FALCONSTOR®

FALCONSTOR
DATA
MASTERY
PLATFORM

FalconStor Data Mastery Platform Cloud Integration Guide

Contents

FalconStor Software, Inc.
701 Brazos Street, Suite 400
Austin TX 78701 USA

Phone: 631-777-5188
Technical Support: 631-777-3332
Web site: www.falconstor.com

FalconStor Software reserves the right to make changes in the information contained in this publication without prior notice. The reader should in all cases consult FalconStor to determine whether any such changes have been made.

xxxx2018

FalconStor Data Mastery Platform Cloud Integration Guide

Overview

The cloud is a relatively new frontier for data management. Why is there so much buzz about moving data management to the cloud? The cloud can offer many benefits, including:

- **Cost Savings**: The cost savings are substantial. Because cloud services typically offer a "pay-as-you-go" model, they are a bargain when compared to the capital costs and overhead of maintaining traditional IT infrastructures. In some instances, cloud providers offer a more cost-efficient storage platform in terms of dollars per gigabyte (GB).
- **Business Continuity and Disaster Recovery (BCDR):** The cloud is ideal for a disaster recovery (DR) solution. Using cloud-service providers for DR can significantly speed up your recovery process so that your organization has minimal disruption to business operations in the event of an outage.
- **Agility:** Working in the cloud allows organizations to access data from anywhere. Having access to data from everywhere allows companies to be more agile. An agile business can be more responsive to business opportunities, can be more customer-centric, and can recover quickly from a disaster.
- **Data tiering:** The cloud can be used as an additional data tier for cold and seldom-used data. As the volume of data grows exponentially in every data center, the need for a tiering solution at different cost points provides flexibility in achieving desired total cost of ownership (TCO) targets.
- **Flexibility:** Although consuming anything as a service can eventually cost more over time, the savings related to capital and people often enable organizations to deploy and move with more flexibility to achieve faster time to revenue. This benefit typically outweighs the longer-term cost factors.

ARCHITECTURAL OVERVIEW

The following components can be deployed in a FalconStor Data Mastery Platform (FDMP) with cloud configuration:

- A FalconStor Storage Server (FSS) that provides storage-virtualization and business-continuity services for continuous availability of your data in a virtual or physical environment. The FSS can be a physical appliance or a virtual appliance. Data replication can be configured from one FSS to another. An FSS virtual appliance (FSSVA) can be deployed as a replica server in a public cloud, such as Amazon Web Services (AWS) or Microsoft Azure.
- The FalconStor Management Server (FMS) is a standalone machine that gathers and consolidates the information coming from your different storage servers into a scalable repository of services, users, and historical data. The FMS includes a web service that allows you to connect to each FSS for management and monitoring purposes.

FSS can provision storage to clients in the data path (in-band) or outside the data path (out-of-band), and it can protect the data.

- An in-band configuration positions an FSS in the data path between a client and its storage. The storage server provisions resources to the client and allows data protection and recovery services.
- An out-of-band configuration positions an FSS outside the direct data path between a client and its storage. FalconStor DiskSafe is installed on the client machine running Linux or Windows applications in order to protect its resources. DiskSafe captures block-level changes made to a protected system disk or data disk on the application server and writes changed data blocks to a “mirror” device managed by the back-end FSS. All data-protection operations – snapshot processing, journaling, and mirroring – are managed on the FSS. Data can be mirrored continuously or at regularly scheduled intervals.

Refer to the DiskSafe User Guide for details about configuring a protection policy for system or data disks.

During the configuration, you will be asked to identify involved servers and clients. If all machines are in the same virtual private network (VPN) and virtual private cloud (VPC), you can use internal IP addresses within that network. Otherwise, you must use public IP addresses.

CLOUD USE CASES

FalconStor Data Mastery Platform provides several efficient ways to restore the data of a protected resource. You can restore either to the original resource or to another resource for migration or duplication.

You also have granular recovery options; you can recover data at a file or folder level, at a device level, or at the whole system level.

Disaster recovery

You can use the cloud as your secondary copy of data for the ultimate in low-cost protection and site recovery.

Using the FSS data-replication feature, data is synchronized from the primary to the replica to ensure consistency. Under normal circumstances, clients do not have access to the replica storage, but when a disaster occurs and data on the replica is needed, the replica can be promoted and become an independent device that can be assigned to clients.

If you use FalconStor DiskSafe to protect your client machine, you will also be able to recover the whole machine.

Some client-recovery scenarios are described further in the "Recovery scenarios" section, below.

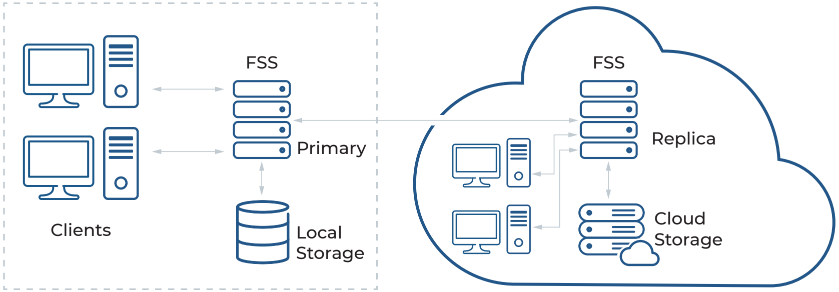

Local-site-to-cloud migration

You can migrate your storage and workspaces to the cloud in order to eliminate data center equipment expenses, maintenance requirements, and support time. The migration procedure is similar to the DR process in that the FSS data-replication feature will be used. Once replica devices are promoted and client machine images are restored on new client virtual machines (VMs) in the cloud, you can retire your local site.

Some client-recovery scenarios are described further in the "Recovery scenarios" section, below.

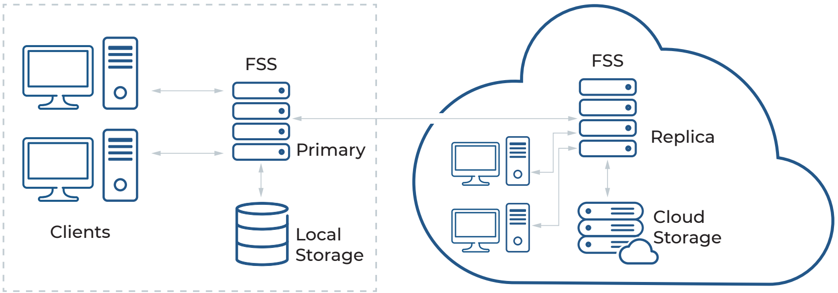

Cloud-to-cloud migration

Using the FSS data-replication feature, you can protect your data in the cloud by replicating to other cloud sites or vendors for maximum security of your data and insurance against cloud pricing changes.

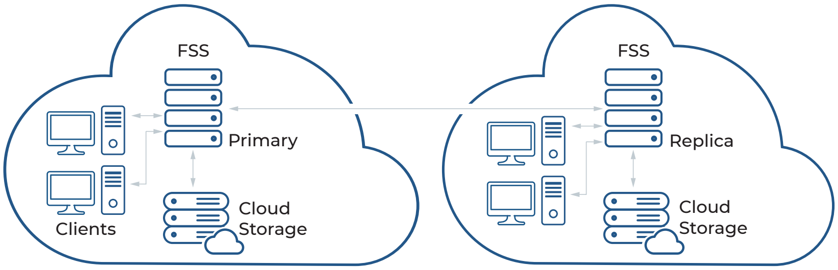

Tape-to-cloud migration

Using the FalconStor Virtual Tape Library (VTL) tape-to-object feature, you can export tapes to the cloud and stop physically moving tapes offsite for storage. The IT industry has been increasingly adopting object storage as one of the storage tiers. Virtual tapes can be migrated when they are ejected by backup software. After migration, disk space used by virtual tapes will be freed and virtual tapes will be converted to stub tapes. In this way, you can use the cloud for archiving purposes.

Refer to the FalconStor VTL User Guide for more details.

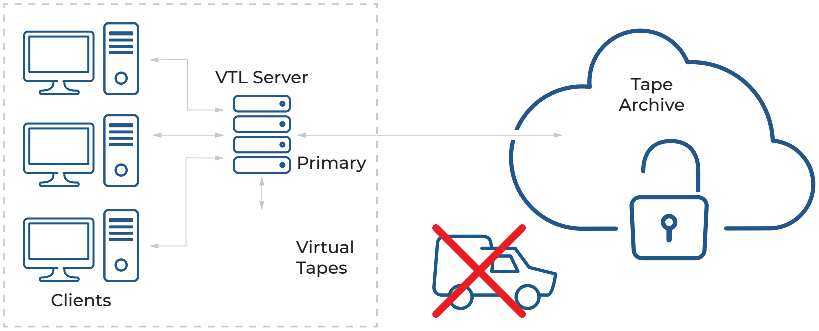

AMAZON WEB SERVICES

Amazon Web Services (AWS) is an infrastructure-as-a-service (IaaS) solution; it is a virtually limitless data center.

AWS uses Amazon Elastic Block Storage (EBS) for block-level storage volumes for use with Amazon Elastic Compute Cloud (EC2) instances in the AWS cloud. For object storage, AWS uses Amazon Simple Storage Service (S3).

There are several possible configurations for FDMP and AWS:

- Amazon EC2 – Uses an FSSVA that is deployed in the AWS public cloud to provide virtualization and DR. This configuration supports:
    - Amazon S3 storage – Provides object storage via an AWS Storage Gateway. This is typically a less expensive solution.
    - Amazon EBS – Provides block storage.
- AWS Storage Gateway – Uses an on-premise AWS Storage Gateway to provide object storage to an FSS.

MICROSOFT AZURE

Microsoft Azure is a cloud computing platform and infrastructure created by Microsoft for building, deploying, and managing applications and services.

Azure uses object storage, also called Azure Blob storage, to store unstructured data in the cloud as objects/blobs. Blob storage can store any type of text or binary data, such as a document, media file, or application installer.

By replicating snapshots from an on-premise FSS to an FSS virtual appliance in the Microsoft Azure cloud, FalconStor offers a simple DR solution.

FDMP and AWS

SUPPORTED CONFIGURATIONS

FalconStor Data Mastery Platform has three supported configurations for AWS:

- Cloud-based FSSVA and Amazon EBS
- Cloud-based FSSVA with AWS Storage Gateway and Amazon S3 object storage
- On-premise AWS Storage Gateway with cloud-based Amazon S3 object storage

Cloud-based FSSVA and Amazon EBS

In the diagram below, the FSS is replicating snapshots to the FSSVA in the AWS cloud with Amazon EBS.

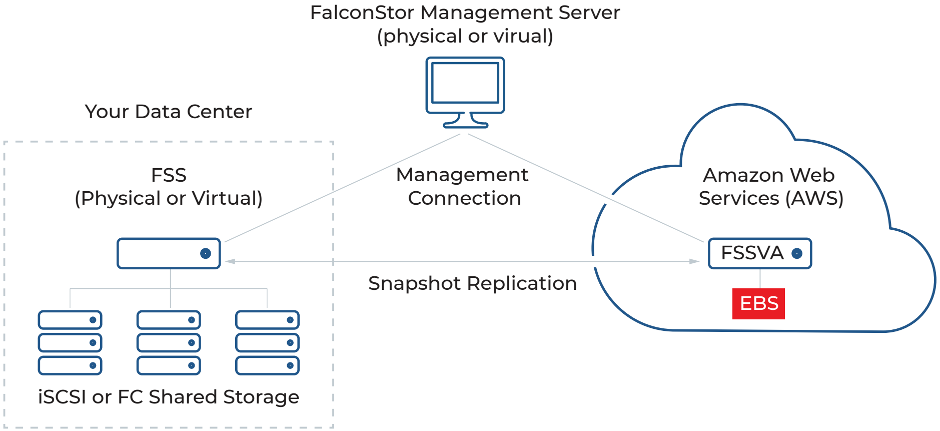

Cloud-based FSSVA with AWS Storage Gateway and Amazon S3 object storage

In the diagram below, the FSS is replicating snapshots to the FSSVA in the AWS cloud with an AWS Storage Gateway and Amazon S3 object storage.

