SecureTransport

Version 5.4

2 April 2024

Azure Installation Guide

This documentation describes the following Axway software:

Axway SecureTransport 5.4

Contents

Introduction 1

The current guide provides a basic outline regarding the installation of SecureTransport on a RedHat Virtual Machine as deployed on the Microsoft Azure portal. You will pass through several simple stages until you can finally proceed with your SecureTransport installation.

Currently SecureTransport deployment on Azure is verified for the RedHat 7.2 or later implementations only.

About SecureTransport

SecureTransport is part of the Axway family of managed file transfer (MFT) products. SecureTransport allows organizations to adeptly control and manage the transfer of files inside and outside of the corporate firewall in support of mission-critical business processes, while satisfying policy and regulatory compliance requirements. SecureTransport serves as a hub and router for moving files between humans, systems and more. SecureTransport also completes tasks related to moving files (push or pull), hosting files in mailboxes or "FTP-like" folders, and provides portal access with configurable workflow for file handling and routing. SecureTransport delivers user-friendly governance and configuration capabilities, including delegated administration and pre-defined and configurable workflows, while providing the highest possible level of security.

For a complete description of SecureTransport features and components, refer to the *SecureTransport Administrator's Guide*.

Overview

This document provides a detailed overview and detailed instructions to set up SecureTransport in the Microsoft Azure Virtual Network (VNet).

The Microsoft Azure Virtual Network service enables Azure resources to securely communicate with each other in a virtual network. A virtual network is a logical isolation of the Azure cloud dedicated to your subscription. You have complete control over your virtual networking environment, including selection of your own IP address range, creation of subnets, and configuration of route tables. You can connect virtual networks to other virtual networks, or to your on-premises network. The following picture shows some of the capabilities of the Azure Virtual Network service:

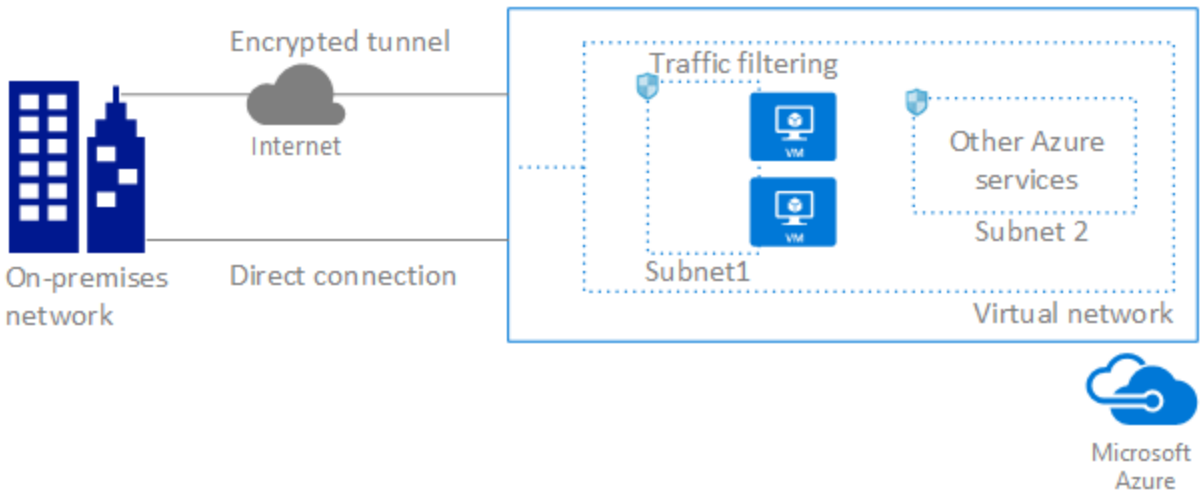

Learn more about Azure Virtual Network.

SecureTransport in Azure Virtual Network 2

Currently, deployment of SecureTransport on Azure VNet has been verified for Red Hat 7.2 only in the following setup:

- Enterprise Cluster of two servers with streaming to Standard Cluster of two edges
- Microsoft SQL Server 2016
- GlusterFS file system with two servers
- A Load Balancer (optional)
- An Administration Host (optional)
- Two private and one public subnets
- A VPN Connection (optional)
- Instances are assigned to five Network security groups (NSG)
- All this grouped in three Availability Sets

Azure operates in multiple datacenters around the world. These datacenters are grouped in to geographic regions, giving you flexibility in choosing where to build your applications. To provide additional scalability and reliability, these data center facilities are distributed in different physical locations, categorized by regions and Availability Sets.

You create Azure resources in defined geographic regions. Within each region, multiple datacenters exist to provide for redundancy and availability. This approach gives you flexibility as you design applications to create VMs closest to your users and to meet any legal, compliance, or tax purposes.

An availability set is a logical grouping of VMs within a datacenter that allows Azure to understand how your application is built to provide for redundancy and availability. We recommended that two or more VMs are created within an availability set to provide for a highly available application.

The availability set is composed of two additional groupings that protect against hardware failures and allow updates to safely be applied - fault domains (FDs) and update domains (UDs).

A fault domain is a logical group of underlying hardware that share a common power source and network switch, similar to a rack within an on-premises datacenter. As you create VMs within an availability set, the Azure platform automatically distributes your VMs across these fault domains. This approach limits the impact of potential physical hardware failures, network outages, or power interruptions.

An update domain is a logical group of underlying hardware that can undergo maintenance or be rebooted at the same time. As you create VMs within an availability set, the Azure platform automatically distributes your VMs across these update domains. This approach ensures that at least

one instance of your application always remains running as the Azure platform undergoes periodic maintenance. The order of update domains being rebooted may not proceed sequentially during planned maintenance, but only one update domain is rebooted at a time.

