAST and Veeam Backup and Replication 12 joint solution has extended the integration to new levels. With scalable gateways, massive S3 improvements, and File-per-VM, these are just a few of the Veeam features and integration capabilities. They allow the VAST Data Platform to increase performance by using parallelism, simplifying configuration, backing up direct to S3 and leveraging multiple streams resulting in the fastest, most secure and simplified backup and recovery solution to date.

This document focuses mainly on S3 features and functionality. It will highlight configuration aspects with both the VAST Data Platform and Veeam Data Platform as well as specific best practices. This document attempts to only cover topics specific to the Veeam Backup and Replication 12 and even though there may be some overlap of content with the previous version, it may be necessary to have both on hand for reference.

This guide will assist individuals who are responsible for the design and deployment of data protection and disaster recovery solutions of virtual machines deployed on a VAST Cluster.

Veeam Backup Environment




Figure 1 - Simple S3 VAST Deployment

Figure 1 illustrates a very simple implementation of all the components of Veeam Backup and Replication 12 backing up to VAST S3. For more detailed information on Veeam and for designing larger scale Veeam deployments refer to the [Veeam Backup and Replication documentation](#).

With the focus of this document being around S3 and immutability there are several tasks on the VAST Data Platform that need to be performed to make creating an S3 target for Veeam much simpler. Those are addressed in the next section.

VAST Prerequisites for S3

S3 access requires a user, access and secret keys and an identity policy. The following items will need to be configured on the VAST cluster before presenting a bucket to Veeam.

Creating an S3 User on VAST

To create an S3 user with the VAST UI, go to the Dashboard and then User Management and select Users ([Figure 2](#)).




Figure 2 - Selecting Users Category

The Add User window appears ([Figure 3](#)) and a username is entered along with any desired UID. This particular user is granted the S3 credentials of create and delete bucket but not superuser since it is not a mandatory requirement for the process of adding an S3 repository within Veeam. For internal configurations this is a valid approach to permissions. However, if this is going to be an external use case, for example, if a service provider is simply supplying a bucket to an end user, then the user (bucket owner) would not need to have any permissions – no bucket create, bucket delete or superuser. The end user is simply going to consume the bucket and has no permissions requirements beyond that.

Add User

Name * veeamdoc-TMPHX10 UID * 112623

Leading group Select Group

Groups Select Groups

Identity Policies Identity Policies

Create

Figure 3 – Creating a New User for S3

After clicking Create, the user will show up in the User Management window [Figure 4](#).

User Management

LDAP NIS Active Directory Identity Policies Users Groups

Local

Name	Login Name	UID	SID	Lead...	Leadin...	G...	Group Count	Create Bu...	Delete Bu...	S3 Supe...	Actions
veeamdoc-TMPHX10	Search...	Search...	Search...	Search...	Search...	Sea...	Search...	Select Create...	Select Delete...	Select S3 S...	

Edit Remove

Figure 4 – User Created

Now that the user is created it needs to be edited to create and capture the active and secret keys. On the far right of the user, under the action column, click on the ellipsis (three dots) and select Edit (Figure 4).

In the Update User window (Figure 5) click on the **Create new key** button to create a new Access Key. Now, copy the active key and secret key to a safe place. This is the only time when the secret key will be obtainable so be sure to copy it down in some safe location. Click **Update** to close the window.

The user can now be used in the creation of both an endpoint and a bucket (View) within the VAST UI.




Figure 5 – Capturing Active and Secret Keys

VAST Identity Policy for Veeam

With the user created, specific S3 API permissions need to be given to that user so that when the user is assigned to an endpoint or a bucket it will have the appropriate settings to interoperate correctly with Veeam. The vehicle to do this is called an Identity Policy.