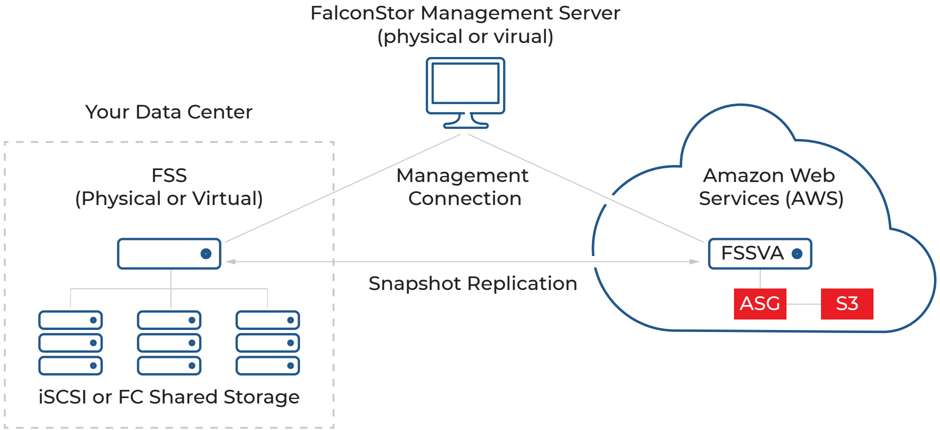

On-premise AWS Storage Gateway with cloud-based Amazon S3 object storage

In the diagram below, the FSS utilizes an on-premise AWS Storage Gateway for object storage. In this configuration, AWS provides snapshot/DR capabilities. The FSS can be physical or virtual.

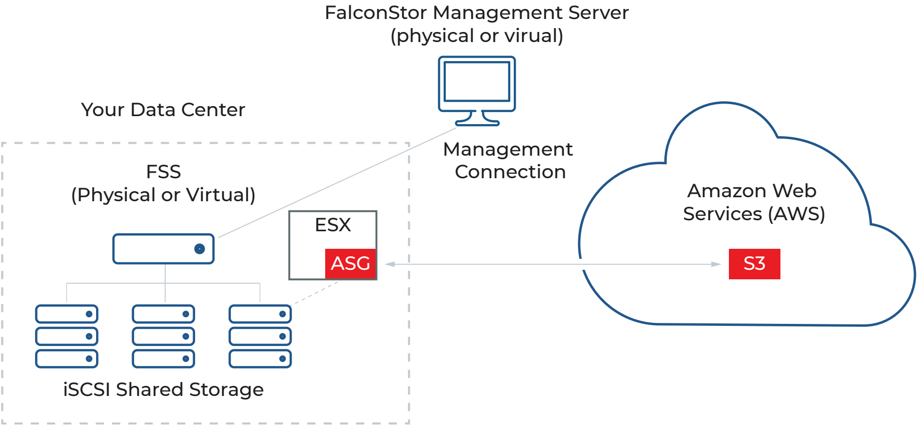

CONFIGURATION OVERVIEW

This section explains how to deploy FSSVA in Amazon EC2. It involves the following steps:

1. Create an FSS VM.
2. Add your new FSS to FMS.
3. Configure storage.
4. Assign storage to the FSS for virtualization.
5. Replicate data from the local FSS to the FSSVA in the AWS cloud.

CREATE AN FSS VM

1. Log into AWS.
2. Launch a new Amazon EC2 instance.
   a. Select All Services / Compute / EC2.

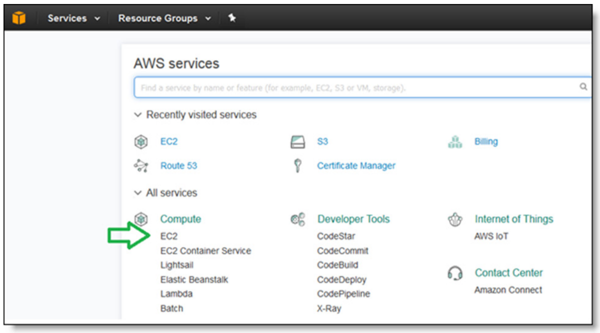

b. Select Instances / Launch Instance.

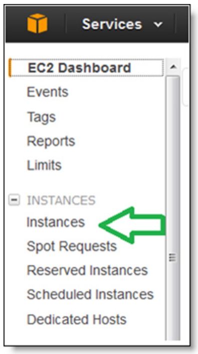

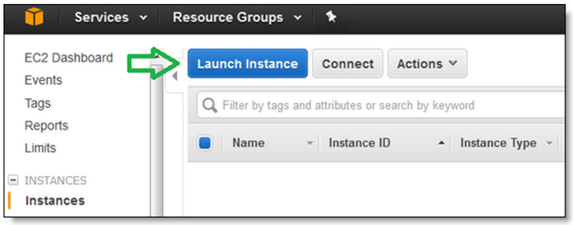

c. Select AWS Marketplace, search for FalconStor to view all offerings, and then select the FalconStor Storage Server.

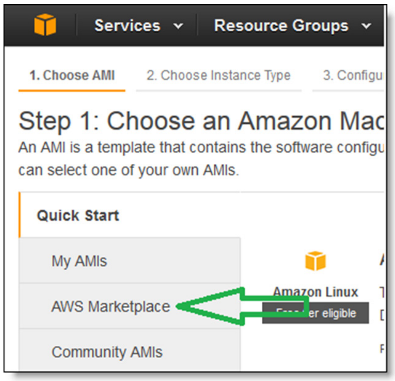

Step 1: Choose an Amazon Mac

An AMI is a template that contains the software configu
can select one of your own AMIs.

d. Choose an Instance Type that is appropriate for your environment, and then click Review and Launch.

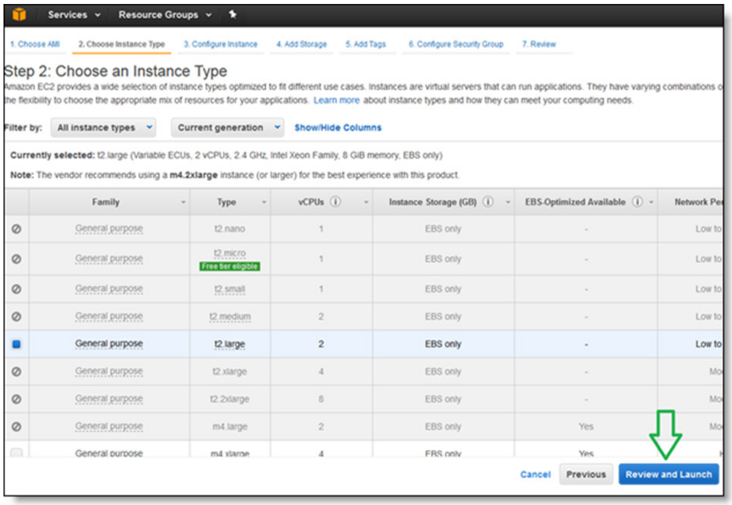

Step 2: Choose an Instance Type

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instances are virtual servers that can run applications. They have varying combinations o
he flexibility to choose the appropriate mix of resources for your applications. Learn more about instance types and how they can meet your computing needs.

Currently selected: t2.large (Variable ECUs, 2 vCPUs, 2.4 GHz, Intel Xeon Family, 8 GiB memory, EBS only)

Note: The vendor recommends using a **m4.2xlarge** instance (or larger) for the best experience with this product.

	Family	Type	vCPUs	Instance Storage (GB)	EBS-Optimized Available	Network Pe
⊘	General purpose	t2.nano	1	EBS only	-	Low to
⊘	General purpose	t2.micro Free tier eligible	1	EBS only	-	Low to
⊘	General purpose	t2.small	1	EBS only	-	Low to
⊘	General purpose	t2.medium	2	EBS only	-	Low to
■	General purpose	t2.large	2	EBS only	-	Low to
⊘	General purpose	t2.xlarge	4	EBS only	-	Mo
⊘	General purpose	t2.2xlarge	8	EBS only	-	Mo
⊘	General purpose	m4.large	2	EBS only	Yes	Mo
□	General purpose	m4.xlarge	4	EBS only	Yes	H

e. Create storage volumes.

A Root volume (70 GB) has been created; to create additional storage volumes for data, click Add New Volume. Block devices related to the new volumes will be automatically configured for FalconStor servers during the first reboot of the server.

If you add storage volumes later, you will need to manually run check_blk to configure block devices for FalconStor servers without the need to restart services.

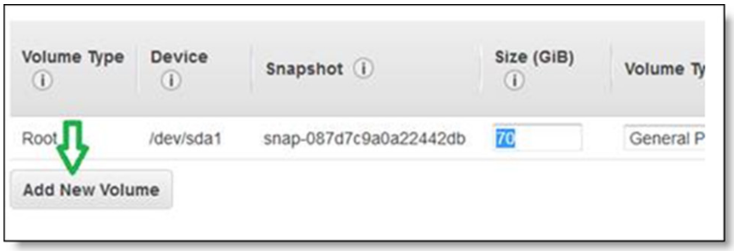

Volume Type	Device	Snapshot	Size (GiB)	Volume Ty
Root	/dev/sda1	snap-087d7c9a0a22442db	70	General P

For example, here the volume /dev/xvdf was created during installation, and the device /dev/sda is configured for FSS during installation:

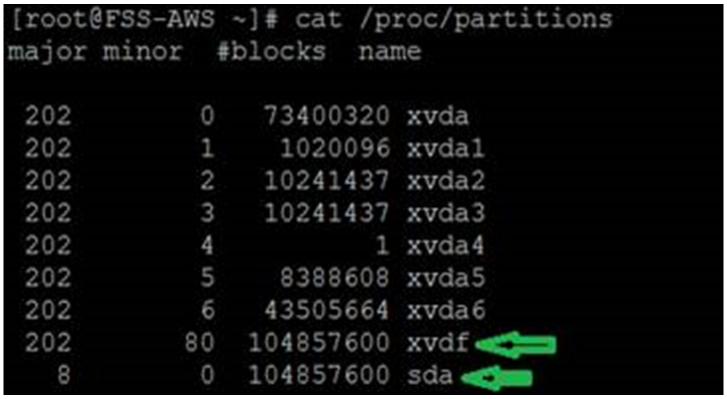

```
[root@FSS-AWS ~]# cat /proc/partitions
major minor  #blocks  name

 202        0   73400320 xvda
 202        1    1020096 xvda1
 202        2   10241437 xvda2
 202        3   10241437 xvda3
 202        4          1 xvda4
 202        5    8388608 xvda5
 202        6   43505664 xvda6
 202       80  104857600 xvdf
   8        0  104857600 sda
```

f. Click Edit security groups.