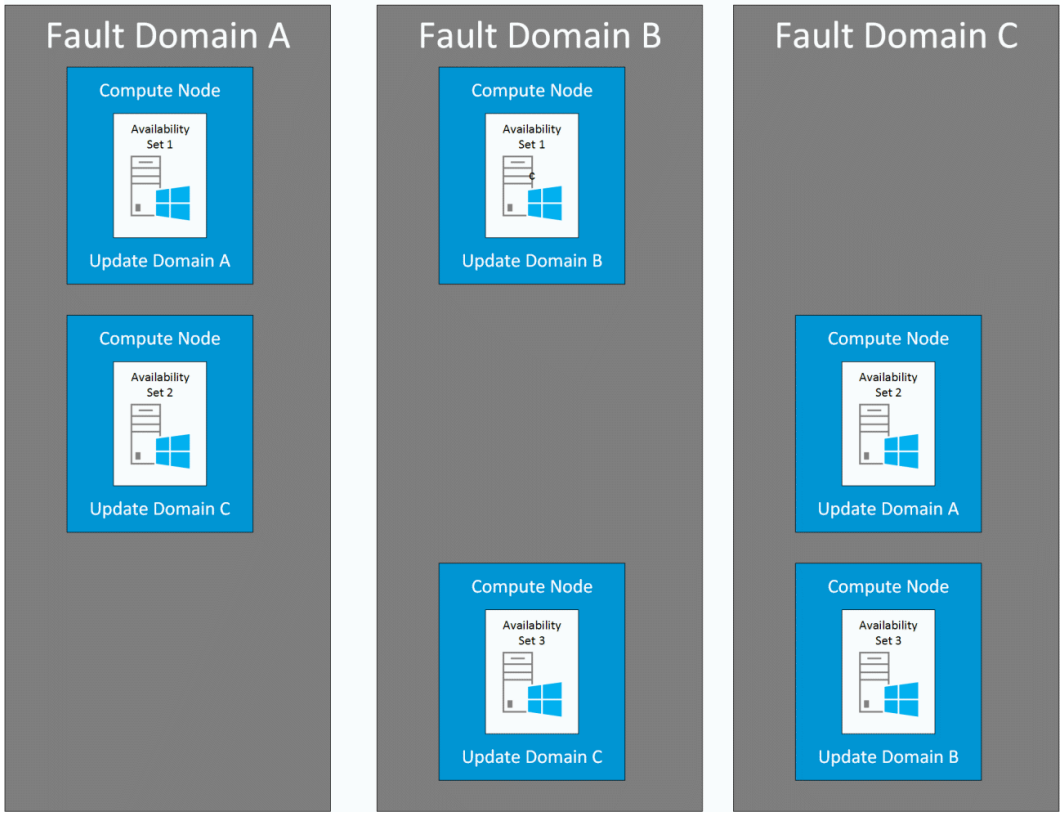

The following diagram represents the location and distribution of instances in the availability sets.

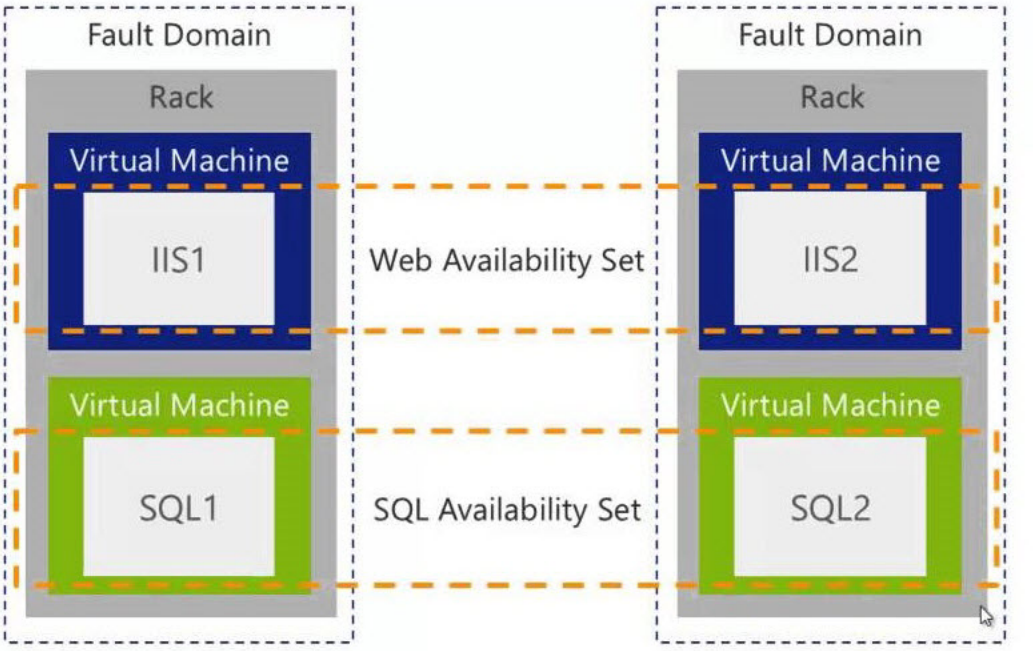

The goal is to provide a highly-available, fault-tolerant and secure setup of SecureTransport in the Microsoft Azure Cloud.

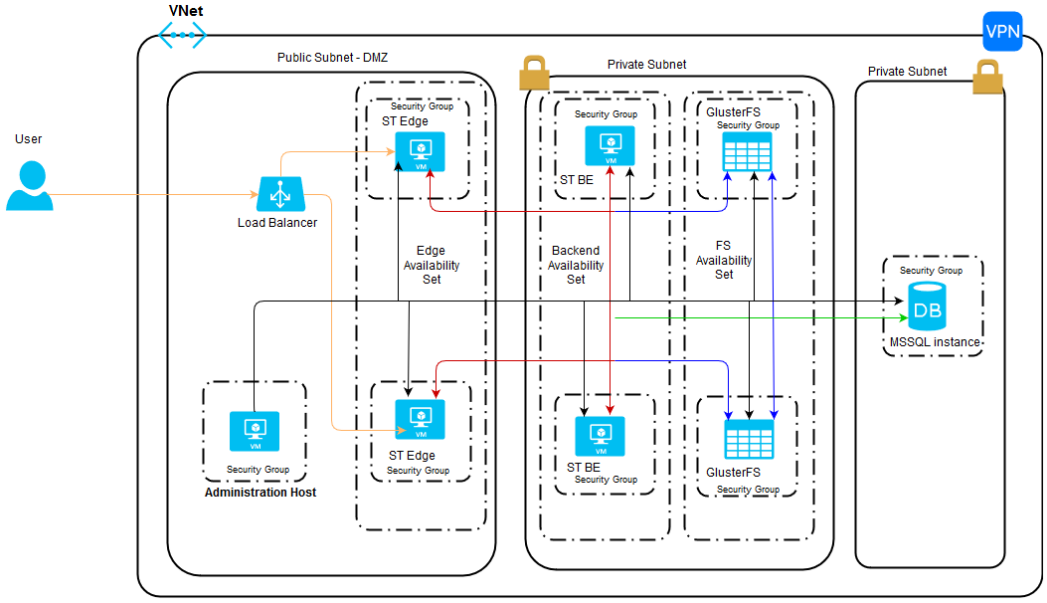

The diagram illustrates a SecureTransport setup of an Enterprise Cluster in streaming mode with two Edge servers joined in a Standard Cluster. Microsoft SQL Server installed on a Virtual Machine is used as a database engine.

The setup contains one public and two private subnets. The public subnet hosts the SecureTransport Edges that are grouped in an Availability Set and the Administration Host. The first private subnet hosts the SecureTransport Servers Availability Set and the external file system Availability Set. The database instance is located in a separate private subnet. The Edge servers are in a Standard cluster setup and are in a constant synchronization connection. The SecureTransport servers are in an Enterprise Cluster setup and in a constant synchronization connection as well. The SecureTransport servers are connected in streaming with the Edge servers as each server establishes streaming connections with both Edge servers. Internet-facing load balancer distributes requests among the Edge servers in the Edge Availability Set.

In case of a system failure in one of the machines in an Availability Set, the other machine remains fully functional with up to date system configuration.

The public subnets in the Azure Cloud are the equivalent of the DMZ (demilitarized zone) in a classic on-premise SecureTransport deployment. The Azure cloud provides security tools like Network Security Groups to protect the public and private subnets in your VNet. These tools act as the firewalls do in the classic SecureTransport on-premise deployments and control both inbound and outbound traffic at an instance and subnet level. Each group (Administration host, Edge servers, Servers, File System, Database) has a corresponding Network Security Group. The Network Security Group is a stateful firewall that works on VM or Subnet level.