The following are Veeam's required S3 API permissions. The permissions below are in JSON format and can be copied and pasted directly into an Identity Policy within VAST.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:GetObjectRetention",
        "s3:PutObjectRetention",
        "s3:PutObjectLegalHold",
        "s3:GetBucketVersioning",
        "s3:GetObjectLegalHold",
        "s3:GetBucketObjectLockConfiguration",
        "s3:PutObject",
        "s3:GetObject",
        "s3:PutBucketObjectLockConfiguration",
        "s3:DeleteObject",
        "s3:GetBucketLocation"
      ],
      "Resource": "*"
    }
  ]
}
```

To create an identity policy, go to **User Management** on the left of the VAST UI and select **Identity Policies** (Figure 6).




Figure 6 - Selecting Identity Policies

In the upper right of the Identity Policies tab click Create Policy (Figure 7).




Figure 7 - Select Create Policy

Add Policy

Name *

VeeamReady

Policy

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:GetObjectRetention",
        "s3:PutObjectRetention",
        "s3:PutObjectLegalHold",
        "s3:GetBucketVersioning",
        "s3:GetObjectLegalHold",
        "s3:GetBucketObjectLockConfiguration",
        "s3:PutObject",
        "s3:GetObject",
        "s3:PutBucketObjectLockConfiguration",
        "s3:DeleteObject",
        "s3:GetBucketLocation"
      ]
    }
  ],
  "Resource": "*"
}
```

Upload JSON file

Create

Figure 8 - Creating Identity Policy

This will open the **Add Policy** window (Figure 8). Give the policy a name and copy/paste the JSON text from above into the policy and click **Create**. The policy will now show as highlighted in Figure 7.

This policy now needs to be assigned to the user created previously. To do that edit the user in the same manner during secret/access key creation. This will open the Update User window. On the right use the pull-down menu to select the identity policy that was created before (Figure 9).

Update User

User: veeamdoc-TMPHX10

Name * veeamdoc-TMPHX10 UID * 112623

Leading group Select Group

Groups Select Groups

Access key TECS091PRXT7264PMWPH Status enabled Actions

Create new key

Identity Policies VeeamReady

Update

Figure 9 - Add Identity Policy to User

This will now give explicit permissions to that user for all the S3 commands in the Identity policy.

Creating an S3 Endpoint on VAST

A good best practice for organizational purposes and ease of use with creating buckets is to create an endpoint at the root of where buckets would be placed. Once an endpoint is established the buckets that are created using the same user and secret key will be placed in an orderly manner underneath that endpoint. Buckets can be created within the VAST UI or through an S3 utility like S3 Browser®.

To create an endpoint, go to the Element Store within the VAST UI, select the Views tab and then click **Create View** as highlighted with the upper right box in Figure 10.

Element Store

Views View Policies Lifecycle Rules Quotas NFS Locks Tenants QOS Policies

Path	Alias	Share	Bucket	Bucket Owner	Policy Name	Protocols	Physical Capa...	Logical Capa...	Live Moni...	Created Time	Tenant	QOS Policy	Actions
doc	Search...	Search...	Search...	Search...	Search...	Select Pro...	Search...	Search...	Select Live...	Choose a date YYYY/MM/DD H...	Search...	Search...	
/veeamdoc			veeamdoc	veeamdoc	S3 ENDPOINT	0 Bytes	0 Bytes	No	2023/11/20 16:39:38 (...)	default			

/veeamdoc veeamdoc S3 ENDPOINT

Figure 10 - Creating an S3 Endpoint on VAST

Add View

General S3 Share-level ACL

Path *
/veeamdoc

Protocols *
S3 Endpoint

SMB share name
Type SMB share name...

NFS alias
Type NFS alias...

S3 Bucket Name *
Type S3 Bucket Name...

Policy name *
veeamdoc

QOS Policy
Select QOS Policy

Create Directory ☒

Create

Figure 11 – S3 Endpoint General Details

This will open an **Add View** wizard as shown in [Figure 11](#). Here the process is similar to bucket creation where a path is entered, a policy is given but only the Endpoint protocol is selected. (Don't forget to slide the Create Directory toggle over). On the S3 tab ([Figure 12](#)) of this wizard ensure the appropriate user is given access to create buckets. Click Create when finished. This will show the endpoint as in [Figure 10](#).

Add View

General S3 Share-level ACL

S3 Owner *
Type...