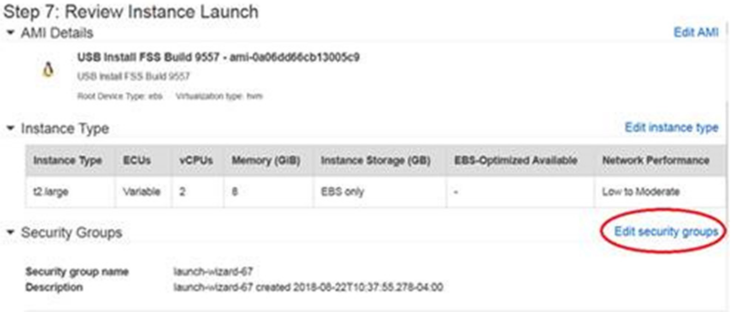

Step 7: Review Instance Launch

AMI Details

USB Install FSS Build 9557 - ami-0a06dd66cb13005c9
USB Install FSS Build 9557
Root Device Type: ebs   Virtualization type: hvm

Instance Type

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance
t2.large	Variable	2	8	EBS only	-	Low to Moderate

Security Groups

Security group name: launch-wizard-67
Description: launch-wizard-67 created 2018-08-22T10:37:55.278-04:00

g. Create a security group or select an existing group. The rules for a security group control the inbound traffic that is allowed to reach the instances associated with the security group and the outbound traffic that is allowed to leave them.

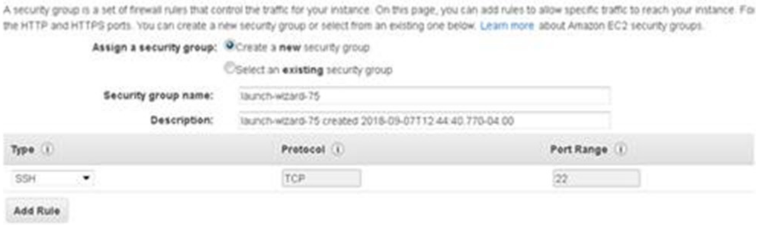

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. Learn more about Amazon EC2 security groups.

Assign a security group: ● Create a **new** security group
○ Select an **existing** security group

Security group name: launch-wizard-75
Description: launch-wizard-75 created 2018-09-07T12:44:40.770-04:00

Type	Protocol	Port Range
SSH	TCP	22

Create the group using the following steps. If you select an existing group, it must have the same rules as the one described below:

i. Click Inbound, and then click Add Rule.

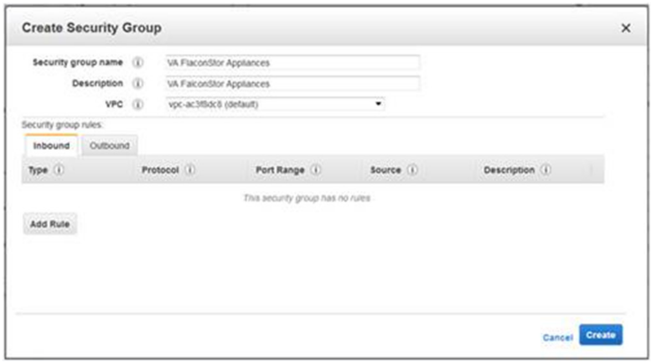

Create Security Group

Security group name: VA FlaconStor Appliances

Description: VA FalconStor Appliances

VPC: vpc-ac3f8dc8 (default)

Security group rules:

Type	Protocol	Port Range	Source	Description

This security group has no rules

ii. Add rules as shown in the following screenshot.

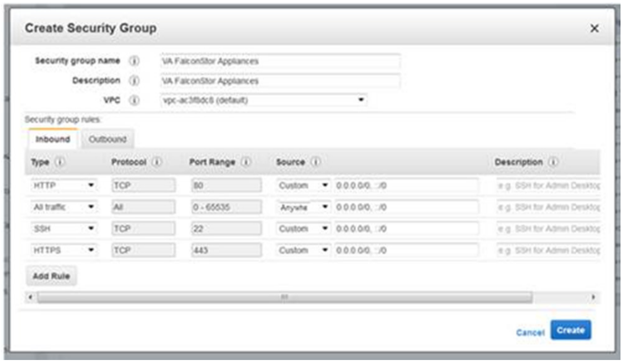

Create Security Group

Security group name: VA FalconStor Appliances

Description: VA FalconStor Appliances

VPC: vpc-ac3f8dc8 (default)

Security group rules:

Type	Protocol	Port Range	Source		Description
HTTP	TCP	80	Custom	0.0.0.0/0, ::/0	e.g. SSH for Admin Desktop
All traffic	All	0 - 65535	Anywhe	0.0.0.0/0, ::/0	e.g. SSH for Admin Desktop
SSH	TCP	22	Custom	0.0.0.0/0, ::/0	e.g. SSH for Admin Desktop
HTTPS	TCP	443	Custom	0.0.0.0/0, ::/0	e.g. SSH for Admin Desktop

iii. Click Outbound, confirm it shows the same as the following screenshot, and then click Create.

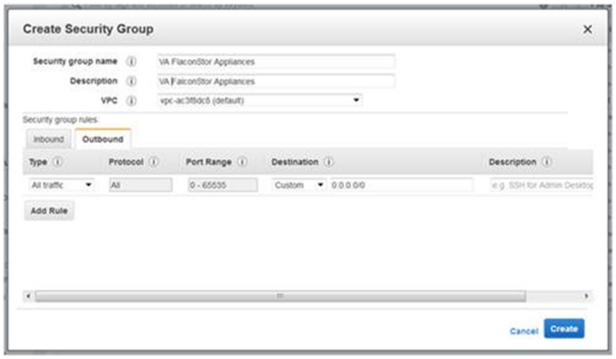

Create Security Group

Security group name: VA FlaconStor Appliances

Description: VA FalconStor Appliances

VPC: vpc-ac3f8dc8 (default)

Security group rules:

Type	Protocol	Port Range	Destination		Description
All traffic	All	0 - 65535	Custom	0.0.0.0/0	e.g. SSH for Admin Desktop

h. Click Launch when the summary is displayed.

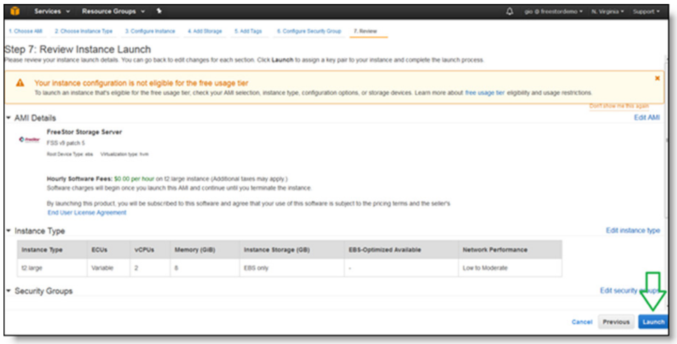

Step 7: Review Instance Launch

Please review your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

Your instance configuration is not eligible for the free usage tier

To launch an instance that's eligible for the free usage tier, check your AMI selection, instance type, configuration options, or storage devices. Learn more about free usage tier eligibility and usage restrictions.

AMI Details

FreeStor Storage Server

FSS v9 patch 5

Root Device Type: ebs Virtualization type: hvm

Hourly Software Fees: $0.00 per hour on t2.large instance (Additional taxes may apply.)

Software charges will begin once you launch this AMI and continue until you terminate the instance.

By launching this product, you will be subscribed to this software and agree that your use of this software is subject to the pricing terms and the seller's End User License Agreement

Instance Type

Instance Type	ECUs	vCPUs	Memory (GiB)	Instance Storage (GB)	EBS-Optimized Available	Network Performance
t2.large	Variable	2	8	EBS only	-	Low to Moderate

Security Groups

3. Choose your key pair.

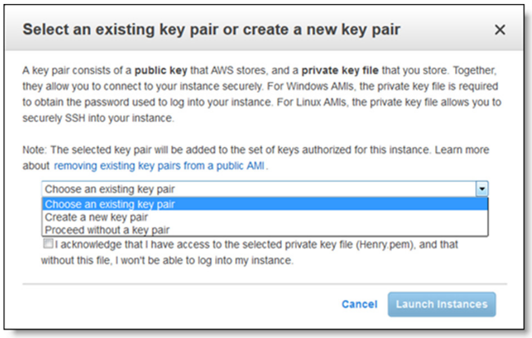

Select an existing key pair or create a new key pair

A key pair consists of a **public key** that AWS stores, and a **private key file** that you store. Together, they allow you to connect to your instance securely. For Windows AMIs, the private key file is required to obtain the password used to log into your instance. For Linux AMIs, the private key file allows you to securely SSH into your instance.

Note: The selected key pair will be added to the set of keys authorized for this instance. Learn more about removing existing key pairs from a public AMI.

Choose an existing key pair
Choose an existing key pair
Create a new key pair
Proceed without a key pair

□ I acknowledge that I have access to the selected private key file (Henry.pem), and that without this file, I won't be able to log into my instance.

If you do not have any predefined key pairs, you must create one. Be sure to download the key and keep it in a safe place as you will not be able to download it again in the future. After this is completed, click the check box to acknowledge that you have downloaded the key, and then click Launch Instances.

The VM will begin to initialize and you will be redirected to the generic screen.

4. Click View Instances to view the newly deployed instance.

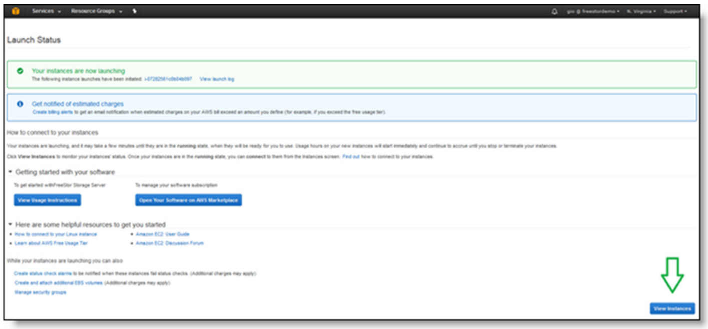

Launch Status

Your instances are now launching

The following instance launches have been initiated: i-07282561c0b04b097 View launch log

Get notified of estimated charges

Create billing alerts to get an email notification when estimated charges on your AWS bill exceed an amount you define (for example, if you exceed the free usage tier).

How to connect to your instances

Your instances are launching, and it may take a few minutes until they are in the **running** state, when they will be ready for you to use. Usage hours on your new instances will start immediately and continue to accrue until you stop or terminate your instances.

Click **View Instances** to monitor your instances' status. Once your instances are in the **running** state, you can **connect** to them from the Instances screen. Find out how to connect to your instances.

Getting started with your software

To get started withFreeStor Storage Server

To manage your software subscription

Here are some helpful resources to get you started

- How to connect to your Linux instance
- Learn about AWS Free Usage Tier
- Amazon EC2: User Guide
- Amazon EC2: Discussion Forum

While your instances are launching you can also

Create status check alarms to be notified when these instances fail status checks. (Additional charges may apply)

Create and attach additional EBS volumes (Additional charges may apply)

Manage security groups

The EC2 instance dashboard will show the new instance with a status of Initializing.

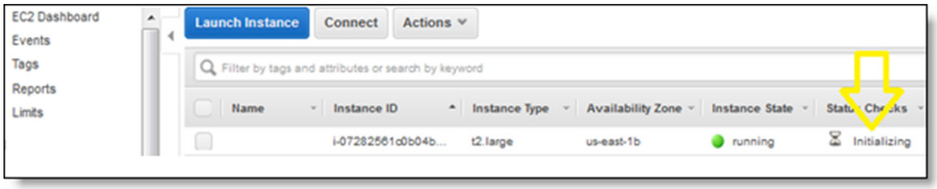

	Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks
□		i-07282561c0b04b...	t2.large	us-east-1b	● running	Initializing

5. When the instance no longer has a status of Initializing, use Secure Shell (SSH) to connect to the instance using either the DNS or IP address.

