vFXT Installation Guide for AWS

Avere Systems, Inc.

Jun 20, 2018

www.averesystems.com

CONTENTS

CHAPTER

ONE

THE VFXT PLATFORM

The Avere Virtual FXT Edge Filer (vFXT) leverages the scalability of cloud computing to make your data accessible when and where it's needed - even for data that's stored in your own hardware storage environment.

Each vFXT edge filer node is a virtual machine that runs Avere OS software. vFXT nodes can use multiple types of storage media including RAM, SSD, and magnetic media.

Three or more nodes make up a vFXT cluster.

In the Amazon Web Services (AWS) environment, the vFXT nodes are virtual machines created within the Amazon Elastic Compute Cloud (EC2) using the vFXT Amazon Machine Image (AMI).

1.1 vFXT instance types

An EC2 *instance* is a virtual machine. vFXT nodes are Amazon EC2 virtual machines with the Avere vFXT AMI.

vFXT nodes in AWS can be created from these supported instance types:

Instance Size	vCPUs	Memory (GiB)	SSD Storage	Networking Performance
r4.2xlarge	8	61	EBS only	Up to 10 Gigabit
r4.8xlarge	32	244	EBS only	10 Gigabit
r3.2xlarge	8	61	1 x 160 GB	High
r3.8xlarge	32	244	2 x 320 GB	10 Gigabit

Note:

- r3 instances are not available in all regions (examples include eu-west-2 and ca-central-1). Contact Avere Global Services if this limitation applies to you.
- Disk cache per node is configurable and can range from 1000GB to 8000GB.

CHAPTER

TWO

PRODUCT REQUIREMENTS

This section outlines the prerequisites for creating a vFXT cluster.

1. Choose the *AWS region and availability zone (AZ)* (page 3) for the vFXT cluster
2. Establish an Amazon Web Services account and note the *account number* (page 4)
3. Create a role that has the required *permissions* (page 4) to create and administer the vFXT cluster
4. Create an *AWS virtual private cloud (VPC)* (page 4)
5. Create a *subnet* (page 5) within the VPC
6. Customize a *security group* (page 5) to use for your cluster instances
7. Ensure that *internet access* (page 6) is available for instances within the subnet

2.1 Choosing the AWS region and availability zone

You must choose which AWS region will host your cluster, and then choose one or more availability zones (AZ) within that region. Regions are geographically based, and availability zones can be thought of as individual data centers within the same region.

Latency is lowest for traffic within one AZ. Traffic between AZs in the same region also have low latency. Traffic between AZs in different regions can have high latency.

2.1.1 Amazon Web Services regions

You must choose which AWS region will host your virtual nodes.

AWS provides resources in global regions - for example, US-West-1 is in northern California, US-East-1 in northern Virginia, and EU-West-1 in Ireland.

Choose a region that is geographically close to the people who need to access your cluster, including data administrators and any client machines that exist outside of the AWS VPC.

2.1.2 AWS availability zones (AZ)

Determine how to place your vFXT cluster (or clusters) in appropriate AWS availability zones.

AWS availability zones are similar to different data centers within a region. For example, the availability zones us-west-1a, us-west-1b, and us-west-1c are all in separate physical buildings within the us-west-1 region.

The recommended configuration is to place all nodes for a single vFXT cluster within a single AZ. Customers can have a cluster in one availability zone and one in another AZ.

It also is possible to create a cluster with nodes in multiple AZs, which provides disaster recovery support and enhanced fault tolerance but increases latency for routine cluster transactions. Read *Multiple availability zone (Multi-AZ) support* (page 23) in this document's appendix to learn about the drawbacks, possible benefits, and additional configuration required to create a multi-AZ vFXT cluster.

2.2 AWS account

The vFXT AMI is a private AMI for customers who are evaluating or purchasing the vFXT series platform.

To access the vFXT AMI and create vFXT nodes, you must provide a valid AWS account number to Avere Systems. Avere provisions access to the AMI based on account ID.

Make sure that you provide an account number that is valid for the region (or regions) where you plan to run your vFXT instances. If you want to create a cluster in a non-public region, you must provide an account with privileges in that region.

2.3 AWS permissions

AWS user permissions are controlled in the Identity and Access Management (IAM) section of the AWS console. Several specific permissions are needed to create a vFXT node.

For a list of IAM permissions needed for vFXT cluster creation in Amazon EC2, please see *Creating the IAM role* (page 11).

2.4 AWS virtual private cloud (VPC)

When creating a vFXT cluster, you must provide it a VPC identifier.

An AWS VPC is a private network infrastructure that allows the segregation of network resources (like servers, routers, and clients) from other networks, both public and private.

A VPC runs within a single AWS region.

When creating a VPC, you provide a CIDR block of IP addresses for the VPC - for example, 10.99.0.0/16.

Because the VPC isn't exposed to the internet, a network administrator must configure how users connect to the cluster.

Connection methods include:

- AWS Direct Connect with all traffic passing through the corporate firewall
- A persistent VPN established between an Availability Zone and the corporate network
- A NAT instance managing internet-based traffic into and out of a VPC and its subnets

This document includes instructions for configuring a NAT instance, since that's the most commonly used method.