S3 Features

S3 Versioning ☐

S3 Object Lock ☐

S3 Retention Period
1d

Retention Mode
None

S3 Access Control

Bucket Creators (Users)
veeamdoc-TMPHX10@... |Type...

Anonymous access ☐

Create

S3 Access Control

Bucket Creators (Users)
veeamdoc-TMPHX10@... |Type...

Anonymous access ☐

Figure 12 – S3 Endpoint S3 Details

Again, a handy feature of having an endpoint is when using any S3 utility along with the same user it will propagate all parameters according to the endpoint and place buckets cleanly under the endpoint path.


For example, the endpoint that was just created has a path of `/veeamdoc`. A new bucket created with the name – **bucket1** will show up like this in the VAST UI: `/veeamdoc/bucket1`. This becomes especially useful when using S3 utilities and software that allows bucket creation within a workflow.

Creating an S3 Bucket on VAST

Now that an endpoint has been created, we'll create a new bucket. Again, this could be done with any S3 utility (leveraging that S3 endpoint that is setup) or directly in the VAST UI. This example will be done with VAST.

The bucket here will initially be created **without** object lock or versioning turned on. If using object lock or versioning, there is a specific workflow for when to turn it on while configuring repositories within Veeam. This is discussed in more detail in the section [Veeam and VAST S3 Repository](#).

Just as before, when creating a view, from the dashboard go to **Element Store** along the left, then **Views** tab and then click **Create** as shown in [Figure 10](#). This View is being configured as a bucket only using just the S3 Bucket protocol (Multiple protocols (NFS, S3) can be used on a View). [Figure 13](#) shows all the settings needed to create the new bucket. Since this is a new View the create directory toggle is selected.



The screenshot shows the 'Add View' dialog box with the 'General' tab selected. The 'S3' sub-tab is active. The 'Path' field contains '/veeamdoc/bucket1'. The 'Protocols' dropdown menu is set to 'S3 Bucket'. The 'SMB share name' field is empty. The 'NFS alias' field is empty. The 'S3 Bucket Name' field contains 'bucket1'. The 'Policy name' dropdown menu is set to 'veeamdoc'. The 'QOS Policy' dropdown menu is set to 'Select QOS Policy'. The 'Create Directory' toggle switch is turned on. A green 'Create' button is located at the bottom center of the dialog.

Figure 13 – S3 Bucket Configuration

The user created previously needs to be given access or ownership of this view. This is done on the S3 tab as shown in [Figure 14](#). Leave S3 Object Lock and Versioning off at this time as this example is just creating the basic bucket. Enabling object lock and versioning will be discussed later. When finished click Create from either tab.

Add View

General

S3

Share-level ACL

S3 Owner*

veeamdoc-TMPHX10@ |Type...

S3 Features

S3 Versioning

S3 Object Lock

S3 Retention Period

1d

Retention Mode

None

S3 Access Control

Bucket Creators (Users)

Type...

Anonymous access

Bucket Creators (Groups)

Type...

Create

General	S3	Share-level ACL
---------	----	-----------------

S3 Owner*

veeamdoc-TMPHX10@ [Type...]

S3 Features

S3 Versioning

S3 Object Loc

None

S3 Retention Period

1d

S3 Access Control

Anonymous access ☐

Create

Figure 14 - Adding Bucket Owner (User)

The view (bucket) now shows up in the Element Store as shown in [Figure 15](#).

The screenshot shows the Veeam Backup & Replication console interface. The 'Buckets' tab is selected, displaying a table of backup buckets. A callout box highlights the path '/veeamdoc/bucket1' in the 'Path' column of the table.

Path	Bucket	Size	Created	Last Modified	Actions
/veeamdoc	Search...	Search...	Search...	Search...	Search...
/veeamdoc	veeam...	53 B...	0 Bytes	0 Bytes	No
/veeamdoc/bucket1	bucket1	veeam...	veeam...	53 B...	0 Bytes

Figure 15 – S3 Bucket Created

Veeam and VAST S3 Repository

A backup repository is a location used by Veeam Backup & Replication jobs to store backup files, VM copies, and metadata for replicated VMs. The following sections provide the appropriate steps to configure Veeam and VAST for an immutable S3 backup repository leveraging VAST S3 Object Lock and versioning.