There is a host placed in one of the public subnets called "Administration host" also known as "Bastion host" or "Jump host" in the networking terminology. Bastion hosts are instances that typically reside within your public subnet and are usually accessed using SSH or RDP. Once remote connectivity is established with the bastion host, it then acts as a *jump* server, allowing you to use SSH or RDP to log in to other instances publicly inaccessible deeper within your network. When properly configured

through the use of security groups, the bastion host essentially acts as a bridge to your private instances via the Internet. The Administration host is used to perform maintenance and administration tasks on the servers in the SecureTransport setup. You should always consider the resiliency and high availability of your services in cloud deployments and the best practice is to have an Availability set of Administration hosts in case one of the hosts goes down. For minimal security risks, you should stop your Administration host instances for the duration of no maintenance work and start them when you need access to your servers in Azure again.

You can connect directly to your VNet using a Point-to-Site connection in case you don't want to have an Administration host in your setup (see *Configure a Point-to-Site connection to your VNet* section in this guide).

The installation of SecureTransport on Red Hat instances in Microsoft Azure cloud follows a flow which is the same as that of an installation on a regular Red Hat machine, including the required installation prerequisites. This process is already described in the SecureTransport Installation Guide.

You will pass through several Azure cloud specific setup and configuration stages until you until you are able to proceed with your SecureTransport installation.

To set up SecureTransport in Azure Virtual Network, you must pass through the following steps:

1. Create a Virtual Network.
2. Create Availability Sets.
3. Create Network Security Groups.
4. Create one public and two private subnets.
5. Set up Microsoft SQL Server.
6. Launch the following Red Hat instances:
    a. Launch an instance for an Administration Host.
    b. Launch instances in the public subnets for SecureTransport Edge installations.
    c. Launch instances in the private subnets for SecureTransport Server installations.
    d. Launch instances in the private subnets for external GlusterFS file system.
7. Establish VPN connection to your Azure VNet.
8. Set up SecureTransport Enterprise Cluster.
9. Set up Load Balancer.

Create a Virtual Network 3

To create a VNet, follow these steps:

1. Log in to the Azure Portal and navigate to the **Create a resource**.
2. Under the "Azure Marketplace" section, choose **Networking**.
3. Under the "Featured" section, choose **Virtual Network**.
4. Fill in the settings and click **Create**.

   Note The image provides an example Address space. You can configure the Address space according to your needs.

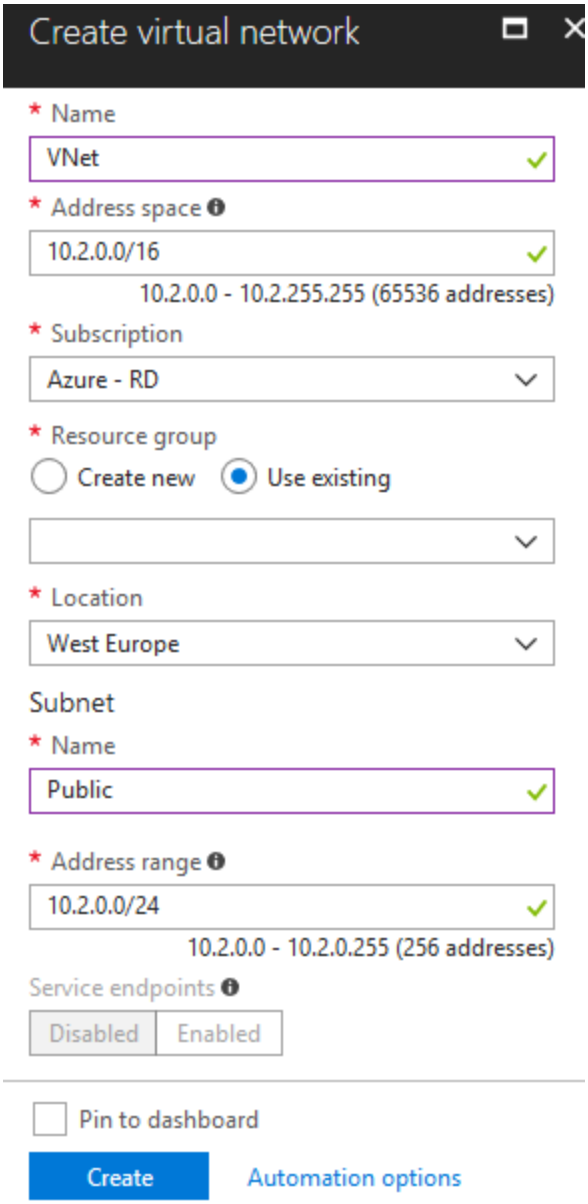

Create Availability Sets 4

An Availability Set is a logical grouping capability that you can use in Azure to ensure that the VM resources you place within it are isolated from each other when they are deployed within an Azure datacenter. Azure ensures that the VMs you place within an Availability Set run across multiple physical servers, compute racks, storage units, and network switches. If a hardware or Azure software failure occurs, only a subset of your VMs are impacted, and your overall application remains up and continues to be available to your customers. Availability Sets are essential when you want to build reliable cloud solutions.

It is recommended to create Availability Sets for SecureTransport Edges, SecureTransport Servers and GlusterFS Servers. In addition, you can also create Availability Sets for the case of multiple Administration Hosts and Database instances, when there are at least two VMs of them.

To create an Availability Set, follow these steps:

1. Navigate to the Azure Portal and click **Create a resource**.
2. Under the "Azure Marketplace" section, choose **Compute**.
3. Under the "Featured" section choose **Availability Set**.
4. On the newly opened dialog box, enter the required information and click **Create**.
    a. Give it a unique name.
    b. Select **Subscription**, **Resource group** and **Location**.
    c. Select the number of **Fault** and **Update Domains**.

Note Place the instances of one Availability Set in different Fault and Update domains.