Each VPC has an identifier, or VPC ID, which is used to reference that unique VPC. The ID is a series of letters and numbers, like vpc-abcd5678. The VPC ID is required when configuring a vFXT cluster.

For more information, refer to Amazon's documentation on VPCs[1].

[1] http://aws.amazon.com/documentation/vpc/

2.5 Subnet

Within the VPC, you need one or more IP subnets for the cluster.

- Each subnet must be contained in a single availability zone.
- If using NAT for VPC access, you must create one public subnet and one private subnet. The public subnet hosts the NAT instance that provides access to the internet gateway. The private subnet serves the vFXT cluster.
- Make sure that the subnet has sufficient IP addresses available to support the instances and services that use it. The vFXT cluster needs at least two IP addresses per node, a few IP addresses for cluster overhead, and a range of client-facing addresses to service requests. Any client instances that reside within the cluster's subnet also will need IPs.

For more information, refer to Amazon's documentation on Subnets for VPCs[2].

2.6 Security groups

You should create a customized security group for your cluster before starting the vFXT node and cluster creation process.

An AWS security group is a firewall that is assigned to an instance. Security groups control incoming and outbound traffic for that instance based on port and on source and destination IP addresses.

Security groups control traffic at the instance level, not at the subnet level.

For example, a security group can include a rule that allows inbound port 22 TCP traffic from 192.168.0.0/16. When that security group is associated with one or more instances, those instances will accept inbound traffic from that IP address range on port 22.

Default security group settings don't allow all of the kinds of traffic that a vFXT cluster needs. During cluster creation, you must either select a preexisting security group or create a new one with the default configuration. If you create a default security group, you will have to reconfigure it before you can use the cluster.

For more information, refer to Amazon's documentation on Security Groups for VPCs[3].

2.6.1 Security group settings

When instantiating a vFXT cluster, you need to specify a security group. However, the inbound traffic settings on this security group must be customized to support the cluster. You can change the security group settings either before or after creating the cluster.

By default, security groups only allow traffic that originates from within the security group.

[2] http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Subnets.html#VPCSubnet
[3] http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_SecurityGroups.html

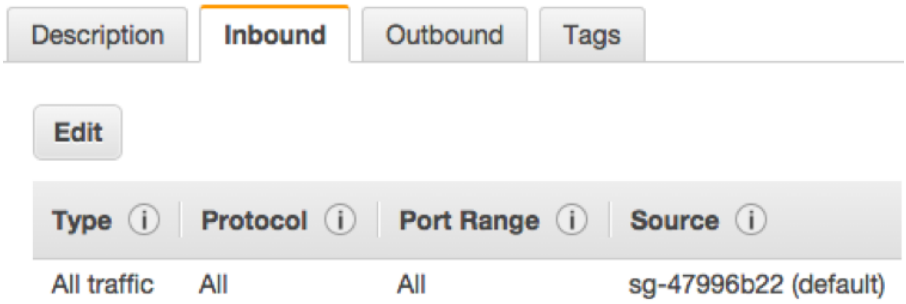

You must configure the security group to allow all traffic originating from that subnet.

To change inbound traffic rules, click the **Edit** button and change the source setting from the security group name (sg-xxxxxxxx) to the subnet range (for example, 10.0.0.0/8).

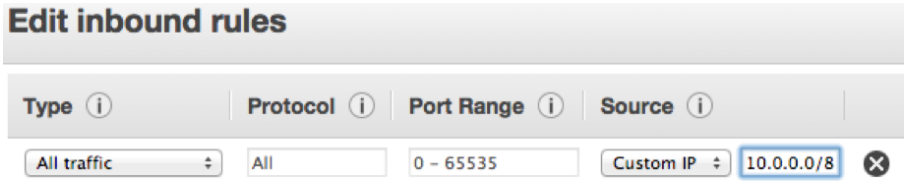

You also must make sure the port rules allow inbound traffic from ports 22, 80, and 443.

Outbound rules also can be customized, but the vFXT cluster does not require any specific configuration for outbound traffic.

2.6.2 Changing security groups

You can switch from one security group to another after the cluster is created and configured. To change the group, right-click on a vFXT instance and select **Change Security Groups**. This change can be done while the nodes are running, but be sure to repeat the change for all of the nodes in the cluster.

2.7 Internet access

vFXT nodes need internet access to communicate with AWS, but they must not be assigned public IP addresses. So the EC2 instance must be able to access the internet but not by using a public address.

This is usually accomplished by providing access using one of these methods:

- AWS Direct Connect
- A VPN tunnel
- A NAT instance

Review the *Internet exposure* (page 10) section of the *Precautions* (page 9) chapter for warnings.

The cluster requires internet access for several infrastructure tasks:

- During cluster setup, it needs to access the latest vFXT updates and to create an S3 bucket (if you are using one).
- During cluster operation, it needs access when IP addresses need to move from one node to another for load-balancing or failover reasons. This movement of IP addresses must be communicated back to AWS through APIs, requiring access to Amazon EC2 in that region.

- The cluster also requires name resolution (DNS) for amazonaws.com addresses. Name resolution for the amazonaws.com domain is typically handled by a private DNS server that is automatically assigned by Amazon at the time of the EC2 instance creation.

2.7.1 Routing

If using NAT for internet access, you will need to configure two route tables in the VPC - one for the public subnet used by the NAT instance, and one for the private subnet used by the cluster vFXT nodes.