Create Immutable Veeam Backup Repository

Creating a backup repository is reasonably straight forward using the wizard. To start the wizard, click on Backup Infrastructure on the left panel and with **Backup Repositories** selected either right click and select **Add Backup repository** or click the **Add Repository** button (Figure 16).




Figure 16 – Start Backup Repository Wizard




Figure 17 – Select Object Storage

This opens the Add Backup Repository window (Figure 17). Here select **Object storage** and it will bring up the **Object Storage** window where S3 Compatible should be selected (Figure 18).




Figure 18 – Select S3 Compatible

The **New Object Storage Repository** window opens (Figure 19). Enter a **Name** to give the repository and click **Next**.

New Object Storage Repository

Name
Type in a name and description for this object storage repository.

Name:
VAST-objlock

Description:
Created by VASTDATA LAB\hwaller at 11/29/2023 10:05 AM.

☐ Limit concurrent tasks to: 2

Use this setting to limit the maximum number of tasks that can be processed concurrently in cases when your object storage is overloaded or cannot keep up with the number of API requests issued by multiple object storage offload tasks.

< Previous **Next >** Finish Cancel

Figure 19 – Give Name to Backup Repository

There is also an additional setting for limiting concurrent tasks (Figure 19) which was left unchecked because there is no need to limit the number of tasks or API requests to the VAST cluster. The next window is where the VAST account information is entered. Enter the appropriate Service point (**VIP Pool.DNS Delegation.Domain Name**) and ensure that the Region is **vast-1**. Select (or add to Veeam, see below) the **Credentials** that were created in the section **Creating an S3 User on VAST** and then click **Next**.

New Object Storage Repository

Account
Specify account to use for connecting to S3 compatible storage system.

Service point:
https://8vips.tmpvhx-10.vastdata.lab

Region:
vast-1

Credentials:
TECS091PRXT7264PMWPH (TMPHX10-veeamdoc, last edited: less than a day a... Add...

Manage cloud accounts

Connection mode:
Direct Choose...

Specify how object storage should be accessed and configure repository access control settings for backup agents.

< Previous **Next >** Finish Cancel

Figure 20 – Enter Service Point and Credentials

Note: Ensure the region is vast-1. If it's not there or there is a default region then it can lead to excess noise in logs.

If the credentials have not been previously added, then they'll need to be added by clicking the Add button and entering the access and secret keys that were created previously (Figure 21). It may be useful to add the VAST username to the description to keep track of multiple cloud accounts.




Figure 21 – Adding S3 User Credentials

Click next when all the fields have been filled. Veeam now reaches out to the VAST cluster with the defined S3 user credentials and scans for a list of accessible buckets. Go ahead and accept the certificate warning if it appears (For more information on SSL certificates see the section – [Root Certificate Authority](#)). When finished here and finished with Account window (Figure 20) simply click Next.

IMPORTANT!

On the **Bucket** screen (Figure 22) there is a particular workflow that needs to be followed to ensure proper interoperability. Recall that the bucket that was created previously was done so without object lock or version on – If object lock, and more specifically when versioning, is enabled at the wrong time there will be an additional slash “/” created at the end of any folder that is created and this is particularly apparent with Veeam. If a folder that has this extra slash is used to create the repository this will cause Veeam backups to fail. *VAST is looking into this issue and a fix will be forthcoming but until then simply follow these steps to ensure that backups work properly with Veeam.*




Figure 22 - Bucket Wizard

Starting with Step 1 in [Figure 22](#) click on the **Browse** button so the bucket that was created previously can be selected. The **Select Bucket** window will show all available buckets to choose from ([Figure 23](#)). Select the bucket and click OK.

Proceeding to Step 2, click on the **Browse** button for the **Folder** entry. This will open a Select Folder window ([Figure 24](#)). Click on **New Folder** to create one within the selected bucket. This creates the folder /Veeam/Archive/folder-name for every new folder that is created in that wizard. Notice that for every new folder, there should be **NO SLASH** at the end.