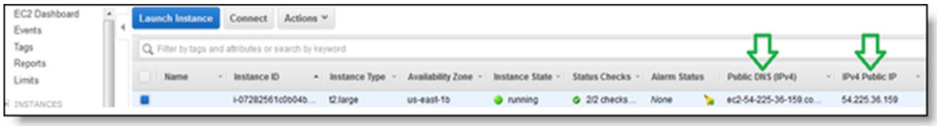

Name	Instance ID	Instance Type	Availability Zone	Instance State	Status Checks	Alarm Status	Public DNS (IPv4)	IPv4 Public IP
	i-07282561c0b04b...	t2.large	us-east-1b	running	2/2 checks...	None	ec2-54-225-36-159.co...	54.225.36.159

You must SSH with authentication using the following syntax:

`ssh -i “<KEY PAIR FILE>” cloud-user@<IP ADDRESS or DNS NAME>`

For example: `ssh -i “KeyFile.pem” cloud-user@30.43.12.45`

6. Prepare the server configuration by executing sudo into the privileged account:

```
# sudo su –
```

You will see screens similar to the following when you connect for the first time as root:

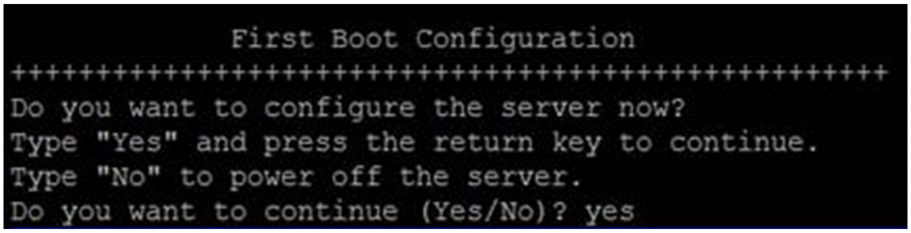

```
                First Boot Configuration
+++++++++++++++++++++++++++++++++++++++++++++++++++
Do you want to configure the server now?
Type "Yes" and press the return key to continue.
Type "No" to power off the server.
Do you want to continue (Yes/No)? yes
```

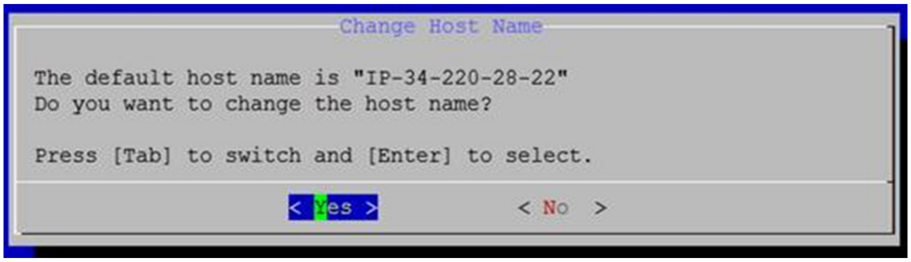

```
                    Change Host Name

 The default host name is "IP-34-220-28-22"
 Do you want to change the host name?

 Press [Tab] to switch and [Enter] to select.

             < Yes >              < No  >
```

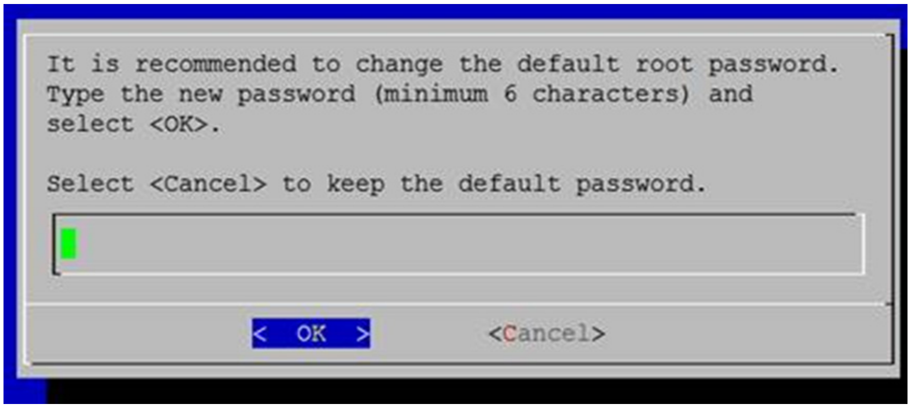

```
 It is recommended to change the default root password.
 Type the new password (minimum 6 characters) and
 select <OK>.

 Select <Cancel> to keep the default password.



            <  OK  >           <Cancel>
```

The server will reboot after changing the password. Confirm that your IP address has not changed from your Amazon EC2 console, wait until the server becomes ready, and then connect again using the cloud user account.

7. Once the instance is online again, add it to the FMS portal and continue with deployment.

The Amazon AWS command-line interface (CLI) is also installed on the FSSVA. Refer to the following link for details: http://docs.aws.amazon.com/cli/latest/userguide/installing.html

ADD YOUR NEW FSS TO FMS

1. Connect to your FMS server and log in as superadmin.
2. Select Administration from the menu bar and then select Servers.
3. Click the "+" icon.
4. Use the public IP address as the server IP address and specify the root password.

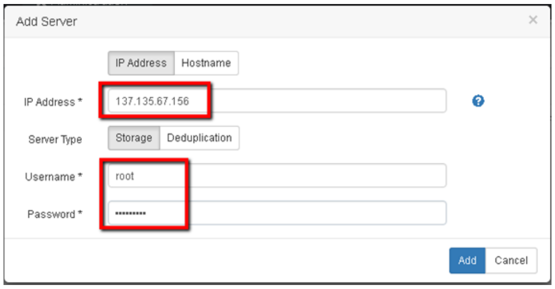

IP Address * 137.135.67.156

Server Type: Storage | Deduplication

Username * root

Password * ••••••••••

5. Click *Add.*

CONFIGURE STORAGE

This section describes how to configure storage in the following configurations:

- Cloud-based FSSVA and Amazon EBS
- Cloud-based FSSVA with AWS Storage Gateway and Amazon S3 object storage
- On-premise storage gateway with cloud-based Amazon S3 object storage

Cloud-based FSSVA and Amazon EBS

1. Log into *Amazon AWS / Services / Compute / Amazon EC2 / Amazon Elastic Block Store / Volume.*
2. Create a volume and make sure it is in the same *Availability Zone* as the FSSVA instance
3. Select and attach the volume to FSSVA by selecting *Volume / Actions / Attach Volume*
4. Once the volumes are attached, check `/var/log/messages` on the FSSVA appliance to verify that the volume label `xvdf` appears.

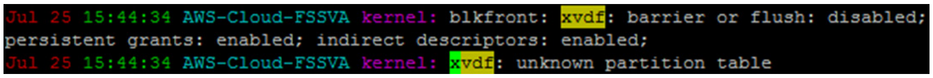

```
Jul 25 15:44:34 AWS-Cloud-FSSVA kernel: blkfront: xvdf: barrier or flush: disabled;
persistent grants: enabled; indirect descriptors: enabled;
Jul 25 15:44:34 AWS-Cloud-FSSVA kernel: xvdf: unknown partition table
```

5. Run `check_blk` to configure block devices for FalconStor servers.
6. Confirm that the `blockscsi` disk has been added to the FSSVA as shown in the example below.

```
# cat /proc/scsi/scsi
```

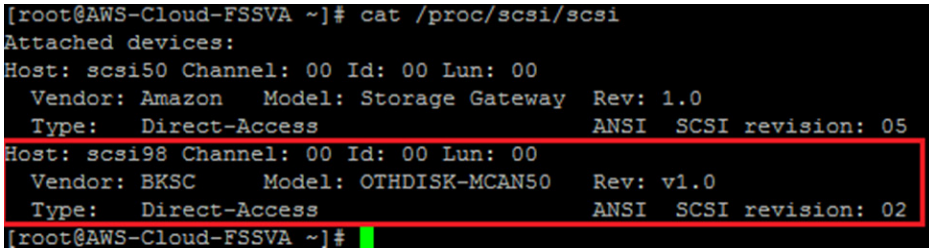

```
[root@AWS-Cloud-FSSVA ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi50 Channel: 00 Id: 00 Lun: 00
  Vendor: Amazon   Model: Storage Gateway  Rev: 1.0
  Type:   Direct-Access                    ANSI  SCSI revision: 05
Host: scsi98 Channel: 00 Id: 00 Lun: 00
  Vendor: BKSC     Model: OTHDISK-MCAN50   Rev: v1.0
  Type:   Direct-Access                    ANSI  SCSI revision: 02
[root@AWS-Cloud-FSSVA ~]#
```

Cloud-based FSSVA with AWS Storage Gateway and Amazon S3 object storage

Refer to the following Amazon link for additional installation details: http://docs.aws.amazon.com/storagegateway/latest/userguide/launch-activate-ec2-ami-common.html

1. Log into *AWS / Services / Compute / Amazon EC2 / Launch Instance / AWS Marketplace,* search for *AWS Storage Gateway,* and then click *Select.*

2. Choose an *Instance Type* that is appropriate for your environment and click *Next: Configure Instance.*
3. Leave the default setting, click *Next: Add Storage*, and then add one disk for *Upload Buffer* and one disk for *Cache Storage*. Refer to the following Amazon link for disk-sizing detail: http://docs.aws.amazon.com/storagegateway/latest/userguide/volume-gateways-architecture.html
4. Click *Review and Launch* / *Launch*, select your setting for a key pair, and then click *Launch Instances.*
5. Once the AWS Storage Gateway appliance appears in the Amazon EC2 console, locate the public IP address.
6. Activate the AWS Storage Gateway appliance by clicking *Services / Storage & Content Delivery / Storage Gateway / Deploy a new Gateway on Amazon EC2 / Gateway-Cache Volume.* Type the public IP address in the *Enter IP Address* field, and then proceed to activation. Follow the steps to complete the activation.
7. Configure the upload buffer and cache storage in the AWS Storage Gateway appliance by expanding *Volume Gateways*, clicking on your new gateway appliance, clicking the *Gateway* tab in the right panel window, clicking *Configure Local Storage*, and then selecting 20GB for the upload buffer and 25GB for cache storage. Do not configure the 10GB disk that comes with the gateway appliance as it is used as a swap disk.
8. Create a new Amazon S3 volume via the newly deployed AWS Storage Gateway for the FSSVA appliance to use.
9. Use SSH to connect to the FSSVA appliance and connect to your new volume on the AWS Storage Gateway appliance via the iscsiadm2 interactive menu.

   ```
   # iscsiadm2
   ```

   - Select (1) *Discover a target.*
   - Enter the private IP address of the AWS Storage Gateway appliance.
   - Press *Enter* to log into the iSCSI target.
   - Press *Enter* again to return to the main menu.
   - Select (9) to log into all targets, and then press *Enter.*
   - Press *Enter* to exit the *iscsiadm2* interactive menu.
10. Confirm that the Amazon S3 iSCSI disks were found and added into /proc/scsi/scsi as shown below.