The public route table should point its default route (0.0.0.0/0) to the Internet Gateway (IGW). The private route table should point its default route to the NAT instance.

After creating the route tables, they must be associated with their respective subnets.

2.7.2 AWS NAT configuration

A popular solution to provide internet access to vFXT nodes is to use Network Address Translation (NAT) to provide internet access to machine instances within the VPC. Instances in the private subnet pass traffic to the NAT gateway. The NAT gateway passes traffic to the VPC's internet gateway (IGW). External traffic will use the NAT gateway's elastic IP (EIP) address, which is a public-facing IP address.

NAT setup and configuration information can be found in the following links:

- Managed NAT Gateway[4] (recommended)
- VPC NAT Instance[5]
- VPC with Public and Private Subnets (Scenario 2)[6]

If using either of the latter two options, you will need to disable the NAT source/destination check to allow communication to the internet. More information about disabling that check can be found here[7].

[4] https://aws.amazon.com/blogs/aws/new-managed-nat-network-address-translation-gateway-for-aws/
[5] http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html
[6] http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_Scenario2.html
[7] http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html#EIP_Disable_SrcDestCheck

2.7.3 Reference diagram

This diagram illustrates relationships in vFXT AWS infrastructure.

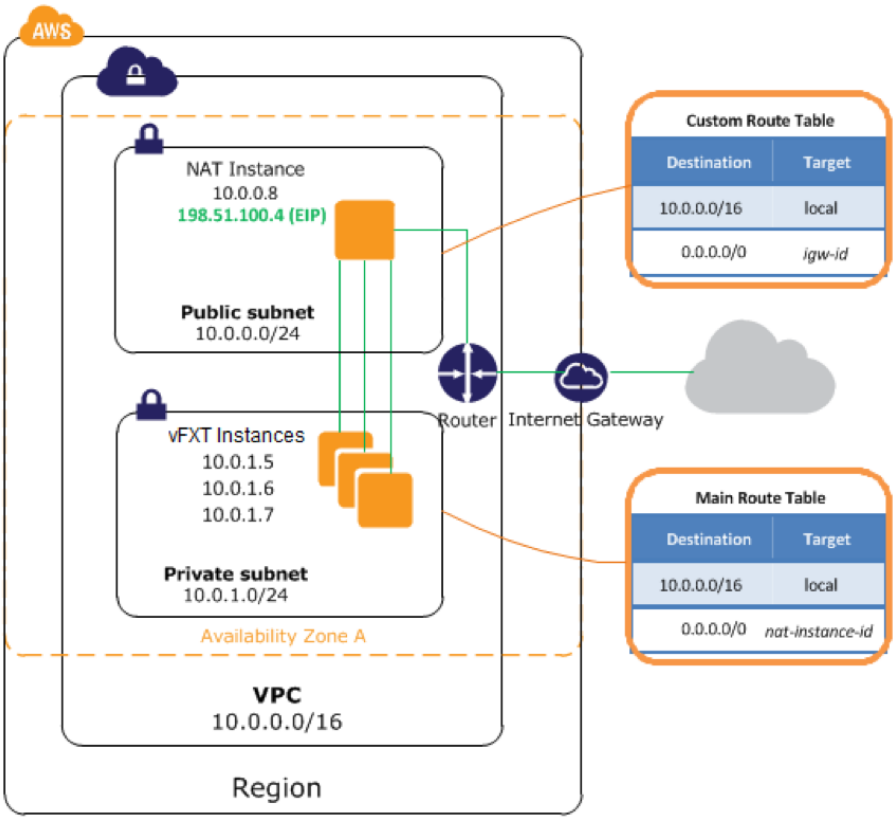

For more information, refer to Amazon's documentation on NAT instances for VPCs[8].

[8] http://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/VPC_NAT_Instance.html

CHAPTER

THREE

PRECAUTIONS

This section highlights some topics to consider when planning your vFXT cluster.

3.1 AWS resource limits

To make sure your vFXT cluster has access to sufficient computing power, plan your installation to avoid exceeding any resource limits.

Consider existing AWS EC2 instances and EBS storage currently in use in your account before attempting to create a vFXT cluster.

Limits are imposed per account on a variety of resources, including *storage* (page 9), *instances* (page 9), and *buckets* (page 9).

3.1.1 Storage limits

Storage on EC2 instances uses Elastic Block Store[1] (EBS) volumes. The vFXT uses EBS general purpose (gp2) SSD volumes. AWS imposes EBS volume limits[2] per account, including 5,000 EBS volumes and 20 TiB.

Limits can be increased by requesting a service limit increase[3].

Each vFXT node requires a minimum amount of EBS storage during node creation. The amount of EBS storage needed depends on the selections made when creating the cluster. For example, if you try to create a three-node cluster with 7000 GB of storage per node, your cluster would require 21 TiB, which is over the 20 TiB limit.

Note that these limits are *per account*. If there are other instances in the account using gp2 volumes, those volumes also count toward the 20 TiB limit even before the first vFXT instance is created.

3.1.2 Instance limits

There also are limits on the number of instances that can be created within an account. For vFXT instance types, the limits are 20 r4.2xlarge or r3.2xlarge instances; and 5 r4.8xlarge or r3.8xlarge instances. (These limits are for on-demand instances; reserved instance limits are 20 for both types.)