Note: A reminder that S3 Object Lock and Versioning on the bucket is still OFF.




Figure 23 - Select VAST Bucket




Figure 24 – Create Folder Within Bucket

Note: Notice the folder does NOT have a slash at the end.

With the bucket selected and the folder created without any slash at the end, the Bucket window should look like [Figure 22](#). Now it's time to turn on S3 Object Lock and Versioning on the VAST bucket. To do this go to the VAST UI and back to the bucket that's highlighted in [Figure 15](#). Click on the three small dots on the far right and select **Edit**. Select the S3 tab ([Figure 25](#)) and turn on both S3 Versioning and Object Lock (Selecting Object Lock will automatically turn on Versioning).




Figure 25 – Turn On VAST S3 Versioning and Object Lock

As an example of the incorrect workflow – if a folder is created now (with Object Lock and Versioning on) the Folders window will show that any new folder created will have a slash (Figure 26). Selecting the folder with a slash will cause Veeam backups to fail.

IMPORTANT!

The workaround is quite simple – do not enable Object Lock and Versioning on a bucket until after bucket and folder are configured within Veeam.




Figure 26 – Folder Created After Turning On Versioning

If folders are created with slashes, they can easily be deleted from within S3 browser as shown in Figure 27.




Figure 27 – Cleaning Up Bad Folders

With a proper bucket and folder created (Figure 22) and with S3 Object Lock and Versioning now enabled the repository configuration can move forward. Select the expiration time by checking the immutability check box and selecting a time period. Then just click next several times through the Mount Server, Review, Apply and Review windows – Stopping on the Review window (Figure 28) to ensure that everything selected looks correct.




Figure 28 – Backup Repository Review

Once the creation of the backup repository finishes it should appear as in [Figure 29](#). The S3 backup repository is now ready to be assigned to a job.




Figure 29 – S3 Backup Repository Successfully Created

Additional S3 Repository Uses

SOBR Backup Repository

With cloud based S3 it is most often categorized as cold storage in many backup and disaster recovery designs. With that in mind the Veeam feature – SOBR or Scale Out Backup Repository offers an intelligent way to use on-prem fast storage as the primary layer of a SOBR repository and then use S3 as the secondary or cold layer.

In an on-prem environment this same design can be deployed using VAST S3 where VAST, due to its performant nature, could be used in either tier of a SOBR repository. Whichever tier VAST is used in, the configuration is simplified in that only a single bucket needs to be supplied to Veeam during configuration. Adding multiple buckets (extents) gives no additional benefit or performance improvement. This is due to a virtual IP (VIP) pool that is used to access VAST as a service point. That VIP pool distributes IPs as streams as they are connected from Veeam and with each connection a different IP is given in a round robin fashion. This process of doling out multiple IPs distributes the traffic across all the front end cNodes of the VAST cluster providing maximum performance.

To help drive performance Veeam provides scalable gateways that can increase the number of streams to a backup repository. So, VAST S3 can be used, and is recommended to be the primary target backup repository but also in the secondary tier within a SOBR repository helping deliver maximum performance and simplified management. VAST makes an S3 backup repository a first-class target with no performance, security or simplicity trade-offs.

Protecting Veeam Backup Infrastructure

Protecting one's data is obviously one of the most important parts of any company but what about the infrastructure? If the backup infrastructure gets compromised rebuilding such an environment could take days not hours to rebuild again. Between accessing proper operating system versions, downloading and installing updates, installing software as well as configuring it properly could cost an immense amount of time.

With the Veeam infrastructure running on VAST it provides immense benefits for both performance and security. During a cyber security event finding the right restore point can take hours for each attempt to locate that clean instance. But with VAST performance and the use of VAST indestructible snapshots – a multitude of restore points can not only be kept but restored quickly to ensure critical systems are up and running.