    ```
    # cat /proc/scsi/scsi
    ```

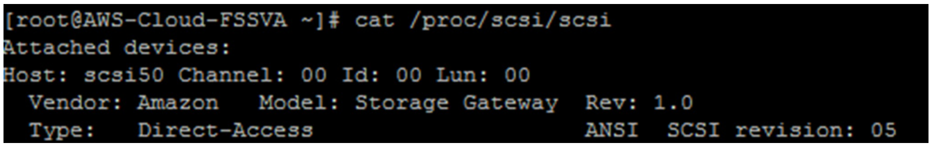

```
[root@AWS-Cloud-FSSVA ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi50 Channel: 00 Id: 00 Lun: 00
  Vendor: Amazon   Model: Storage Gateway  Rev: 1.0
  Type:   Direct-Access                    ANSI  SCSI revision: 05
```

On-premise AWS Storage Gateway with cloud-based Amazon S3 object storage

The following procedure is for a gateway-cached volume, which retains a copy of frequently accessed data locally. Refer to AWS for information about configuring gateway-stored volumes.

1. Download the gateway-cached volumes by logging into *Amazon AWS / Services / Storage & Content Delivery / Storage Gateway / Deploy a new Gateway / Gateway-Cached volumes / Continue / Continue*, and then choosing the correct virtualization platform to run the AWS Storage Gateway and saving the .ova image.
2. Deploy the AWS Storage Gateway .ova image on your local VMware environment and be sure to select the *Thick provisioned format* storage option. Also, in *VM setting / Options / VMware Tools*, check the box *Synchronize guest time with host.*

3. Allocate sufficient storage for the gateway. Change the SCSI controller type of the VM to *Paravirtual*. Refer to the following Amazon link for detailed instructions: http://docs.aws.amazon.com/storagegateway/latest/userguide/GettingStartedPLDSMain-vm-common.html
4. Power on the gateway appliance and locate the IP address. If the gateway appliance is unable to obtain an IP address via DHCP, log into the gateway appliance with the user "sguser" and the default password "sgpassword" to set the network configuration.
5. Enter the IP address for the local gateway appliance and proceed to activate it.
6. Configure the local storage to *Use for Cache Storage* and *Use for Upload Buffer*.
7. Click on your gateway appliance and create an Amazon S3 volume.
8. Configure your local FSSVA to connect to the newly created Amazon S3 volume.
   - Use SSH to connect to your local FSSVA.
   - Launch the iscsiadm interactive utility.

     ```
     # iscsiadm2
     ```
   - Select option 1 to enter the target IP address of the AWS Storage Gateway, and then press *Enter*.
   - Select option 9 to log in to all targets, and then press *Enter* again to exit.
   - Confirm that the Amazon S3 volume shows up.

     ```
     # cat /proc/scsi/scsi
     ```

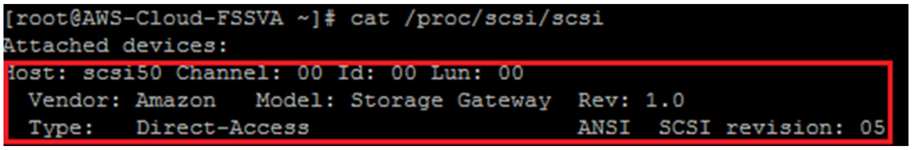

```
[root@AWS-Cloud-FSSVA ~]# cat /proc/scsi/scsi
Attached devices:
Host: scsi50 Channel: 00 Id: 00 Lun: 00
  Vendor: Amazon   Model: Storage Gateway  Rev: 1.0
  Type:   Direct-Access                    ANSI  SCSI revision: 05
```

ASSIGN STORAGE TO THE FSS FOR VIRTUALIZATION

This section explains how to assign storage from AWS to the FSSVA and virtualize it using FMS.

1. Connect to your FMS and log in as superadmin.
2. Select *Manage* from the menu bar, select your server, click the *Physical Resources* tab, and then click the *Physical Devices* tab. If a device does not show up in the list, click the *Rescan* button to rescan the physical resources.

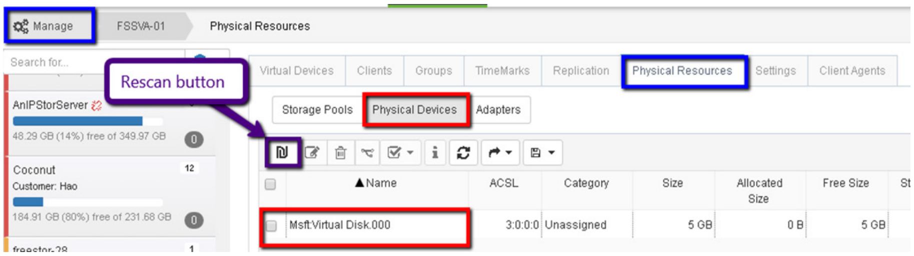

Name	ACSL	Category	Size	Allocated Size	Free Size
Msft Virtual Disk.000	3:0:0:0	Unassigned	5 GB	0 B	5 GB

3. Highlight the physical device in the list, and then click the *Edit Physical Device* icon to set its properties to *Virtual*. You can then start creating resources on it as if it was any physical storage device. Refer to the *FalconStor Data Mastery Platform User Guide* for more detail about this and to configure clients to access the virtual devices on the FSS appliance.

REPLICATE DATA FROM THE LOCAL FSS TO THE FSSVA IN THE AWS CLOUD

Replication can be based upon a defined schedule or watermark, or it can be set to occur continuously. For performance purposes and added protection, data can be compressed or encrypted during replication. FalconStor MicroScan can be used to reduce the amount of data being replicated by only transmitting changed blocks.

Refer to the *FalconStor Data Mastery Platform User Guide* for details about creating a replication policy for your virtual devices.

During the replication configuration, you will be asked to identify the replica server. If all machines are in the same VPN and VPC, you can use internal IP addresses within that network. Otherwise, you must use public IP addresses in order to connect. Also, if the FSS at your local site is behind a firewall or a NAT router, it also needs a public-facing address in order for replication to succeed.

If you have different virtualized network interface controllers (vNICs) and virtual local area networks (VLANs), and the source and replica servers use isolated networks to connect to FMS, where there is no connectivity between those networks, you will need to specify the source and replica IP addresses for the replication handshake.

FMS builds higher than 3200 and FSS patch `update-is955703` provide the option for specific IP addresses to use for handshakes between source and replica servers.

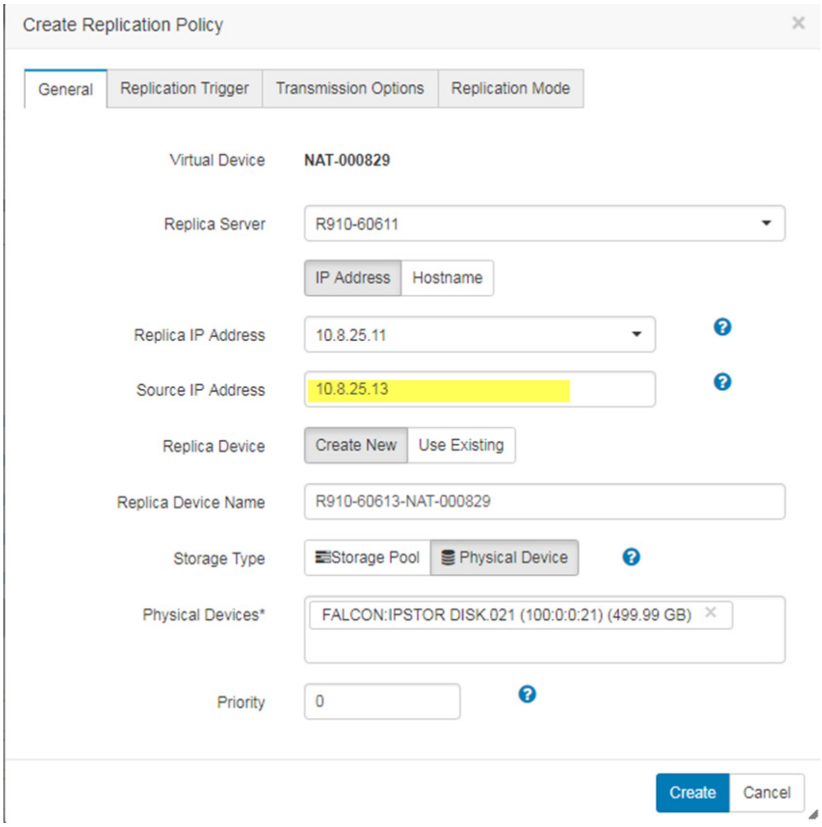

Create Replication Policy

Virtual Device: NAT-000829

Replica Server: R910-60611

IP Address | Hostname

Replica IP Address: 10.8.25.11

Source IP Address: 10.8.25.13

Replica Device: Create New | Use Existing

Replica Device Name: R910-60613-NAT-000829

Storage Type: Storage Pool | Physical Device

Physical Devices*: FALCON:IPSTOR DISK.021 (100:0:0:21) (499.99 GB)

Priority: 0

RECOVERY SCENARIOS

This section provides some recovery scenarios for a client machine.

Recover a VM protected with DiskSafe in cloud

In this scenario, you have a VM in a cloud. The machine image is protected by FalconStor DiskSafe and FSSVA. If the VM fails, you can follow the steps below to recover the VM image on another VM in the cloud:

1. Create a new VM instance with the same operating system and the same or a similar amount of CPU and memory as the VM that you are planning to recover. Be sure to choose the same boot disk size.
2. Power off the VM instance. Detach the VM boot disk and attach it to FSSVA, which will detect the boot disk as a block device.
3. On FSSVA, run the following command to see the major and minor number of the newly attached block disk device:

   ```
   # cat /proc/partitions
   ```
4. Convert the block device to an SCSI device by running the FalconStor `add_blkscsi` utility and select option 3 to convert the block device to a FalconStor blockscsi device:

   ```
   # add_blkscsi
   ```
5. Run the following command to allow the *blockscsi* device to be used as a Service-Enabled Device (SED) for layout preservation. This command requires FSS patch `update-is955702`:

   ```
   # echo forcesedallowed > /tmp/.falconstor
   ```
6. Change the category of the blockscsi device to an SED from the FDMP Portal or via the CLI.
7. From the FDMP Portal, create a copy of the TimeMark that corresponds to the DiskSafe snapshot image of the boot disk to recover: *Manage* / select a virtual device / select a TimeMark / click the *TimeMark* icon / select *Copy TimeMark*. While the copy process is occurring in the background, a new TimeMark copy virtual device appears.
8. Mirror the TimeMark copy virtual device to the SED device representing the boot disk: *Manage* / select the newly created TimeMark virtual device copy / click the *Mirror* icon / select *Create Virtual Device Mirror* / select the SED category / click *Create*.
9. Once the data mirroring is complete and the boot disk contains the snapshot image, perform the following steps for cleanup:
   - Delete the mirror.
   - Change the SED device category to *unassigned*.
   - Run the FalconStor `add_blkscsi` utility and select option 4 to remove the *blockscsi* device.
   - Perform a rescan so that the SED device shows as *offline*, and then delete the *offline entry*.

10. From AWS, detach the disk from the FSSVA, and then reattach it back to the original instance. Make sure the device is attached as /dev/sda1, not the default device ID.

12. Boot up the instance.

13. If you still want to protect the new VM with DiskSafe, check and redefine the protection policies using new IP addresses, disk IDs, and the storage server.

Recover a machine protected with DiskSafe to the cloud with AWS EC2

In this scenario, you have a physical or virtual machine in your local site that is protected by FalconStor DiskSafe and FSS. If the machine fails, you can follow the steps below to recover the local machine image on a VM in the cloud. In the following example, the steps are performed from the local FSS and local mirror device. If replication is enabled and the replica disk is available on an FSSVA in Amazon EC2, you can execute the same steps from the replica FSSVA, using the replica device ID:

1. Run the following command to expose the virtual device that is the mirror image of the machine protected by DiskSafe. This example uses virtual device ID 12 and makes it available for backup for two hours using the existing TimeMark:

   ```
   # iscli enablebackup -s 127.0.0.1 -v 12 -l 2H -n use-existing
   ```

2. Use the following command in the background to make a raw copy of the virtual device to a file for upload to Amazon S3. Make sure there is enough free local disk space, which should be at least twice the size of the raw boot disk size.