For instance, you cannot create two three-node clusters with r4.8xlarge nodes within the same account unless you have received a service limit increase.

3.1.3 Bucket limits

If your cluster uses S3 buckets as core filers, also note that there is a limit of 50 buckets per AWS account.

[1] http://aws.amazon.com/ebs/
[2] http://docs.aws.amazon.com/general/latest/gr/aws_service_limits.html#limits_ebs
[3] https://console.aws.amazon.com/support/home#/case/create?issueType=service-limit-increase&limitType=service-code-ebs

3.2 AWS account charges

Amazon Web Services charges are incurred for (but are not limited to) the following types of use:

- Running EC2 instances
- EBS volumes (EC2 storage), even when the machine is not running
- S3 storage
- Data transfer into and out of AWS
- Data transfer between availability zones if using a multi-AZ configuration

Avere Systems recommends that administrators monitor all AWS charges and set up billing alerts.

For more information, refer to Amazon's pricing page[4] as well as Amazon's documentation on monitoring estimated charges[5].

3.3 Encryption key management

After the vFXT cluster has been created, it is strongly recommended that you create a new cloud encryption key and save the key file before using your new cluster.

Instructions for creating a new cloud encryption key can be found in the Core Filer > Cloud Encryption Settings[6] section of the *Cluster Configuration Guide*.

3.4 Internet exposure

vFXT nodes require internet access, but they should not be directly exposed to the internet.

Caution: Avere vFXT nodes are not hardened for direct internet access.

The nodes must sit behind a firewall to protect them against attacks. This requirement also applies to any clients or servers within your network.

Most Avere customers use an EC2-based NAT instance to allow designated traffic to traverse public and private subnets within a VPC. Other customers extend their corporate network infrastructure to AWS by using a VPN or AWS Direct Connect. Read *Internet access* (page 6) for details about configuring NAT for your cluster VPC.

[4] http://aws.amazon.com/pricing/
[5] http://docs.aws.amazon.com/AmazonCloudWatch/latest/DeveloperGuide/monitor_estimated_charges_with_cloudwatch.html
[6] http://library.averesystems.com/ops_guide/4_7/gui_cloud_encryption_settings.html#cloud-encryption-settings

CHAPTER

FOUR

INSTALLING A VFXT SERIES EDGE FILER IN AMAZON EC2

To create the vFXT cluster, follow these basic steps:

1. Create and configure network infrastructure (itemized below in *Network infrastructure checklist* (page 11) and described in more detail in the *Product requirements* (page 3) section)
2. Create the IAM role for the cluster controller (details in *Creating the IAM role* (page 11))
3. Create the cluster controller instance and configure it with the required software (described in *Creating a cluster controller* (page 13))
4. From the cluster controller instance, run the vfxt.py script and create the cluster (details in *Creating a vFXT cluster* (page 15))

4.1 The vfxt.py script

Avere Systems has created a script that automates vFXT cluster creation. The vfxt.py script can be run from a machine within the cloud (recommended) or from a local machine. Complete details for using the script can be found in the vfxt.py Usage Guide[1].

4.2 Network infrastructure checklist

As the first step in creating a vFXT cluster, create and configure the following network infrastructure in your AWS project.

- VPC - read *AWS virtual private cloud (VPC)* (page 4)
- Public and private subnets - read *Subnet* (page 5)
- Public and private route tables associated with the corresponding subnets - read *Routing* (page 7)
- Security group that allows access to the VPC's IP address range and ports 22, 80, and 443 - read *Security groups* (page 5)
- Internet access through an internet gateway (IGW), NAT, or VPN - read *Internet access* (page 6)

4.3 Creating the IAM role

A cluster controller instance is used to perform several types of operations within the AWS cloud in order to create the cluster. The cluster controller requires permission to create and modify AWS entities. To assign it these permissions, you must create a role that includes the needed permissions, and then use a profile to attach the role to the instance.

Avere Systems recommends creating an IAM role specifically for the cluster controller instance.

[1] http://library.averesystems.com/#vfxt

There are two parts to creating the IAM role. First, define the policy with the permissions for cluster creation. Second, add the policy to a role so that it can be connected to the cluster controller instance.

Take these steps in the AWS **Identity & Access Management** console.

4.3.1 Creating the policy

Follow these steps to create the permissions policy for the cluster controller role.

1. Navigate to the **Identity & Access Management** service within AWS.
2. Click **Policies** on the left.
3. Click the blue **Create Policy** button.
4. Choose **Select** next to **Create Your Own Policy**.
5. Provide a policy name - for example, `vfxt-policy`
6. Copy and paste the following policy:

   (This text also appears in the *appendix* (page 25) without page breaks.)

```
{
    "Statement": [
    {
        "Resource": "*",
        "Action": [
            "ec2:Describe*",
            "ec2:RunInstances",
            "ec2:TerminateInstances",
            "ec2:RebootInstances",
            "ec2:StartInstances",
            "ec2:StopInstances",
            "ec2:CreateTags",
            "ec2:DeleteTags",
            "ec2:ModifyInstanceAttribute",
            "ec2:CreateVolume",
            "ec2:DeleteVolume",
            "ec2:AttachVolume",
            "ec2:DetachVolume",
            "ec2:CreateSnapshot",
            "ec2:DeleteSnapshot",
            "ec2:RegisterImage",
            "ec2:DeregisterImage",
            "ec2:CreateImage",
            "ec2:DeleteRoute",
            "s3:CreateBucket",
            "s3:DeleteBucket",
            "s3:SetTag",
            "s3:ListBucket",
            "iam:AddRoleToInstanceProfile",
            "iam:CreateInstanceProfile",
            "iam:CreateRole",
            "iam:DeleteInstanceProfile",
            "iam:DeleteRole",
            "iam:DeleteRolePolicy",
            "iam:GetRolePolicy",
            "iam:ListRolePolicies",
```

(continues on next page)

(continued from previous page)

```
            "iam:GetInstanceProfile",
            "iam:PutRolePolicy",
            "iam:RemoveRoleFromInstanceProfile",
            "iam:GetRole",
            "iam:PassRole"
        ],
        "Effect": "Allow"
    }
    ],
    "Version": "2012-04-18"
}
```

7. Click the **Validate Policy** button.
8. Click the blue **Create Policy** button.

4.3.2 Creating the cluster controller role

After the cluster creation policy has been created, you must create a new role that uses the policy.

1. Navigate to the **Identity & Access Management** service within AWS.
2. Click **Roles** on the left.
3. Click **Create New Role**.
4. Enter a name for the role - for example, `vfxt-role`
5. Click **Next Step**.
6. Choose **Select** next to **Amazon EC2**.
7. Click the checkbox next to the cluster controller policy.
8. Click **Next Step**.
9. Click **Create role**.

When you create the cluster controller instance, you will assign it the role that you created in this step.

4.4 Creating a cluster controller

Avere Systems recommends installing and running the vfxt.py script on a Linux-based instance in the cloud. This instance, called a cluster controller instance, can be a pre-existing instance, or you can create a new one. In either case, you must assign it the cluster controller role that you defined in the previous step.

After you create the cluster controller instance, you will download the required software for creating the vFXT cluster.

To create a new cluster controller instance:

1. Navigate to the EC2 service within AWS.
2. Click **Instances** on the left.
3. Click **Launch Instance**.
4. Click **Community AMIs** on the left.
5. You must find and install a current, stable Debian release.

At the time this document was prepared, the stable distribution was Debian 9 and nicknamed "Stretch".

In the search field, type `stretch 2018` and press **Enter**. Click **Select** next to an appropriate image. (Usually, the first search result is appropriate.)

6. Leave the instance type at `t2.micro`.
7. Click **Next: Configure Instance Details**.
8. Choose the network and private subnet for the cluster controller instance.

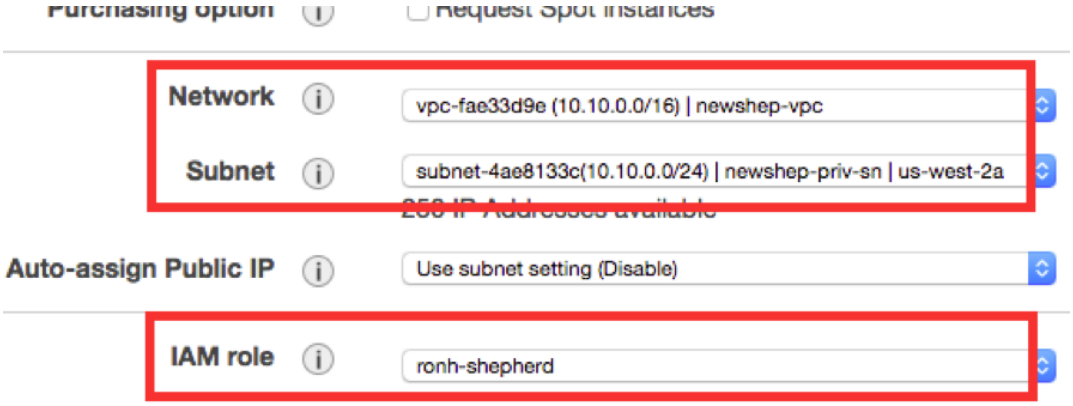

10. Choose a public IP address (**Enable**).
11. Choose the **IAM role** that was created for the cluster controller (for example, `vfxt-role`).
12. Click **Next: Add Storage**.
13. Click **Next: Tag Instance**.
14. Enter **Name** for a tag and provide a name for the cluster controller instance.
15. Click **Next: Configure Security Group**.
16. Choose **Select an existing security group**.
17. Click the box next to the security group that you customized for the cluster (see *Security groups* (page 5)).
18. Click **Review and Launch**.
19. Review the information on the page and click **Launch**.
20. Choose an existing key pair or create a new key pair.
21. Click **Launch Instances**.

4.4.1 Accessing the cluster controller instance

Use SSH to connect to the cluster controller.

You will need the cluster controller's IP address and key file in order to access it with SSH. Click on the instance in EC2 and make note of its IP address (at the bottom). Download its key file to a convenient location.

SSH example:

```
ssh -i ~/Downloads/key.pem admin@54.186.72.16
```

4.4.2 Configuring the cluster controller

After logging in, download the necessary files needed to run the vfxt.py script. Refer to the vfxt.py Usage Guide[2] for details.