Indestructible Snapshots

Indestructible snapshots are analogous to Object Lock in that once enabled they cannot be deleted (without multi-factor secure authentication with VAST customer support). By default, protection policies and their snapshots are not deleteable during the expiration period. After that period expires, then a snapshot may be deleted and if all snapshots have reached their expiration period then the protection policy may also be deleted.

With the indestructible option enabled, neither the policy nor any snapshot can be modified or deleted without interacting with VAST customer support and multiple layers of security. This creates an excellent vehicle for protecting the Veeam backup infrastructure, from the Backup and Replication server to the proxies and the gateway servers. This allows instant VM recovery so the infrastructure is ready to start restoring the backed up object-locked S3 data.

The protection policy and manual snapshot is found in the VAST UI by going to menu column on the left of the UI and selecting Data Protection and then the snapshot method of choice ([Figure 30](#)).



Figure 30 – Selecting Data Protection Method

Add Protection Policy

Policy name *
Bucket2-Indestructible

Snapshot prefix *
bucket2-

Indestructible
WARNING! This setting cannot be disabled without providing a challenge token to VAST support.

Every * 15 Period * Minutes Start at * 2023/12/11 12:00:00 Keep local copy for * 7 Period * Days Keep remote copy for * 30

Create

Figure 31 – Configuring Protection Policy

Using the Protection Policy is a typical way of establishing a systematic backup time increment for protecting a path (Figure 31). A good approach before enabling the Indestructible feature is to craft the policy and let it run for a while to confirm the proper frequency of the snapshot schedule. Once that has been determined then the Indestructible setting can be enabled because, again, once it's enabled the only way to modify the protection policy is to go through VAST customer support. Figure 32 shows a protection policy where the Indestructible feature has not yet been enabled.




Figure 32 - Protection Policy Created

Advanced Settings

The following configuration sections are finer detail concepts that could be overlooked but are an integral part of the VAST and Veeam solution.

Block Size Setting

One configuration change that can help improve performance is the block size that Veeam uses during backups. Changing this setting is done per backup job and is accessed in the advanced properties of a particular job. To find the setting during a backup job creation or editing a current job it can be found on the Storage window under the Advanced button (Figure 33).




Figure 33 - Select Advanced on Storage Window of Backup Job

On the Advanced Settings window select the Storage tab and select the setting labeled **Storage Optimization** as shown in Figure 34.

The default setting is a **1MB** block size.

VAST recommends changing this to 8MB. Normally, the tradeoff would be worse data reduction from Veeam, and some potential storage bloat since even the smallest change is recorded as 8MB. On VAST this won't have any negative impact, and all dedupe and compression missed by Veeam is picked up by VAST anyway. This is because even at Veeam's smallest block size setting of 256K, VAST has a much finer grained data reduction, as low as 16K, but on average around 32K.




Figure 34 - Selecting Storage Optimization

Note: When checking the storage optimization property, it most likely won't show the 8MB option as it requires a registry change on the Veeam server. Here is the key information in text form so the regedit entry can be created:

```
[HKEY _ LOCAL _ MACHINE\SOFTWARE\Veeam\Veeam Backup and Replication]
Value Name: UIShowLegacyBlockSize
Type: DWORD (32bit)
Value: 1
```

When finished the Windows regedit utility shows the setting as shown in [Figure 35](#).




Figure 35 - Adding UI Key to See 8MB Block Option

Even though the Veeam UI shows a 'not recommended' configuration option VAST recommends it because it provides both improved performance, data reduction, and less CPU cycles.

This enhancement was added as part of the Veeam 11a update. For more information on this registry setting and UI capability visit: <https://www.veeam.com/kb4215>. For the UI functionality to appear it is either a UI restart or Veeam services restart. It's unclear as the information is limited to this statement:

Object storage compatibility: to help object storage devices with limited scalability, you can again enable your backup jobs to use the 8MB block size to reduce the number of objects created. To make this option visible, create the `UShowLegacyBlockSize` (DWORD, 1) registry value under the 'HKLM\SOFTWARE\Veeam\Veeam Backup and Replication' key on the backup server. Keep in mind that this option significantly increases incremental backups size and requires an active full backup to be activated.