   ```
   # nohup dd if=/dev/isdev/kisdev<deviceID> of=<BootImageName.raw>
   bs=65536 &
   ```

   For example:

   ```
   # nohup dd if=/dev/isdev/kisdev12 of=myimage.raw bs=65536 &
   ```

3. For Linux machines, modify the boot parameters. Refer to the Modify boot parameters for Linux DiskSafe protected machines section of this document (on the next page) for more information.

4. Once the raw copy is complete, log into your AWS region account via the following command:

   ```
   # aws configure
   ```

5. Run the following AWS CLI command to copy the image file to the Amazon S3 bucket:

   ```
   # aws s3 cp myimage.raw s3://falconstorimages/
   ```

6. Create a containers.json file with the information below. Be sure to set the Amazon S3 key name to the name used in the previous step:

   ```
   [{
       "Description": "Boot Image",
       "Format": "raw",
       "UserBucket": {
           "S3Bucket": "falconstorimages",
           "S3Key": "myimage.raw"
       }
   }]
   ```

7. Run the following AWS CLI command to import the image:

   ```
   # aws ec2 import-image --description "Boot Image" --disk-container
   file://containers.json
   ```

8. Run the following AWS CLI command to check the status of the import in the queue. If it has completed, this image will appear in your Amazon Machine Image (AMI) account:

   ```
   # aws ec2 describe-import-image-tasks
   ```

9. Launch a VM from this image.

10. If you still want to protect the new VM with DiskSafe, check and redefine protection policies using new IP addresses, disk IDs, and the storage server.

MODIFY BOOT PARAMETERS FOR LINUX MACHINES PROTECTED WITH DISKSAFE

1. Run the following commands to mount the raw Linux boot image file on a loop device for access:

   ```
   # losetup /dev/loop0 myimage.raw
   # kpartx -a /dev/loop0
   # mount /dev/mapper/loop0p1 /mnt
   ```

2. Change the default boot menu to the one before the DiskSafe image was installed.

   ```
   # vi /mnt/grub/grub.conf (Linux 6)
   ```
   or
   ```
   # vi /boot/grub2/grub.cfg (Linux 7)
   ```

3. Unmount the boot image.

   ```
   # umount /mnt
   ```

4. The root partition becomes the Logical Volume Manager (LVM) partition after DiskSafe is installed. Run the following LVM commands to mount the partition:

   ```
   # vgscan
     Reading volume groups from cache.
     Found volume group "vg_rh6u5" using metadata type lvm2
   # vgchange -ay vg_rh6u5
     2 logical volume(s) in volume group "vg_rh6u5" now active
   # mount /dev/vg_rh6u5/lv_root /mnt
   # vi /mnt/etc/fstab
   ```

   Remove references to `disksafe/` in `fstab.`

   ```
   # vi /mnt/etc/lvm/lvm.conf
   ```

   Remove references to `disksafe/` in the filter section of `lvm.conf`:

   ```
   “filter=[“a|^/dev/disksafe/sda.*$|”,”r|.*|”]”
   # umount /mnt
   # vgchange -an vg_rh6u5
     0 logical volume(s) in volume group "vg_rh6u5" now active
   # kpartx -d /dev/loop0
   # losetup -d /dev/loop0
   # vgscan
     Reading volume groups from cache.
   ```

FDMP and Microsoft Azure

The diagram below illustrates a simple DR solution using FSS with the Microsoft Azure cloud. In this case, the FSS is replicating snapshots to the FSSVA in the Microsoft Azure cloud using page blob storage (block storage).

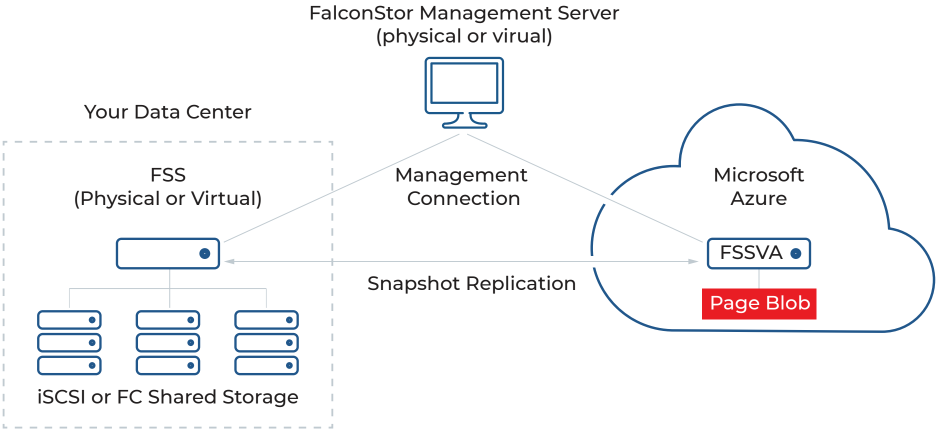

DEPLOY AN FSS FROM THE AZURE MARKETPLACE

Note: You must have a Microsoft Azure account and knowledge about administrating Microsoft Azure resources to complete these steps.

1. Log in to portal.azure.com.
2. Select Virtual Machines from the left pane.
3. Click Add to access the Azure Marketplace offerings.
4. Type FalconStor in the search bar, and then click the FalconStor Storage Server icon from the results.
5. Click Create.
6. When prompted, provide a username and password for login, along with the Azure resources that are to be allocated for the new VM being created.

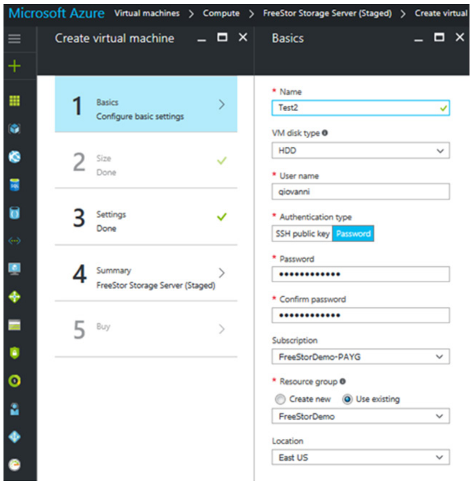
Microsoft Azure
Virtual machines > Compute > FreeStor Storage Server (Staged) > Create virtual
Create virtual machine
Basics

1 Basics — Configure basic settings
2 Size — Done ✓
3 Settings — Done ✓
4 Summary — FreeStor Storage Server (Staged)
5 Buy

* Name: Test2 ✓

VM disk type: HDD

* User name: giovanni

* Authentication type: SSH public key [✓ Password]

* Password: ••••••••••••

* Confirm password: ••••••••••••

Subscription: FreeStorDemo-PAYG

* Resource group: ○ Create new ● Use existing — FreeStorDemo

Location: East US

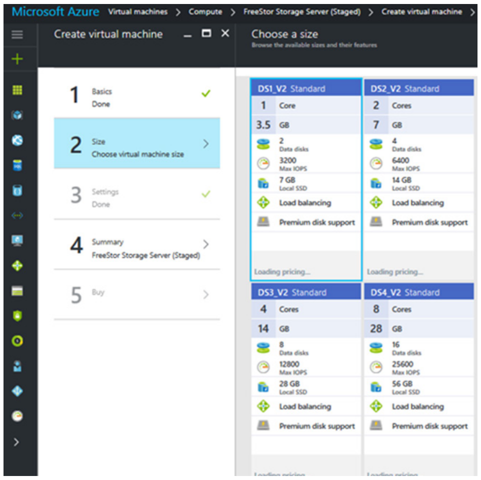
Microsoft Azure
Virtual machines > Compute > FreeStor Storage Server (Staged) > Create virtual machine >
Create virtual machine
Choose a size
Browse the available sizes and their features

1 Basics — Done ✓
2 Size — Choose virtual machine size
3 Settings — Done ✓
4 Summary — FreeStor Storage Server (Staged)
5 Buy

DS1_V2 Standard	DS2_V2 Standard	DS3_V2 Standard	DS4_V2 Standard
1 Core	2 Cores	4 Cores	8 Cores
3.5 GB	7 GB	14 GB	28 GB
2 Data disks	4 Data disks	8 Data disks	16 Data disks
3200 Max IOPS	6400 Max IOPS	12800 Max IOPS	25600 Max IOPS
7 GB Local SSD	14 GB Local SSD	28 GB Local SSD	56 GB Local SSD
Load balancing	Load balancing	Load balancing	Load balancing
Premium disk support	Premium disk support	Premium disk support	Premium disk support
Loading pricing...	Loading pricing...	Loading pricing...	Loading pricing...

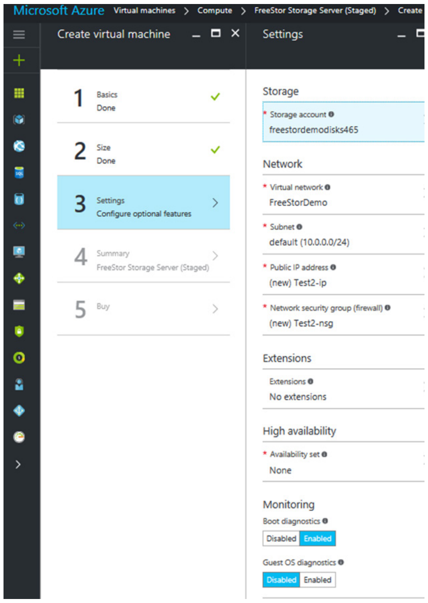

Microsoft Azure
Virtual machines > Compute > FreeStor Storage Server (Staged) > Create
Create virtual machine
Settings

1 Basics Done ✓

2 Size Done ✓

3 Settings Configure optional features

4 Summary FreeStor Storage Server (Staged)

5 Buy

Storage

* Storage account: freestordemodisks465

Network

* Virtual network: FreeStorDemo

* Subnet: default (10.0.0.0/24)

* Public IP address: (new) Test2-ip

* Network security group (firewall): (new) Test2-nsg

Extensions

Extensions: No extensions

High availability

* Availability set: None

Monitoring

Boot diagnostics: Disabled [✓ Enabled]

Guest OS diagnostics: [✓ Disabled] Enabled

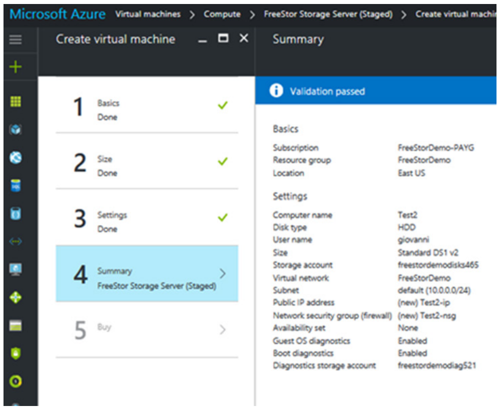

Microsoft Azure
Virtual machines > Compute > FreeStor Storage Server (Staged) > Create virtual machi
Create virtual machine
Summary

1 Basics Done ✓

2 Size Done ✓

3 Settings Done ✓

4 Summary FreeStor Storage Server (Staged)

5 Buy

Validation passed

Basics

Subscription	FreeStorDemo-PAYG
Resource group	FreeStorDemo
Location	East US

Settings

Computer name	Test2
Disk type	HDD
User name	giovanni
Size	Standard DS1 v2
Storage account	freestordemodisks465
Virtual network	FreeStorDemo
Subnet	default (10.0.0.0/24)
Public IP address	(new) Test2-ip
Network security group (firewall)	(new) Test2-nsg
Availability set	None
Guest OS diagnostics	Enabled
Boot diagnostics	Enabled
Diagnostics storage account	freestordemodiag521

On the final step, prior to the VM being deployed, Microsoft provides an estimate of how much it will cost per hour to run the VM with the previously selected resources. FalconStor is not collecting a fee during this process, but a license is required to manage the system. For additional information about this, contact salesinfo@falconstor.com.