For example:

```
sudo apt-get update
sudo apt-get install python-pip
sudo easy_install pip
sudo pip install --upgrade boto requests google-api-python-client awscli vFXT
```

4.5 Creating a vFXT cluster

After creating and configuring the cluster controller instance, use it to run the vfxt.py script to create the vFXT cluster. Complete instructions are included in the vfxt.py Usage Guide[3].

You will need to supply the following information for the command:

- AWS region - for example, `us-west-2a`
- private subnet ID - for example, `subnet-6e82b427`
- security group ID - for example, `sg-3d75aa47`

Note: If using AWS S3 as a core filer (your cluster's backend storage), you will want to enable S3 endpoints, which optimize throughput and availability for S3 traffic into and out of the VPC. See *Creating AWS S3 endpoints* (page 22) in the appendix of this document.

Refer to the vfxt.py Usage Guide[4] for additional information on how to use the cluster controller to create a vFXT cluster.

(The sample vfxt.py command below includes line breaks for readability; enter the command on one line or with escape characters before line breaks in actual usage.)

Sample vfxt.py command:

```
vfxt.py --create --cloud-type aws --on-instance --region us-west-2a
        --subnet subnet-6e82b427 --instance-type r3.2xlarge
        --cluster-name vfxt-demo-cluster --admin-password P@$$w0rd
        --node-cache-size 1000 --nodes 3 --debug
        --security-group sg-3d75aa47
```

Make note of the *management IP address* printed at the end of the script. You will use this IP address and the administrative password you set in `admin-password` to access the cluster.

This sample message shows how the management IP address is reported from a newly created vFXT cluster:

```
2017-07-12T15:46:49+0000 - vfxt:INFO - vfxt-cluster management address: 10.10.1.4
```

[2] http://library.averesystems.com/#vfxt
[3] http://library.averesystems.com/#vfxt
[4] http://library.averesystems.com/#vfxt

4.6 Post-creation configuration

The vfxt.py script creates and configures the vFXT cluster. To administer the finished cluster, use the Avere Control Panel. Although many settings are set appropriately for your cloud service by the creation script, there are others that you might want to customize.

The FXT Cluster Creation Guide[5] is designed for clusters of physical hardware nodes, but some information in the document is relevant for vFXT clusters as well. In particular, these sections can be useful for vFXT cluster administrators:

- Logging In to the Avere Control Panel[6] explains how to connect to the Avere Control Panel and log in. However, note that you must use a VPN or SSH tunnel to access the cluster nodes inside the AWS VPC. Read *Accessing vFXT nodes (HTTPS and SSH)* (page 17) for details.
- Configuring VServers and Global Namespace[7] has information about creating a client-facing namespace.
- Adding Backend Storage[8] documents how to add storage.

  (Note: You might need to update the cluster IAM policy to add a second cloud core filer to your AWS cluster - read the cloud core filers tip in the *AWS Cluster Settings* section of the vfxt.py Usage Guide[9].)
- Customizing Support and Monitoring Settings for the Avere Cluster[10] explains how to customize support settings and remote monitoring.

These additional documents also might be helpful:

- The Cluster Configuration Guide[11] is a complete reference of settings and options for an Avere cluster. A vFXT cluster uses a subset of these options, but many of the same configuration pages apply.
- The Dashboard Guide[12] explains how to use the cluster monitoring features of the Avere Control Panel.

Current documents can always be found on the documentation website at http://library.averesystems.com.

[5] http://library.averesystems.com/#fxt_cluster
[6] http://library.averesystems.com/create_cluster/4_8/html/initial_config.html#gui-login
[7] http://library.averesystems.com/create_cluster/4_8/html/config_vserver.html#config-vserver
[8] http://library.averesystems.com/create_cluster/4_8/html/config_core_filer.html#add-core-filer
[9] http://library.averesystems.com/#vfxt
[10] http://library.averesystems.com/create_cluster/4_8/html/config_support.html#config-support
[11] http://library.averesystems.com/#operations
[12] http://library.averesystems.com/#operations

CHAPTER

FIVE

MANAGING VFXT NODES

This section explains how to connect to vFXT nodes and stop or start instances that are part of the cluster.

For additional cluster configuration, log in to the cluster's web-based Avere Control Panel. Through this interface you can add backend storage (core filers), customize caching parameters, and change many other settings. The *Post-creation configuration* (page 16) section of this document gives links for learning how to connect to the Avere Control Panel and about more configuration tasks for vFXT clusters.

5.1 Accessing vFXT nodes (HTTPS and SSH)

Because a vFXT cluster is inside a VPC, you must connect through a VPN or gateway that will allow access from outside addresses.

- If you use a VPN to access AWS, you can enter the management IP into your browser.
- If your have AWS Direct Connect access, use that to access the cluster's management IP address.
- If you use NAT to provide internet access to the cluster VPC, you can use an SSH tunnel to access the vFXT nodes.

5.1.1 SSH tunnel access

For an SSH tunnel, you will need the management IP address of the cluster and the public IP address of another instance (for example, the cluster controller instance).