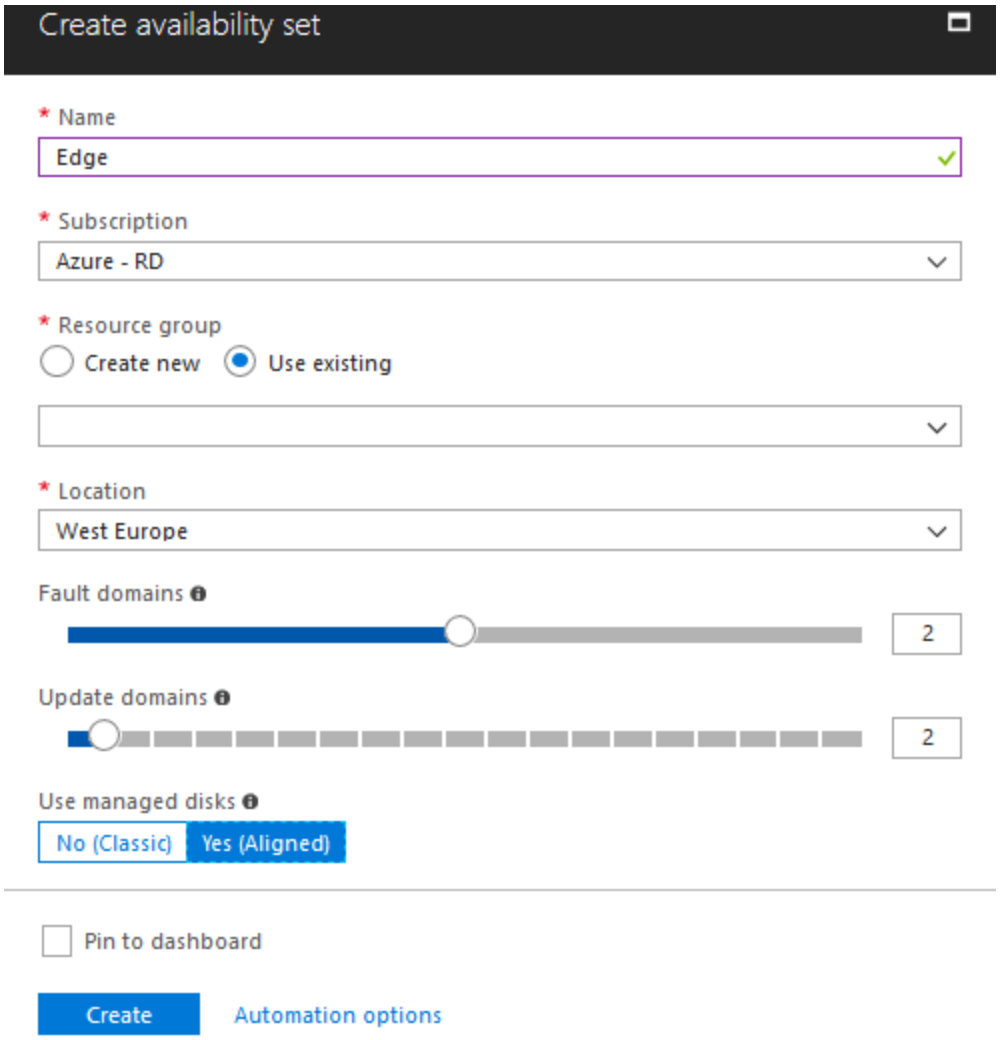
Create availability set
* Name
Edge
* Subscription
Azure - RD
* Resource group
Create new
Use existing
* Location
West Europe
Fault domains
2
Update domains
2
Use managed disks
No (Classic)
Yes (Aligned)
Pin to dashboard
Create
Automation options

Network Security Groups 5

Azure provides two features that you can use to increase security in your VNet: *network security groups* and *endpoint ACLs*. Security groups control inbound and outbound traffic for your instances or subnets, and endpoint ACLs control inbound and outbound traffic for your VM endpoint. For more information, see the Azure Security Documentation:

- Network Security groups - these act as a firewall for associated VMs or Subnets, controlling both inbound and outbound traffic.
- Endpoint access control lists (ACLs) - provides the ability to selectively permit or deny traffic for a virtual machine endpoint. This packet filtering capability provides an additional layer of security. You can specify network ACLs for endpoints only. You cannot specify an ACL for a virtual network or a specific subnet contained in a virtual network. It is recommended to use network security groups (NSGs) instead of ACLs, whenever possible.

For more details on security options check the Azure Security Best Practices guide.

Note The firewall of the instance is with higher priority than the instance's Network Security Group. So, if you would like to use NSG, the instance's firewall must be configured in the same way as the NSG's rules.

Note The suggested configurations with Network Security Groups as stated in the current guide do not include outbound rules. For more information on best practices for the outbound rules to apply for your setup, please refer to the *SecureTransport Installation guide*.

Create Network Security groups

When you launch a virtual machine in a VNet, you can associate it with one network security group that you have created. If you do not specify a network security group when you launch an instance, all traffic to this instance is allowed.

It is recommended that you create the network security groups you need for your SecureTransport infrastructure in Azure as a first stage before proceeding with rest of the setup. You need to group (assign) the instances and components into the following security groups:

- SecureTransport Edge network security group
- SecureTransport Server network security group
- External Database network security group
- External File System network security group
- Administration Host security group

It is recommended to have all components prepared prior the launch process of the instances. In this way, during your SecureTransport setup, you will just select the network security group you need from the list. You can always create the Network Security Groups on a later stage but we suggest you adhere to the flow as described.

To create a network security group in Azure VNet:

1. Navigate to the Azure Portal and go to **Create a resource**.
2. Under the "Azure Marketplace" section, choose **Networking**.
3. Under the "Featured" section, choose **Network Security Group**.
4. On the newly opened dialog box, enter the required information and click **Create**.

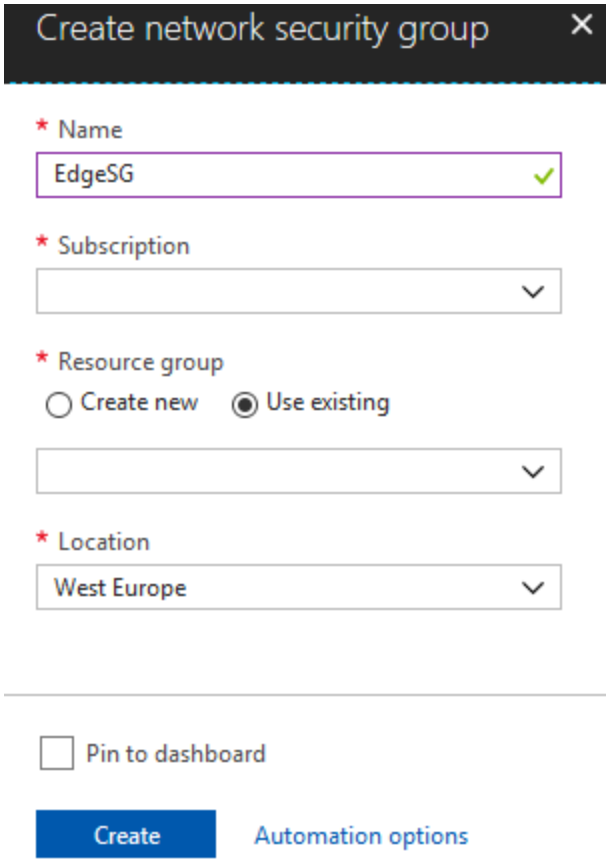

5. Now that your NSG is created, you can select it on the "All resources" list.
6. Edit the **"Inbound security rules"** section of the network security group.

7. Use **"Add"** to enable/disable access to/from your instances on specific ports/source.

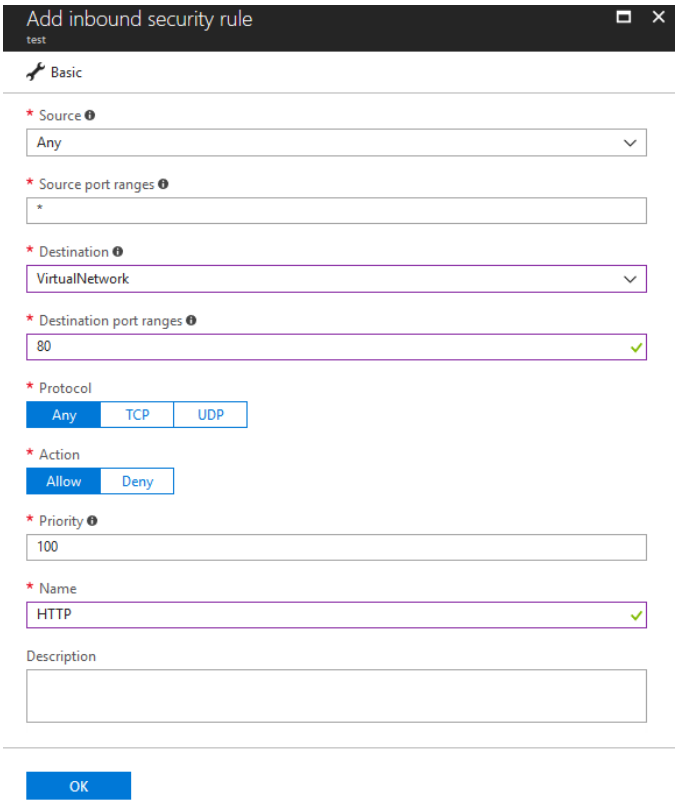

You must add the necessary inbound/outbound rules to the network security groups and the instances assigned to each group will have the required for the group level of connectivity and security. Please refer to the firewall specific information already provided in *SecureTransport Administrator guide*.