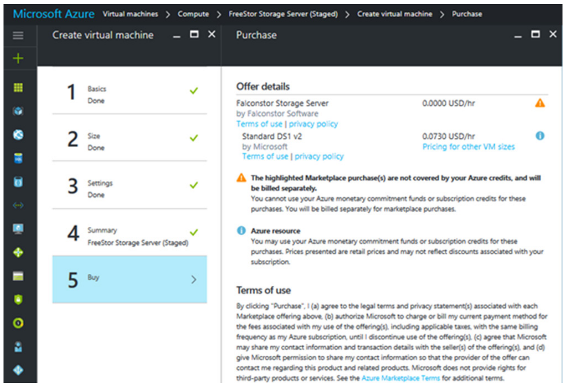

Microsoft Azure
Virtual machines > Compute > FreeStor Storage Server (Staged) > Create virtual machine > Purchase
Create virtual machine
Purchase

1 Basics Done ✓

2 Size Done ✓

3 Settings Done ✓

4 Summary FreeStor Storage Server (Staged) ✓

5 Buy

Offer details

Falconstor Storage Server by Falconstor Software Terms of use | privacy policy	0.0000 USD/hr
Standard DS1 v2 by Microsoft Terms of use | privacy policy	0.0730 USD/hr Pricing for other VM sizes

The highlighted Marketplace purchase(s) are not covered by your Azure credits, and will be billed separately.
You cannot use your Azure monetary commitment funds or subscription credits for these purchases. You will be billed separately for marketplace purchases.

Azure resource
You may use your Azure monetary commitment funds or subscription credits for these purchases. Prices presented are retail prices and may not reflect discounts associated with your subscription.

Terms of use

By clicking "Purchase", I (a) agree to the legal terms and privacy statement(s) associated with each Marketplace offering above, (b) authorize Microsoft to charge or bill my current payment method for the fees associated with my use of the offering(s), including applicable taxes, with the same billing frequency as my Azure subscription, until I discontinue use of the offering(s), (c) agree that Microsoft may share my contact information and transaction details with the seller(s) of the offering(s), and (d) give Microsoft permission to share my contact information so that the provider of the offer can contact me regarding this product and related products. Microsoft does not provide rights for third-party products or services. See the Azure Marketplace Terms for additional terms.

After clicking Purchase on the final screen of the wizard, you will be returned to the dashboard and an icon will be displayed to show that the new VM is being deployed:

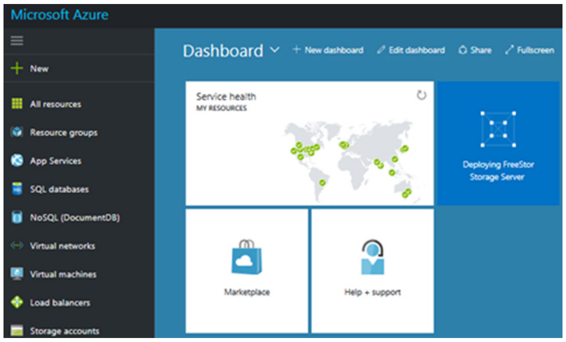

Once the VM has been successfully deployed, you will be redirected to the newly deployed VM.

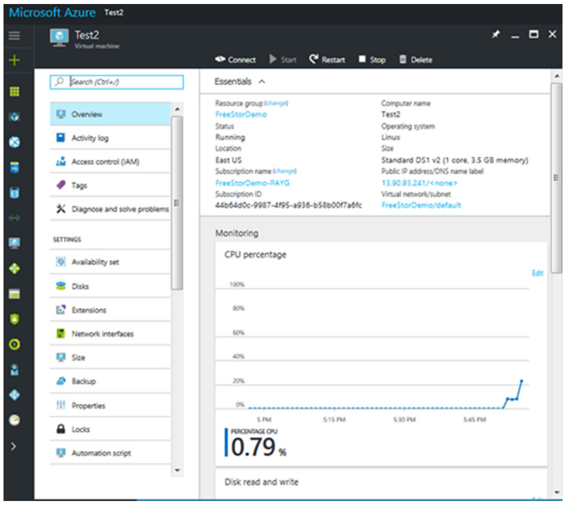

Resource group (change): FreeStorDemo
Status: Running
Location: East US
Subscription name (change): FreeStorDemo-PAYG
Subscription ID: 44b64d0c-9987-4f95-a936-b58b00f7a6fc
Computer name: Test2
Operating system: Linux
Size: Standard DS1 v2 (1 core, 3.5 GB memory)
Public IP address/DNS name label: 13.90.93.241/<none>
Virtual network/subnet: FreeStorDemo/default

7. Enable the root account (it is disabled by Azure), and run ipstor setup to populate ipstor.conf. To do this, use SSH to connect to the dynamically assigned public IP address provided by Azure.

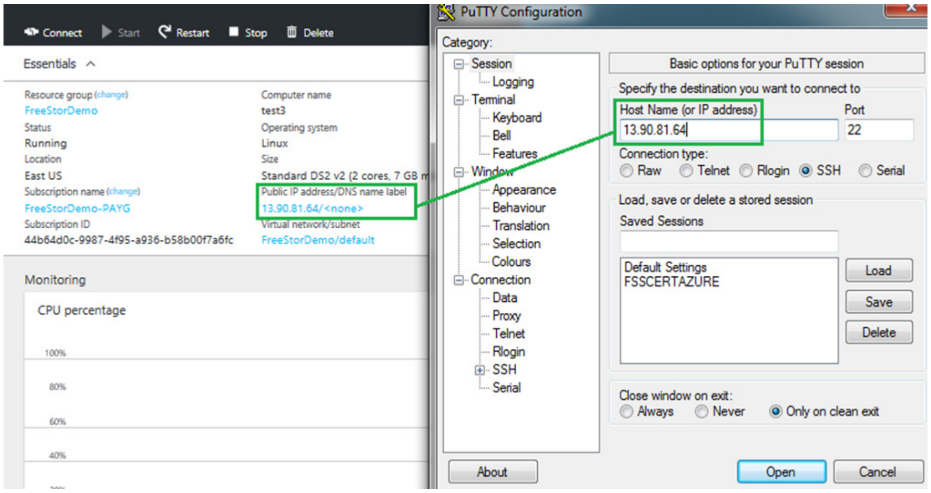

Resource group (change): FreeStorDemo
Status: Running
Location: East US
Subscription name (change): FreeStorDemo-PAYG
Subscription ID: 44b64d0c-9987-4f95-a936-b58b00f7a6fc
Computer name: test3
Operating system: Linux
Size: Standard DS2 v2 (2 cores, 7 GB m
Public IP address/DNS name label: 13.90.81.64/<none>
Virtual network/subnet: FreeStorDemo/default

Host Name (or IP address): 13.90.81.64
Port: 22
Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial
Close window on exit: ○ Always ○ Never ● Only on clean exit

Note: You must use the credentials supplied during step 6 (VM creation wizard) to log in.

8. After logging in, perform the following actions:

   a. Use “Su” to become a super user.
   [testuser@testserver ~]$ ***sudo su –***

   b. Create a password for the root account:
   [root@testserver ~]# ***passwd root***
   Changing password for user root.
   New password:
   Retype new password:
   passwd: all authentication tokens updated successfully.

   c. Stop FSS services.
   [root@testserver ~]# ipstor stop
   FalconStor IPStor Server version 9.00 (Build 9048)
   All virtual devices are going to be detached.

```
We recommend stopping all IPStor client services prior to shutdown.
Otherwise, data loss may occur.
Are you sure you want to continue? (y/n) [n]: y
```

d. Delete the following default directory:

```
[root@testserver ~]# rm -fr $ISHOME/etc/$HOSTNAME
```

e. Rename the following existing directory:

```
[root@testserver ~]# mv $ISHOME/etc/localhost $ISHOME/etc/$HOSTNAME
```

f. Create the base configuration:

```
[root@testserver ~]# ipstor setup
FalconStor IPStor Server version 9.00 (Build 9048)
Starting IPStor Block Device Module                    [  OK  ]
FalconStor IPStor Server v9.00 - (Build 9048) Setup
Supported SCSI device(s) found: 2
FalconStor IPStor Server version 9.00 (Build 9048)
Stopping IPStor Block Device Module                    [  OK  ]
```

g. Start FSS services:

```
[root@testserver ~]# ipstor start
```

9. Install the Azure CLI for Linux on the VM.
10. Add your newly deployed FSS to FMS and manage as needed. Refer to the following section of this document for more information.

ADD YOUR NEW FSS TO FMS

1. Connect to your FMS server and log in as *superadmin.*
2. Select *Administration* from the menu bar, and then select *Servers.*
3. Click the "+" icon.
4. Use the public IP address provided to the VM from Azure as the server IP address and specify the root password.

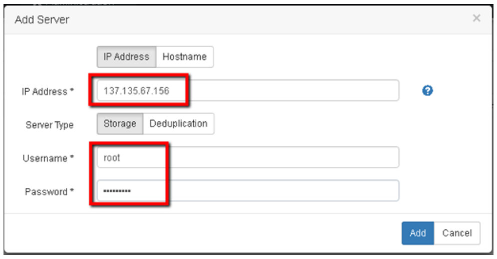

Add Server

IP Address * 137.135.67.156

Server Type: Storage

Username * root

Password * •••••••••

5. Click *Add.*

ASSIGN STORAGE TO THE FSS FOR VIRTUALIZATION

This section explains how to assign page blob storage (block storage) from the Azure cloud to the FSSVA and how to virtualize it using the FMS portal.

1. From Azure Resource Manager, click on your VM, click Disks, and then click Attach New.
2. Use the default settings. The size of this device can be up to 1023GB (1TB). If you need a smaller device, change the size. If you need a larger device, you can allocate multiple disks and FSS will virtualize them into one larger device.
3. Connect to your FMS server and log in as superadmin.

4. Select Manage from the menu bar, select your server, click the Physical Resources tab, and then click the Physical Devices tab. If a device does not show up in the list, click the Rescan button to rescan the physical resources.

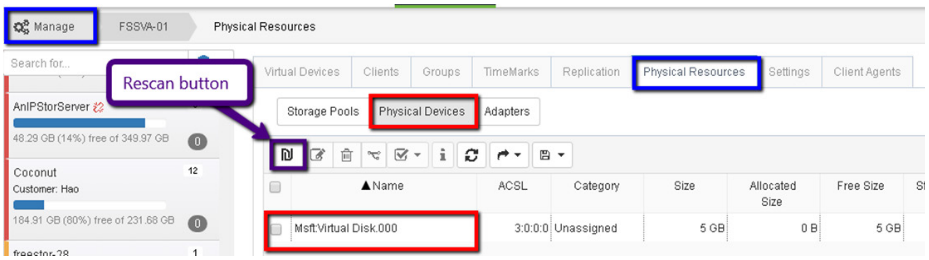

Rescan button

☐	▲Name	ACSL	Category	Size	Allocated Size	Free Size	St
☐	Msft Virtual Disk.000	3:0:0:0	Unassigned	5 GB	0 B	5 GB	

5. Highlight the physical device in the list, and then click the Edit Physical Device icon to set its properties to Virtual. You can then start creating resources on it as if it was any physical storage device. Refer to the FalconStor Data Mastery Platform User Guide for more details about this and to configure clients to access the virtual devices provided by FSS.