To create an SSH tunnel for vFXT access:

1. Open a terminal session on your local machine.
2. Enter the tunnel command, which has the following form:

   ```
   ssh -L <localPort>:<managementIPaddress>:443
       <user>@<publicIPaddress>
       -i <pathToKeyFile>
   ```

   Example:

   ```
   ssh -L 8443:10.10.2.4:443 admin@52.38.22.162 -i ~/Downloads/key.pem
   ```

3. Open a new tab in your browser.
4. Type `https://127.0.0.1:<localPort>` and press **Enter**.

   Example:

   ```
   https://127.0.0.1:8443
   ```

5. Click **Advanced** and click **Proceed** to bypass the warning and access the login page.

6. Enter `admin` for the username and enter the password you chose when running the `vfxt.py --create` command.

5.2 Managing instances

This section explains how to stop, restart, and destroy cloud instances that serve as vFXT cluster nodes.

5.2.1 Stopping instances

If you need to stop an instance (one node) or the entire cluster and intend to restart it later, Avere Systems recommends using the Avere Control Panel.

The **FXT Nodes** settings page has controls for shutting down or rebooting individual nodes. (Note that IP addresses might move among cluster nodes when the number of active nodes changes.) Read Cluster > FXT Nodes[1] for more information.

To stop or reboot the entire cluster, use the **System Maintenance** settings page. Read Administration > System Maintenance[2] for details.

If you need to stop an instance or the entire cluster but do not intend to restart it, you can terminate the instance by using tools within the Amazon EC2 console. See *Terminating instances* (page 18) for more information.

Note: Although EC2 charges are not incurred while instances are stopped, storage charges will continue for any S3 buckets and EBS volumes associated with the vFXT node.

5.2.2 Restarting instances

If you need to restart a stopped instance, you must use the AWS console. Navigate to **EC2 > Instances** and right-click to select one or more instances that you want to restart. Under the **Actions** section near the bottom, choose **Start**.

5.2.3 Terminating instances

Caution: Terminated instances cannot be restarted or retrieved. Instance termination is a permanent action and cannot be undone.

Before terminating a vFXT instance, remove it from the cluster or shut down the cluster as described below in *Terminating one node* (page 19) and *Terminating all nodes in the vFXT cluster* (page 19).

To permanently destroy one or more instances used as vFXT node, use the AWS console. Navigate to **EC2 > Instances** and right-click to select one or more instances that you want to destroy. Under the **Actions** section near the bottom, choose **Terminate**.

[1] http://library.averesystems.com/ops_guide/4_7/gui_fxt_nodes.html#gui-fxt-nodes
[2] http://library.averesystems.com/ops_guide/4_7/gui_system_maintenance.html#gui-system-maintenance

Terminating one node

If you want to terminate one node from the vFXT cluster but keep the remainder of the cluster, you must first remove the node from the cluster using the Avere Control Panel.

> **Caution:** If you terminate a node without first removing it from the vFXT cluster, data might be lost.

After removing the node, follow the instructions in *Terminating instances* (page 18) above to destroy it.

Terminating all nodes in the vFXT cluster

To if you are finished using the vFXT cluster and want to permanently delete it, you should shut down the cluster by using the Avere Control Panel first. A graceful shutdown allows any unsaved client changes to be written to permanent storage, ensuring data integrity.

Use the Administration > System Maintenance[3] settings page to power down the cluster. After the cluster has stopped posting messages to the **Dashboard** tab, the Avere Control Panel session will stop responding and you will know that the cluster has been shut down.

After shutting down the cluster, use the AWS console to destroy all of the node instances as described above in *Terminating instances* (page 18).

[3] http://library.averesystems.com/ops_guide/4_7/gui_system_maintenance.html#gui-system-maintenance

CHAPTER

SIX

APPENDIX

This section has additional information that might be needed for vFXT cluster configurations.

Click a link to jump to the following sections:

6.1 Required ports

The ports listed in this section are used for vFXT inbound and outbound communication.

Never expose the vFXT or the cluster controller instance directly to the public internet.

6.1.1 API

Inbound:		
TCP	22	SSH
TCP	443	HTTPS

Outbound		
TCP	80	HTTP
TCP	443	HTTPS

6.1.2 NFS

Inbound and Outbound		
TCP/UDP	111	RPCBIND
TCP/UDP	2049	NFS
TCP/UDP	4045	NLOCKMGR
TCP/UDP	4046	MOUNTD
TCP/UDP	4047	STATUS

6.1.3 SMB/CIFS

Inbound		
TCP	445	SMB
TCP	139	SMB
UDP	137	NETBIOS
UDP	138	NETBIOS

Outbound		
TCP/UDP	53	DNS
TCP/UDP	389	LDAP
TCP	686	LDAPS
TCP/UDP	88	Kerberos
UDP	123	NTP
TCP	445	SMB
TCP	139	SMB
UDP	137	NetBIOS
UDP	138	NetBIOS

6.2 Domain whitelist

Make sure that these web URLs are accessible from your cluster.

verisign.com[1]

ocsp.verisign.com[2]

SVRSecure-G3-crl.verisign.com[3]

s3.amazonaws.com[4]

sd.symcd.com[5]

download.averesystems.com[6]

You will also want to whitelist the AWS S3 and EC2 endpoints used by your instances. Refer to these documents for specifics:

https://docs.aws.amazon.com/general/latest/gr/rande.html#s3_region

https://docs.aws.amazon.com/general/latest/gr/rande.html#ec2_region

6.3 Creating AWS S3 endpoints

If using AWS S3 for storage as a core filer, you should enable S3 endpoints, which allows S3 traffic to pass through a customized gateway instead of using the public gateway to your VPC. This access strategy optimizes throughput for S3 traffic, making your storage system more available.