As mentioned, for the change to work, a new ACTIVE FULL is required. Doing this will double the logical capacity of backed up data because a new .vbk will be written for each VM. If the move is simply from one .vbk per job to a .vbk per VM, the timing is perfect as a new active full would be required anyway and it makes sense to get this change in at the same time.

Root Certificate Authority

There are many ways to get a certificate of authority. The following sections provide one method for creating a certificate for the VAST VMS, VAST S3 and shows how to install them within Veeam Backup and Replication. If there are questions about creating the certificate for VMS or S3 contact your IT department for proper guidance.

This process uses OpenSSL for most of the key generation – again there are many methods to achieve the same results.

VAST VMS Certificate Creation

Create a Private Key

The first step in this process is to create a private key for the VAST cluster using the following command:

```
openssl genpkey -algorithm RSA -out tmpvx20-vms-priv.pem -pkeyopt rsa_keygen_bits:2048
```

Create a certificate request

In order to use SANs (Subject Alternative Names) you need to use a .cnf file. For this VMS cert it's not necessary unless there are aliases but it's very useful for S3 SSL certs so one can be assigned per FQDN of the VIP pool. For simplicity and consistency, the CNF file is used here with no added complexity.

Below are the contents of the .cnf file used for a VMS SSL cert:

```
[ req ]
distinguished_name      = req_distinguished_name
req_extensions          = req_ext

[ req_distinguished_name ]
countryName             = US
countryName_default     = US
stateOrProvinceName     = CA
stateOrProvinceName_default = CA
localityName            = Los Gatos
localityName_default    = Los Gatos
organizationName        = VAST Data
organizationName_default = VAST Data
commonName              = tmphx20.vastdata.lab
commonName_default      = tmphx20.vastdata.lab
commonName_max          = 64

[ req_ext ]
subjectAltName          = @alt_names

[alt_names]
DNS.1                  = tmphx20.vastdata.lab
```

Use the above config file to create a Certificate Signing Request (csr) along with the private key that was created using this command:

```
openssl req -new -key tmphx20-vms-priv.pem -out tmphx20-vms-certrequest.csr -config
openssl-tmphx20-vms.cnf
```

Process Certificate Request

Now, open the Certification Authority app (this is the Domain Controller in this scenario – IT may need to assist):

1. Submit new request and point to the .csr file created previously (Figure 36):




Figure 36 – Submitting New Certificate Request

2. Go to “pending requests” in the CA app.
3. Right-click the certificate request and select **Issue** (Figure 37).




Figure 37 – Issue Certificate Request

4. Under **Issued Certificates** find the last one and select Open (Figure 38). (do NOT right-click and export as that is the wrong kind of export). On the Details tab of the **Certificate** window → **Copy to File** (Figure 39).




Figure 38 - Open Issued Certificate




Figure 39 - Copy Certificate Details to File

5. Open the file in notepad and copy the entire portion from [----Begin Certificate-----] through [----End Certificate---] and paste into VMS. To do this, go to the VAST UI and from the left menu select Settings and Certificates as shown in [Figure 40](#).




Figure 40 – Select VMS From Menu

Then, paste the certificate into the left section – appropriately labeled *Certificate* ([Figure 41](#)).




Figure 41 – Paste Certificate and Private Key

6. Now, open the .pem file (private key) and paste it into the **Key** section within VMS ([Figure 41](#)).

VAST S3 Certificate Creation

The steps to create the S3 certificate are the same as the VMS certificate but the .cnf file used is tailored towards the use of multiple VIP pools. Notice in the .cnf file contents below that the CommonName is using an asterisk in front of the domain. Also note that the first DNS entry also has a wild card as well as explicitly calling out several known VIPs that will be used.