REPLICATE DATA FROM THE LOCAL FSS TO THE FSSVA IN THE AZURE CLOUD

FSS can provision storage to clients in the data path (in-band) or outside the data path (out-of-band), and it can protect the data.

An in-band configuration positions an FSS in the data path between a client and its storage. The storage server provisions resources to the client and allows data-protection and recovery services. In an in-band configuration, data replication can be configured from one FSS to another.

An out-of-band configuration positions an FSS outside the direct data path between a client and its storage. FalconStor DiskSafe is installed on the client machine running Linux or Windows applications in order to protect its resources. DiskSafe captures block-level changes made to a protected system disk or data disk on the application server and writes changed data blocks to a “mirror” device managed by the back-end FSS. All data protection operations – snapshot processing, journaling, and mirroring – are managed on the FSS. Data can be mirrored continuously or at regularly scheduled intervals.

Refer to the *FalconStor Data Mastery Platform User Guide* for details about creating a replication policy for your virtual devices. During the configuration, you will be asked to identify the replica (target) server. You must use the public IP address provided by Azure in order to connect. Also, if the FSS at your local site is behind a firewall, it also needs a public-facing address in order for replication to succeed.

Refer to the *DiskSafe User Guide* for details about configuring a protection policy for system or data disks.

RECOVERY SCENARIOS

This section provides some recovery scenarios for a VM.

Recover a Windows VM protected with DiskSafe from a local VMware configuration to the Azure cloud

On the target FSS, prepare the DiskSafe boot image for the DiskSafe client that you want to recover to Azure.

Prerequisite: Install the Azure CLI on the target FSSVA. You must also have the Azure PowerShell and Azure CLI installed on a Window Server environment that supports Windows PowerShell commands.

1. Identify the VID of the DiskSafe operating system mirror disk and enable raw disk backup on this VID.

   ```
   # iscli enablebackup -s 127.0.0.1 -v <vid_ds_resource> -n use-existing -r <RawBootDiskName> -l 2H
   ```

   For example:

   ```
   # iscli enablebackup -s 127.0.0.1 -v 2 -n use-existing -r BootDisk1 -l 2H
   ```

2. Use the `dd` command to dump the mirror disk into a file. Make sure there is enough free local disk space, which should be at least twice the size of the raw boot disk size.

   ```
   # nohup dd if=/dev/isdev/kisdev<RawBootDiskName> of=<RawBootDiskName.raw> bs=65536 &
   ```

   For example:

   ```
   # nohup dd if=/dev/isdev/BookDisk1 of=BootDisk1.raw bs=65536 &
   ```

3. Disable the `backupenabler` once `dd` has completed.

   ```
   # iscli disablebackup -s 127.0.0.1 -v <vid_ds_resource>
   ```

4. If the DiskSafe boot image you are attempting to recover to Azure is a Linux VM, complete the steps in the Boot a Linux disk image section of this paper. If the image is Windows VM, skip this step and continue with step 5.

5. Use the third-party q*emu-img* utility to convert the raw disk into a *vhd* format.

   ```
   # nohup qemu-img convert -f raw -o subformat=fixed < RawBootDiskName.raw> -O vpc
   <RawBootDiskName.vhd> &
   ```

6. From your FSSVA, log into your Azure account and select your subscription.

   ```
   # azure login
   # azure account set <Azure-sub-2-Name>
   ```

7. Upload the converted *vhd* file into your Azure storage account with *nohup*, and then run it in the background.

   ```
   # nohup azure storage blob upload <local file to upload> <destination container name> <destination blob file name> -a <accountName2> -k <accountKey2> &
   ```

8. Once the upload is completed, log into your Azure account and select your subscription ID via a Windows PowerShell command.

   ```
   # login-azurermaccoun
   # select-azurermsubscriptionID “subscription ID xxxxxxx”
   ```

9. Run the following commands from your Windows Server environment with PowerShell installed to create and launch the VM.

```
$pipName = "<Give a public IP name for the DS protected VM that you
are recovering>"
$rgName = "<Enter your Resource Group Name>"
$vnetName = "<Enter your virtual Network Name>"
$location = "<Enter the location of the VM to be running>"
$nicname = "<Give your network interface name>"
$pip = New-AzureRmPublicIpAddress -Name $pipName -ResourceGroupName
$rgName -Location $location -AllocationMethod Dynamic
$vnet = Get-AzureRmVirtualNetwork -Name "FreeStorDemo"
-ResourceGroupName FreeStorDemo
$nic = New-AzureRmNetworkInterface -Name $nicname -ResourceGroupName
$rgName -Location $location -SubnetId $vnet.Subnets[0].Id
-PublicIpAddressId $pip.Id
$cred = Get-Credential
$storageAccName = "<Enter your Azure storage account name>"
$storageAcc = Get-AzureRmStorageAccount -ResourceGroupName $rgName
-AccountName $storageAccName
$vmName = "<Give a name for your DS Virtual Machine>"
$vmConfig = New-AzureRmVMConfig -VMName $vmName -VMSize "<VM Size>"
$osDiskUri = "<Put in the https path of the vhd disk that you just
uploaded>"
$vm = Add-AzureRmVMNetworkInterface -VM $vmConfig -Id $nic.Id
$vm = Set-AzureRmVMOSDisk -VM $vm -Name $vmName -VhdUri $osDiskUri
-CreateOption "Attach" -Windows
New-AzureRmVM -ResourceGroupName $rgName -Location "East US" -VM $vm
```

Example of the above commands:

```
PS C:\Users\Administrator> $pipName = "dsbootvm2"
PS C:\Users\Administrator> $rgName = "MyResourceGroupname"
PS C:\Users\Administrator> $vnetName = "MyVMNetworkName"
PS C:\Users\Administrator> $location = "eastus"
PS C:\Users\Administrator> $nicname = "dsbootvm2-nic"
PS C:\Users\Administrator> $pip = New-AzureRmPublicIpAddress
-Name $pipName -ResourceGroupName $rgName -Location $location
-AllocationMethod Dynamic
PS C:\Users\Administrator> $vnet = Get-AzureRmVirtualNetwork -Name
"FreeStorDemo" -ResourceGroupName FreeStorDemo
PS C:\Users\Administrator> $nic = New-AzureRmNetworkInterface -Name
$nicname -ResourceGroupName $rgName -Location $location -SubnetId
$vnet.Subnets[0].Id -PublicIpAddressId $pip.Id
PS C:\Users\Administrator> $cred = Get-Credential
PS C:\Users\Administrator> $storageAccName =
"demostorageaccountxxxx"
PS C:\Users\Administrator> $storageAcc = Get-AzureRmStorageAccount
-ResourceGroupName $rgName -AccountName $storageAccName
```

```
PS C:\Users\Administrator> $vmName = "dsbootvm2"
PS C:\Users\Administrator> $vmConfig = New-AzureRmVMConfig -VMName
$vmName -VMSize "Standard_A1"
PS C:\Users\Administrator> $osDiskUri = https:// demostorageac-
countxxxx.blob.core.windows.net/fss/RawBootDiskName.vhd
PS C:\Users\Administrator> $vm = Add-AzureRmVMNetworkInterface
-VM $vmConfig -Id $nic.IdPS C:\Users\Administrator> $vm = Set-
AzureRmVMOSDisk -VM $vm -Name $vmName -VhdUri $osDiskUri
-CreateOption "Attach" -Windows
PS C:\Users\Administrator> New-AzureRmVM -ResourceGroupName $rgName
-Location "East US" -VM $vm
```

10. Refresh your Azure browser. You should see a new VM is being created. Azure assigns a new DHCP IP for the newly created VM. You can use Microsoft Remote Desktop Protocol (RDP) to access it via the new DHCP IP.

    If you have data disks used by the applications within this new VM, you can create TimeViews or a TimeView copy and then assign the virtual device to the new VM via iSCSI protocol. Refer to the *FalconStor Data Mastery Platform User Guide* for details on how to create TimeViews and TimeView copies, in addition to how to configure an iSCSI client and assign virtual devices.

BOOT A LINUX DISK IMAGE

1. Convert the raw device to a loopback device on the target FSSVA.

   ```
   # losetup /dev/loop0 linux6u5azure01.raw
   ```

2. Use the Linux kpartx utility to read partition tables on the specified device and create device maps over the partition segments detected.

   ```
   # kpartx -a /dev/loop0
   ```

3. Mount the boot partition to /mnt.

   ```
   # mount /dev/mapper/loop0p1 /mnt
   ```

4. Edit grub.conf and change the default boot image to the one from before DiskSafe was installed.

   ```
   # vi /mnt/grub/grub.conf
   ```

5. Umount the boot partition.

   ```
   # umount /mnt
   ```

6. Use vgscan to scan for the DiskSafe LVM volume.

   ```
   # vgscan
   ```

7. Activate the volume group that was found, vg_rh6u5 in this example.

   ```
   # vgchange -ay vg_rh6u5
   ```

8. Mount the lv_root volume to /mnt.

   ```
   # mount /dev/vg_rh6u5/lv_root /mnt
   ```

9. Edit fstab and lvm.conf to remove the DiskSafe reference from the two files below. Refer to KB article 1548 for more information.

   ```
   # vi /mnt/etc/fstab
   # vi /mnt/etc/lvm/lvm.conf
   ```

10. Update the ifcfg-eth0 file to this so that it can obtain the IP via DHCP.

```
# vi /mnt/etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
ONBOOT=yes
BOOTPROTO=dhcp
TYPE=Ethernet
USERCTL=no
PEERDNS=yes
IPV6INIT=no
```

11. Make sure networking is enabled.

```
# vi /mnt/etc/sysconfig/network
NETWORKING=yes
HOSTNAME=localhost.localdomain
```

12. Make sure the resolv.conf file is empty.

```
# vi /mnt/etc/resolv.conf
```

13. Remove the Ethernet information from the udev net rule so that it will add a new entry during the next boot.

```
# vi /mnt/etc/udev/rules.d/70-persistent-net.rules
Umount the lv_root partition
# umount /mnt
```

14. Deactivate the volume group, delete the partition mapping, delete the loopback device, and then run a volume group scan.

```
# vgchange -an vg_rh6u5
# kpartx -d /dev/loop0
# losetup -d /dev/loop0
# vgscan
```

Information in this document is provided "As Is" without warranty of any kind, and is subject to change without notice by FalconStor, which assumes no responsibility for any errors or claims herein. Copyright 2018 FalconStor Software. All rights reserved. FalconStor Software, FalconStor and RecoverTrac are trademarks or registered trademarks of FalconStor Software, Inc. in the United States and other countries. All other company and product names contained herein are or may be trademarks of the respective holder. PIDCIG01209

FALCONSTOR®

Corporate Headquarters
701 Brazos Street, Suite 400
Austin, TX 78701
Tel: +1.631.777.5188
salesinfo@falconstor.com

Europe Headquarters
Rosa-Bavarese-Strasse 3
80639 München, Germany
Tel: +49 (0) 89.41615321.10
salesemea@falconstor.com

Asia Headquarters
Room 1901, PICC Office Tower
No. 2 Jian Guo Men Wai Avenue
Chaoyang District Beijing 100020 China
Tel: +86.10.6530.9505
salesasia@falconstor.com