[1] http://verisign.com/
[2] http://ocsp.verisign.com/
[3] http://SVRSecure-G3-crl.verisign.com/
[4] http://s3.amazonaws.com/
[5] http://sd.symcd.com/
[6] http://download.averesystems.com/

Use the AWS console to create an endpoint:

1. Navigate to the VPC service within AWS.
2. Click **Endpoints** on the left.
3. Click the blue **Create Endpoint** button at the top.
4. Choose the VPC, service, and policy.
5. Click the blue **Next Step** button in the lower right.
6. Click the checkbox next to the private route table used by the vFXT cluster.
7. Click the blue **Create Endpoint** button in the lower right.

To verify that the endpoint was correctly created and applied to the route table:

1. Navigate to **Route Tables** within the VPC service console.
2. Click the route table at the top.
3. Click the **Routes** tab at the bottom.
4. Verify that there is a route at the bottom that begins with "pl."

Destination	Target	Status	Propagated
172.31.0.0/16	local	Active	No
pl-68a54001 (com.amazonaws.us-west-2.s3)	vpce-11639e78	Active	No

6.4 Multiple availability zone (Multi-AZ) support

The recommended Avere vFXT cluster configuration is to have all nodes running within a single AWS Availability Zone (AZ). This arrangement provides the best performance and lowest cost.

For more information, see the AWS Availability Zones documentation[7].

For disaster recovery scenarios, the vFXT cluster can support nodes in multiple AZs. This configuration increases latency and costs as node data is transferred across AZs within the same region. It allows a node in one zone to take over serving requests from a node in a zone that has gone down. The operating vFXTs reprogram routes and allow clients to connect to the same IP address used on the failed vFXT.

You can create a Multi-AZ cluster that has one node in each of three different availability zones. This is the only supported configuration.

6.4.1 Multi-AZ requirements and restrictions

To use a Multi-AZ configuration, administrators must ensure that the cluster environment meets these requirements:

- Include sufficient permissions in the vFXT IAM role for managing instances in multiple zones. In addition to the cluster controller permissions, each node must have the failover permissions
- Have an available CIDR range outside of the VPC CIDR range for the management, cluster, and client-facing addresses

[7] https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html

- Configure routes for directing traffic to and from that additional CIDR range
- Ensure that clients can route to the client-facing addresses (a typical method is to have the EC2 instances share the same route table with the vFXT nodes)

The cluster itself is subject to the following limitations:

- The cluster can use a maximum of three Availability Zones.
- Each AZ can host only one vFXT node. (Using multiple zones limits the cluster to a size of three vFXT nodes.)
- Subnets must all be within the same VPC.
- All vFXT nodes should share the same route table.

6.5 Customizing the cluster node role

The vfxt.py script creates a role for the cluster nodes. You have the option to create your own role for cluster nodes and supply it to the vfxt.py script when you create the cluster.

The policy below is the default role assigned to vFXT nodes. If you use a custom role, make sure that it includes this information:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "ec2:AssignPrivateIpAddresses",
                "ec2:UnassignPrivateIpAddresses",
                "ec2:DescribeInstance*",
                "ec2:DescribeRouteTables",
                "ec2:DescribeSubnets",
                "ec2:ReplaceRoute",
                "ec2:CreateRoute",
                "ec2:DeleteRoute"
            ],
            "Resource": "*",
            "Effect": "Allow"
        }
    ]
}
```

6.6 Sample IAM policy

The following is the policy that should be cut and pasted to create the permissions for the cluster controller.

```
{
    "Statement": [
    {
        "Resource": "*",
        "Action": [
            "ec2:Describe*",
            "ec2:RunInstances",
            "ec2:TerminateInstances",
            "ec2:RebootInstances",
            "ec2:StartInstances",
            "ec2:StopInstances",
            "ec2:CreateTags",
            "ec2:DeleteTags",
            "ec2:ModifyInstanceAttribute",
            "ec2:CreateVolume",
            "ec2:DeleteVolume",
            "ec2:AttachVolume",
            "ec2:DetachVolume",
            "ec2:CreateSnapshot",
            "ec2:DeleteSnapshot",
            "ec2:RegisterImage",
            "ec2:DeregisterImage",
            "ec2:CreateImage",
            "ec2:DeleteRoute",
            "s3:CreateBucket",
            "s3:DeleteBucket",
            "s3:SetTag",
            "s3:ListBucket",
            "iam:AddRoleToInstanceProfile",
            "iam:CreateInstanceProfile",
            "iam:CreateRole",
            "iam:DeleteInstanceProfile",
            "iam:DeleteRole",
            "iam:DeleteRolePolicy",
            "iam:GetRolePolicy",
            "iam:ListRolePolicies",
            "iam:GetInstanceProfile",
            "iam:PutRolePolicy",
            "iam:RemoveRoleFromInstanceProfile",
            "iam:GetRole",
            "iam:PassRole"
        ],
        "Effect": "Allow"
    }
    ],
    "Version": "2012-04-18"
}
```

INDEX

W