Create a Private Key

The same private key can be used here, or a new one can be created specifically for S3 using the following command

```
openssl genpkey -algorithm RSA -out tmphx20-priv.pem -pkeyopt rsa_keygen_bits:2048
```

Create a certificate request

```
openssl req -new -key tmphx20-priv.pem -out tmphx20-certrequest.csr -config openssl-tmphx-20.cnf
```

Contents of .cnf file:

```
[ req ]
distinguished_name = req_distinguished_name
req_extensions = req_ext

[ req_distinguished_name ]
countryName = US
countryName_default = US
stateOrProvinceName = CA
stateOrProvinceName_default = CA
localityName = Los Gatos
localityName_default = Los Gatos
organizationName = VAST Data
organizationName_default = VAST Data
commonName = *.tmphx-20.vastdata.lab
commonName_default = *.tmphx-20.vastdata.lab
commonName_max = 64

[ req_ext ]
subjectAltName = @alt_names

[alt_names]
DNS.1 = *.tmphx-20.vastdata.lab
DNS.2 = 8vips.tmphx-20.vastdata.lab
DNS.3 = 4vips.tmphx-20.vastdata.lab
DNS.4 = vsphere-1.tmphx-20.vastdata.lab
DNS.5 = vsphere-2.tmphx-20.vastdata.lab
DNS.6 = vsphere-3.tmphx-20.vastdata.lab
DNS.7 = vsphere-4.tmphx-20.vastdata.lab
```

With the certificate request created now follow the same process as was shown in the section above – [Process Certificate Request](#). When pasting the certificate and private keys into the VAST UI ensure that the S3 certificate has been selected in the pull-down menu as shown in [Figure 42](#).




Figure 42 – Selecting S3 Certificate Definitions

Veeam S3 Certificates

During the creation of the S3 backup repository the VAST S3 certificate is trusted and installed into Veeam and there should be no certificate pop up warning. If the VAST S3 Certificate ever needs to be changed then here is the process for updating the certificate within Veeam S3 Backup repositories.

Windows Servers

Updating the certificate might be possible by right clicking a repository and choosing Rescan but this method wasn't tested. The solution used here was to open each S3 repository that had its certificate updated and select **Properties** ([Figure 43](#)).




Figure 43 - Selecting Properties to Edit S3 Backup Repository

Then simply click **Next** through all the screens and then **Finish** at the end (Figure 44).




Figure 44 - Finishing Edit of S3 Backup Repository

In the case of creating a new S3 repo once the VAST cluster has a proper certificate installed there should not be any pop-ups or errors. As long as the Veeam VBR server (Windows) is a member of the domain from which the certificate was issued, it will already have the root CA included as a Trusted Root Certification Authority for User, Local Computer and Services certificate stores. If the Veeam Backup and Replication server (Computer) is already part of the same domain as the VAST cluster is joined, then by default it's trusted, and no pop-up will appear. This can all be verified via the Certificates MMC snapin ([Figure 45](#)).




Figure 45 - MMC Snap-in Confirming Trusted Root Authority

Notice in [Figure 45](#) that this is only looking at Current User there are also the Service and Computer snap-ins to add and confirm.

Linux Servers

In this example Linux Rocky 9.2 is used and the root CA also needs to be trusted or certificate errors will appear. This will prevent the servers from taking on a gateway role (one per VM being backed up). To check if the cert is trusted or not, run:

```
wget https://8vips.tmphx-20.vastdata.lab
```

If either index.html is downloaded, then the root CA is trusted. If not, there will be a validity error 20 in which case the cert is untrusted. Another method is to run:

```
openssl s_client -connect 8vips.tmphx-20.vastdata.lab:443
```

For a trusted cert the output should indicate - verify errors 0. If there is any other output, then the cert is untrusted.

Assuming untrusted and the certs need to be trusted – to do this run the following on each server:

1. Using the same .pem file of the root CA (from previous steps), copy it to

```
/etc/pki/ca-trust/source/anchors
```

2. Run:

```
sudo update-ca-trust
```

Once the ca-trust is updated, you should be able to find the cert at the top of the following file:

```
/etc/pki/ca-trust/extracted/pem/tls-ca-bundle.pem.
```

Re-run the wget or OpenSSL checks and confirm it's trusted. Now backups using the S3 repo should be successful.



For more information on Universal Storage and how it can help you solve your application problems, reach out to us at hello@vastdata.com.