You must give unique **Name** and **Priority** to each rule. Please, keep in mind that rules are applied according to their priority. The lower the number, the higher the priority.

Note There are three default rules in each Network security group with the lowest priority that cannot be modified and deleted. If you do not want them to take action, you can create a rule with higher priority that denies all traffic.

Note The rules defined in the following Network Security Groups cover the most basic security scenario. If you would like more restrictive rules, you can add more Inbound and Outbound rules.

Note The rules below use default ports or ports specific for the test setup. Please change/add rules according to your specific setup. Check the *FTP does not work through the firewall* section in the *SecureTransport Administrator Guide* if you want to configure FTP.

SecureTransport Edge Security Group

Allow inbound traffic according to the *Firewall rules* section as described in the *SecureTransport 5.4 Administrator's Guide*.

Inbound traffic - this setup refers to traffic from the load balancer to the Edge serves.

- The **Destination** is a comma separated list of SecureTransport Edge Private IPs.

Note Specifying **Any** for **Source** and the fact that the SecureTransport Edge VMs do not have a public IPs means that only the Load Balancer and VMs from the Virtual Network can access the Edges on the specified ports.

Type	Protocol	Port / Port Range	Source	Description
ST Edges	TCP	80	Any	HTTP
ST Edges	TCP	443	Any	HTTPS
ST Edges	TCP	10022	Any	SSH (SFTP and SCP)
ST Edges	TCP	21	Any	FTP (secure and non-secure) control channel (For secure connections: the firewall must allow bidirectional communication)
ST Edges	TCP	20	Any	FTP (secure and non-secure) active-mode data channel
ST Edges	TCP	User-defined range	Any	FTP (secure and non-secure) passive-mode data channel
ST Edges	TCP	10080	Any	AS2 (non-SSL)
ST Edges	TCP	10443	Any	AS2 (SSL)
ST Edges	TCP	17617	Any	PeSIT (non-SSL)
ST Edges	TCP	17627	Any	PeSIT over secure socket (Transfer CFT compatible)

Type	Protocol	Port / Port Range	Source	Description
ST Edges	TCP	17637	Any	PeSIT over secure socket (CFT compatible)
ST Edges	TCP	19617	Any	PeSIT over pTCP plain socket
ST Edges	TCP	19627	Any	PeSIT over pTCP Secured Socket

Streaming – refers to traffic from the ST servers to the Edge serves.

- The **Destination** is a comma separated list of SecureTransport Edge Private IP addresses.
- The **Source** is comma separated list of SecureTransport Server Private IP addresses.

Type	Protocol	Port / Port Range	Source	Description
ST Edges	TCP	20080	SecureTransport Servers	Streaming HTTP Server
ST Edges	TCP	20022	SecureTransport Servers	Streaming SSH Server
ST Edges	TCP	20021	SecureTransport Servers	Streaming FTP Server
ST Edges	TCP	21080	SecureTransport Servers	Streaming AS2 Server
ST Edges	TCP	20444	SecureTransport Servers	Streaming Administration Tool Server
ST Edges	TCP	27617	SecureTransport Servers	Streaming PeSIT Server

Internal communication – refers to traffic between the Edge serves and traffic from the Administration host.

- The **Destination** is a comma separated list of SecureTransport Edge Private IP addresses.

Type	Protocol	Port / Port Range	Source	Description
ST Edges	TCP	33060	Administration Host Private IP	Database Administration
ST Edges	TCP	33060	SecureTransport Edges	MySQL communication
ST Edges	TCP	22	Administration Host Private IP	SSH for Administration
ST Edges	TCP	444	Administration Host Private IP	Administration Tool (HTTPS)
ST Edges	TCP	444	SecureTransport Edges	Cluster Synchronization
ST Edges	TCP	8005	SecureTransport Edges	Tomcat shutdown
ST Edges	TCP	8006	SecureTransport Edges	AS2 shutdown
ST Edges	TCP	7800-7802	SecureTransport Edges	Hibernate second level cache

SecureTransport Server Security Group

Allow inbound traffic according to the *Firewall rules* section as described in the *SecureTransport 5.4 Administrator's Guide*.

Note Inbound traffic from 8088-8093 range should be allowed for both TCP and UDP protocols from one SecureTransport Server to another (SecureTransport Server Security Group).

- The **Destination** is a comma separated list of SecureTransport Edge Private IPs.
- The **Source** is comma separated list of SecureTransport Edge Private IPs.

Type	Protocol	Port / Port Range	Source	Description
SecureTransport Servers	TCP	80	SecureTransport Servers	SNMP

Type	Protocol	Port / Port Range	Source	Description
SecureTransport Servers	TCP	443	SecureTransport Servers	Cluster cache management
SecureTransport Servers	UDP	10022	SecureTransport Servers	Cluster cache management
SecureTransport Servers	TCP	444	Administration Host Private IP	Administration Tool (HTTPS)
SecureTransport Servers	TCP	44431	SecureTransport Servers	Cluster Listener
SecureTransport Servers	TCP	9999	SecureTransport Servers	TM JMX Port
SecureTransport Servers	SSH	22	SecureTransport Servers	SSH for Administration
SecureTransport Servers	TCP	8005	SecureTransport Servers	Tomcat shutdown port
SecureTransport Servers	TCP	8009	SecureTransport Servers	Tomcat JK connector
SecureTransport Servers	TCP	20444	SecureTransport Servers	Administration Tool
SecureTransport Servers	TCP	7	SecureTransport Servers	Coherence
SecureTransport Servers	TCP	7800-7802	SecureTransport Servers	Hibernate second-level cache
SecureTransport Servers	TCP	*	Any	Deny all TCP Traffic
SecureTransport Servers	UDP	*	Any	Deny all UDP traffic
SecureTransport Servers	Any	*	SecureTransport Servers	Allow all ICMP traffic for Cluster communication

SecureTransport Server nodes communicate via ICMP protocol which is not available for selection as a separate Protocol. If we want to enable it, we must enable the traffic for all protocols which we do not prefer. So, we create two rules that deny all TCP and UDP traffic on all ports with lower priority than the above rules. The rule with the lowest priority allows traffic on all protocols and ports. This means that only the ICMP traffic will be allowed.

External Database Security Group

Allow the following inbound traffic:

Type	Protocol	Port / Port Range	Source	Description
Database VM Private IP	Any	1433	Administration Host Private IP, SecureTransport Servers	Access to the database
Database VM Private IP	Any	3389	Administration Host Private IP	RDP access to the database VM

GlusterFS Security Group

Allow the following inbound traffic:

Type	Protocol	Port / Port Range	Source	Description
GlusterFS Servers	SSH	22	Administration Host Private IP	Administration Host Security Group
GlusterFS Servers	TCP	24007	GlusterFS Servers, SecureTransport Servers	Gluster Daemon
GlusterFS Servers	TCP	111	GlusterFS Servers, SecureTransport Servers	Portmapper
GlusterFS Servers	TCP	49152-49251	GlusterFS Servers, SecureTransport Servers	Each brick for every volume on your host requires its own port

Type	Protocol	Port / Port Range	Source	Description
GlusterFS Servers	TCP	2049	GlusterFS Servers, SecureTransport Servers	NFS

Administration Host Security Group

Allow the following inbound traffic:

Type	Protocol	Port / Port Range	Source	Description
RDP	TCP	3389	Your IP address	Remote connection to the Administration Host

The inbound rule is applicable for Windows instance in our case. Enter the source address from which you plan to connect to the Administration Host.

Access your servers using Administration host

When designing the Administration host for your Azure infrastructure, you should use it only for maintenance and administration. You need to keep it locked down as much as possible and avoid opening unnecessary security holes. You could look into hardening your chosen operating system for even tighter security. In order to minimize security risks, you should start your Administration host instances only when you need access to your servers in Azure.

Here are the basic steps for creating a bastion host for your Azure infrastructure (see section **Launch an instance for the Administration Host**):

1. Launch a VM.
2. Apply your OS hardening as required.
3. Set up the appropriate network security groups (NSG).
4. Implement either SSH-Agent Forwarding (Linux connectivity) or Remote Desktop Gateway (Windows connectivity).

Network Security groups are essential for maintaining tight security and play a big part in making this solution work. First, you need to create a network security group or update an existing one that will be used to allow connectivity from the Administration host for your existing private instances (see the SecureTransport Server Network Security Group in the *Network Security groups* section of the guide). This NSG should only accept SSH or RDP inbound requests from your Administration hosts. Apply this group to all your private instances that require connectivity.

Next, create a network security group to be applied to your Administration host. Inbound and outbound traffic must be restricted at the protocol level as much as possible. The inbound rule base should accept SSH or RDP connections only from the specific IP addresses (usually those of your administrators' work computers). See the *Administration Host Network Security Group* in the *Network Security groups* section of the guide. Your outbound connection should again be restricted to SSH or RDP access to the private instances of your Azure infrastructure. An easy way to do this is to populate the 'Destination' field with the IP of your private instances.

Subnets

After creating a VNet, you can add one or more subnets. Both Public and Private Subnets have outbound access to the Internet by default. If you want inbound Internet traffic to you rVM, you give it a Public IP, no matter of the Subnet. When you create a subnet, you specify the CIDR block for the subnet, which is a subset of the VNet CIDR block.

Learn more about Subnets.

In the following example, the VNet has a single subnet.

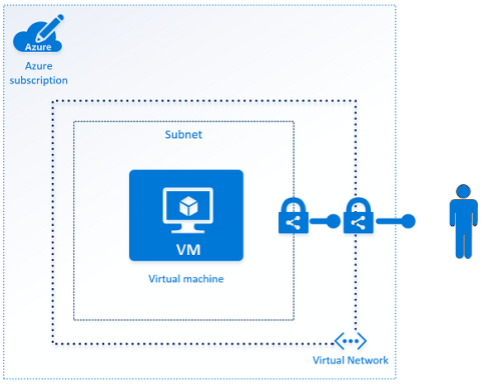

Create subnets

Create three subnets (two private and one public) as follows:

1. Navigate to the Azure Portal, go to "All resources" and select your VNet.
2. Under the "Settings" section, choose **Subnets**.
3. Click **+Subnet**.
4. Fill the following settings.
    - **Name:** specify unique name for each subnet
    - **IPv4 CIDR block:** for each subnet specify a different block
    - **Network Security Group:** None

- **Route table:** None
  For example:
  - 1.0.0/24 - Public
  - 1.1.0/24 - Private1
  - 1.2.0/24 - Private2

    Note: This is an example CIDR block size. You can configure the CIDR block according to your needs.

5. Click **OK** to add your subnet.
6. Repeat the steps for the additional subnets you wish to create.

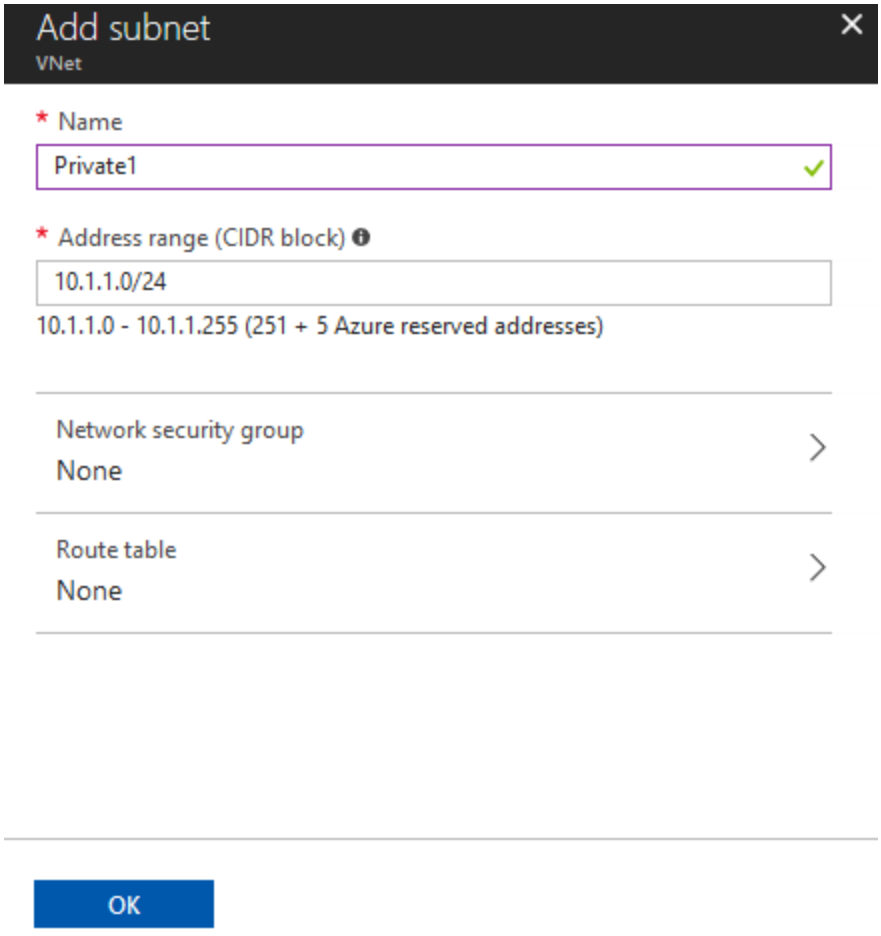

Launch VM instances 6

For the purpose of the described setup, we need seven RHEL instances in total:

- Two SecureTransport Servers – in the first Private Subnet
- Two SecureTransport Edges – in the Public Subnet
- Two GlusterFS Servers – in the first Private Subnet
- One instance to administer and access the above instances – Administration Host in one of the public subnets

To launch a VM instance, follow these steps:

1. Navigate to the Azure Portal and click **Create a resource**.
2. Under the "Azure Marketplace" section, select **Compute**.
3. Under the "Featured" section, select **Red Hat Enterprise Linux 7.2**.
4. Fill in the **Basics**:
    a. Give your VM a unique name
    b. Select disk type: HDD or SSD
    c. Select Authentication type: SSH public key or Password
    d. Select your Subscription
    e. Select existing or create a new Resource group
    f. Select your Location

5. Click **OK**.

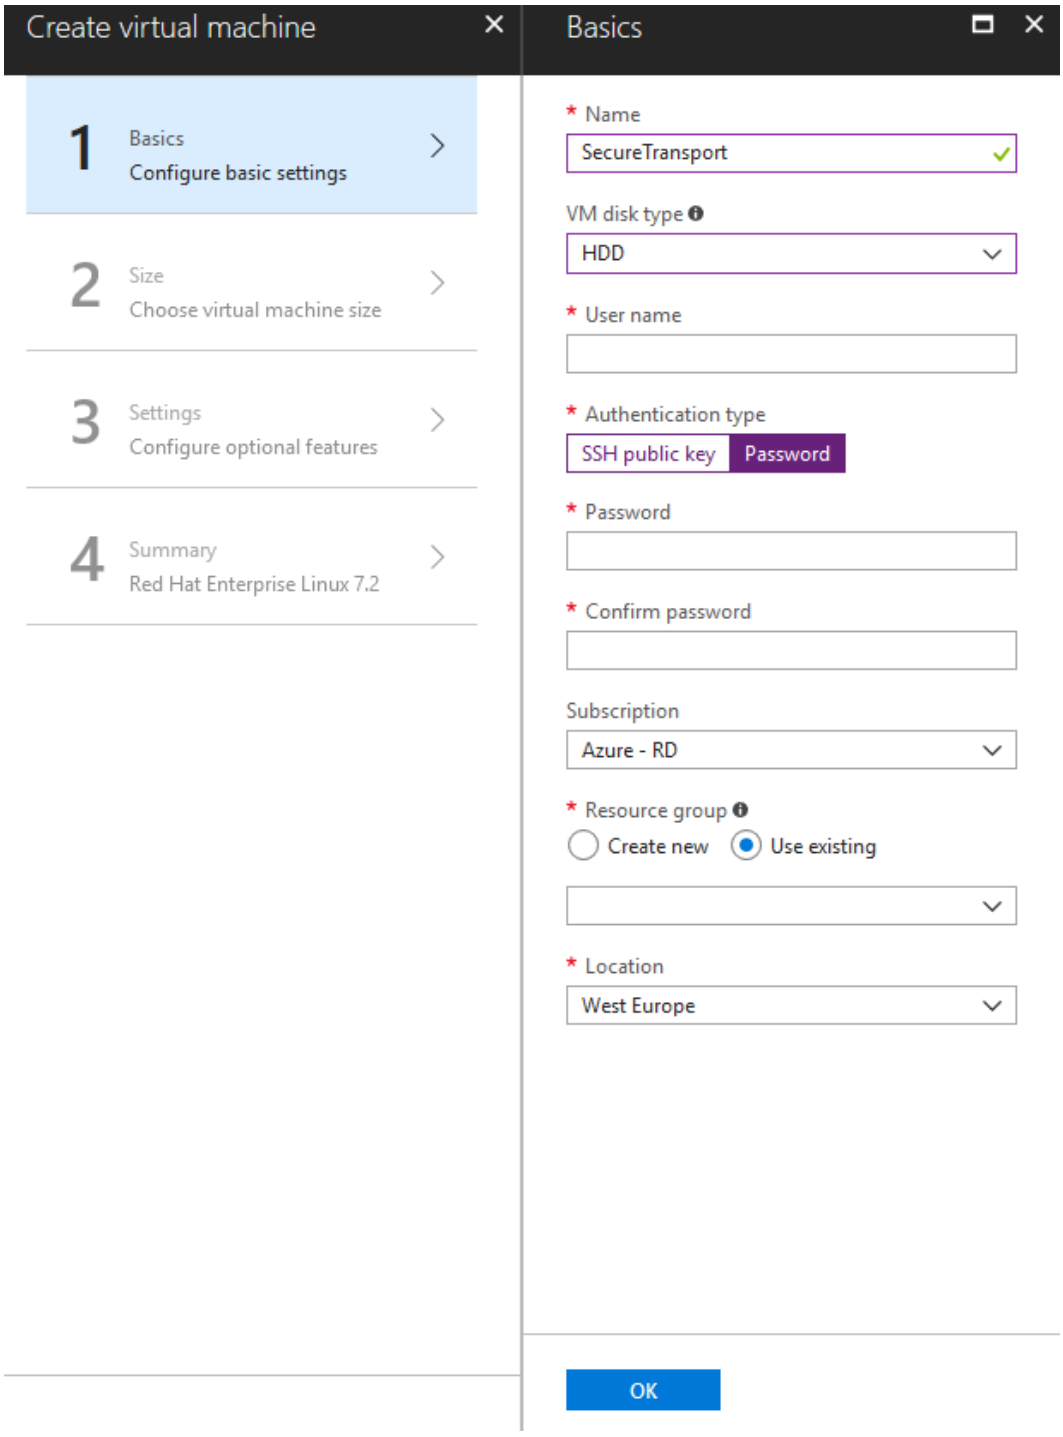

6. Choose a Size according to the *Minimum UNIX hardware requirements* in the SecureTransport 5.4 Installation Guide.

   Note: SecureTransport in Azure has been verified with *D4S_V3 Standard* image for all RHEL instances.

7. Configure Settings
    - **High availability:** Select the previously created Availability Set if you would like to place your VM in it.
    - **Disk type:** Choose between *HDD* and *SSD*.
    - **Use managed disks:** Choose between *Yes* and *No*.
    - **Virtual Network:** Select your previously created Virtual Network or create a new one.
    - **Subnet:** Select your previously created subnet.
    - **Public IP address:** Select *None*.
    - **Network security group:** Select previously created NSG or create a new one.
    - **Extensions:** You can add new features, like configuration management or antivirus protection.
    - **Enable Auto-shutdown:** Choose between *Off* and *On*.
    - **Boot diagnostics:** Choose between *Enabled* and *Disabled*.
    - **Guest OS diagnostics:** Choose between *Enabled* and *Disabled*.
    - **Diagnostics storage account:** Choose existing or create a new one.
8. Review the Summary and click **Create**.

Launch an instance for the Administration Host 7

The Administration Host is an instance of a virtual machine you launch in the public subnet in your VNet. You can use this machine to perform the following activities:

- SecureTransport installations on launched instances
- Configurations on your launched instances
- Maintenance and administration tasks on your instances

Follow the steps for launching RHEL Instance with the following specifics:

1. Choose an Azure image and select the operating system that will be most suitable for your administering needs.
2. Choose an Instance Size that will be most suitable for your administering needs.
3. Configure Instance **Settings -> Virtual network**.
4. Choose your Virtual Network.
5. Configure Instance **Settings -> Subnet**.
6. Choose the public subnet in your VNet.
7. Configure Instance **Settings -> Public IP address** (depending on the connectivity you would like to have to this host).
8. Configure Instance **Settings -> Network security group**.
9. Assign the instance with the previously created Administration Host Network Security Group as described in *Network Security Groups*.

Launch SecureTransport Edge instances

Follow the steps for launching RHEL Instance with the following specifics:

1. Configure Instance **Settings-> High availability**.
2. Select the previously created SecureTransport Edges Availability Set.
3. Configure Instance **Settings-> Subnet**.
4. Choose the public subnet in your VNet.
5. Configure Instance **Settings** -> **Public IP address:**
6. Configure Instance **Settings** -> **Network security group**.
7. Assign both SecureTransport Edges to the previously created SecureTransport Edge Network security group as described in *Network Security Groups* .

Launch SecureTransport Server instances

Follow the steps for launching RHEL Instance with the following specifics:

1. Configure Instance **Settings-> High availability.**
2. Select the previously created SecureTransport Servers Availability Set.
3. Configure Instance **Settings-> Subnet**.
4. Choose the first private subnet in your VNet.
5. Configure Instance **Settings** -> **Public IP address:**
6. Configure Instance **Settings** -> Network security group.
7. Assign both SecureTransport Servers to the previously created Security Group as described in *Network Security Groups*.

Set up GlusterFS servers

Follow the steps for launching RHEL Instances with the following specifics:

1. Configure Instance **Settings-> High availability**.
2. Select the previously created GlusterFS Availability Set.
3. Configure Instance **Settings-> Storage**.
4. Select *Yes* for **Use managed disks**.
5. Configure Instance **Settings-> Subnet**.
6. Choose the first private subnet in your VNet.
7. Configure Instance **Settings** -> **Public IP address**:
8. Configure Instance **Settings** -> Network security group.
9. Assign both GlusterFS servers to the previously created GlusterFS Network Security Group as described in *Network Security Groups* .

Attach additional volumes

You will need to attach additional volumes to the GlusterFS instances, as they need two or more virtual disks.

Create two Managed Disks:

1. Navigate to Azure Portal and click **Create a resource**.
2. in the search field, type "Managed Disks" and click the search icon.
3. Fill in the Settings.

For each GlusterFS Server do the following:

1. Navigate to **Azure Portal ->All resources**.
2. Select your GlusterFS Instance.
3. Under**Settings** section select**Disks -> +Add data disk**.
4. Select the previously created data disk.
5. Click **Save**.

Install GlusterFS

Follow the instructions for installing GlusterFS in the GlusterFS Documentation.

Configure MS SQL Server 8

SecureTransport in Microsoft Azure has been verified with database Microsoft SQL Server 2016 installed on a Windows Virtual machine placed in the second private subnet in a Virtual Network.

Follow these steps to set up a Microsoft SQL Server on a Windows VM instance:

1. Follow the steps for launching a VM instance with the following specifics:
    a. Select Windows Server 2016 for image.
    b. Select an appropriate size - for example SecureTransport with D4_V3 Standard size for VM instance.
    c. Configure Instance **Settings**-> **Subnet**.
    d. Choose the second private subnet in your VNet.
    e. Configure Instance **Settings** -> **Public IP address:**
    f. Configure Instance **Settings** -> Network security group.
    g. Assign the Windows VM to the previously created Database Network security group as described in Network *Security Groups*
2. Install Microsoft SQL Server on the Windows instance.
3. Set up the database as described in *Database requirements* in SecureTransport 5.4 Installation Guide.

Note: For a more fault-tolerant setup, you can launch two Windows Servers in a Database Availability Set and configure a Microsoft SQL Cluster.

Configure a Point-to-Site connection to your VNet 9

A Point-to-Site (P2S) configuration allows you to create a secure connection from a client computer to your virtual network. So, you need a P2S SSTP tunnel for a client to connect to your VM.

On the Azure portal, open your newly created VM and follow the process as defined in the detailed documentation as provided by Microsoft.

First-timers must go through all steps as defined in the Microsoft documentation:

1. Create a virtual network.
2. Specify address space and subnets.
3. Add a gateway subnet.
4. Specify a DNS server (optional).
5. Create a virtual network gateway.
6. Generate certificates:
    - Obtain the *.cer* file for the root certificate: generate a network-side certificate and paste the public part in the Gateway >P2S configuration.
    - Generate a client certificate.
7. Add the client address pool.
8. Upload the root certificate *.cer*
9. Install the VPN client configuration package.
    - Download the client configuration package: for example, VPN client (e.g. x64).
    - Install the client configuration package.
10. Install the client certificate - generate a client certificate for secure connection to VPN.
11. Connect to Azure - verify that you can connect to your VM using its private IP address.
12. Verify your connection.

Set up Enterprise Cluster with streaming 10

Prerequisites

The following list outlines the prerequisites necessary for the described SecureTransport Enterprise cluster setup.

1. Two RHEL Instances for SecureTransport Servers.
2. Two RHEL Instances for SecureTransport Edges.
3. Two RHEL Instances with GlusterFS Servers installed.
4. One Administration Host (optional).
5. One Microsoft SQL Server.

For correct setup of all your instances, please refer to the *Prerequisites -> UNIX-based platforms* topic in the *SecureTransport5.4 Installation Guide*.

Note Make sure to have the following prerequisite installation package added:
ld-linux.so.2 library package

Install SecureTransport

1. Install SecureTransport Servers as described in the SecureTransport Installation Guide: *Install SecureTransport Server in an Enterprise Cluster or to use an external database*.
2. Install SecureTransport Edges as described in the SecureTransport Installation Guide: *Install SecureTransport Server to use the embedded database* and in the SecureTransport Administrator's Guide: *SecureTransport Edge synchronization*
3. Install GlusterFS client on both SecureTransport Servers and mount the GlusterFS volume on the client side.

For further reference, see GlusterFS Documentation.

Set up Basic Load Balancer 11

Azure Load Balancer delivers high availability and network performance to your applications. It is a Layer 4 (TCP, UDP) load balancer that distributes incoming traffic among healthy instances of services defined in a load-balanced set. Azure Load Balancer supports two different types: Basic and Standard. We have verified Basic Load Balancer.

When you create a load balancer in a VNet, you can make it an internal load balancer or an Internet-facing load balancer. You create an Internet-facing load balancer.

When you create your load balancer, you configure Frontend IP, Load Balancing rules, health probes, and register back-end instances. You configure a Load Balancing rule by specifying a protocol and a port for front-end (client to load balancer) connections, and a protocol and a port for back-end (load balancer to back-end instances) connections. You can configure multiple rules for your load balancer.

Follow these steps to create a Basic Load Balancer redirecting requests to SecureTransport Edge instances in Azure VNet:

1. Log in to **Azure portal -> Create a resource**.
2. Under **Azure Marketplace** section select **Networking**.
3. Under **Featured** section select **Load Balancer**.
4. Fill in the Basic settings:
    a. Enter Load Balancer name.
    b. Select **Type**: *Public*.
    c. **Choose a Public IP address** - Select existing or create a new one.
    d. Select your **Subscription**, **Resource group** and **Location**.
    e. Click **Create**.

After creating the Load Balancer, you must configure Health probes, Backend pools and Load Balancing rules.

Configure Health probes

1. Navigate to your Load balancer settings -> **Health probes**.
2. Click **+Add**.
3. Give your probe a **Name**, **Protocol**, **Port**, **Interval** and **Unhealthy threshold**.

Note: Create a Health probe for each Load Balancing rule specifying the port and protocol.

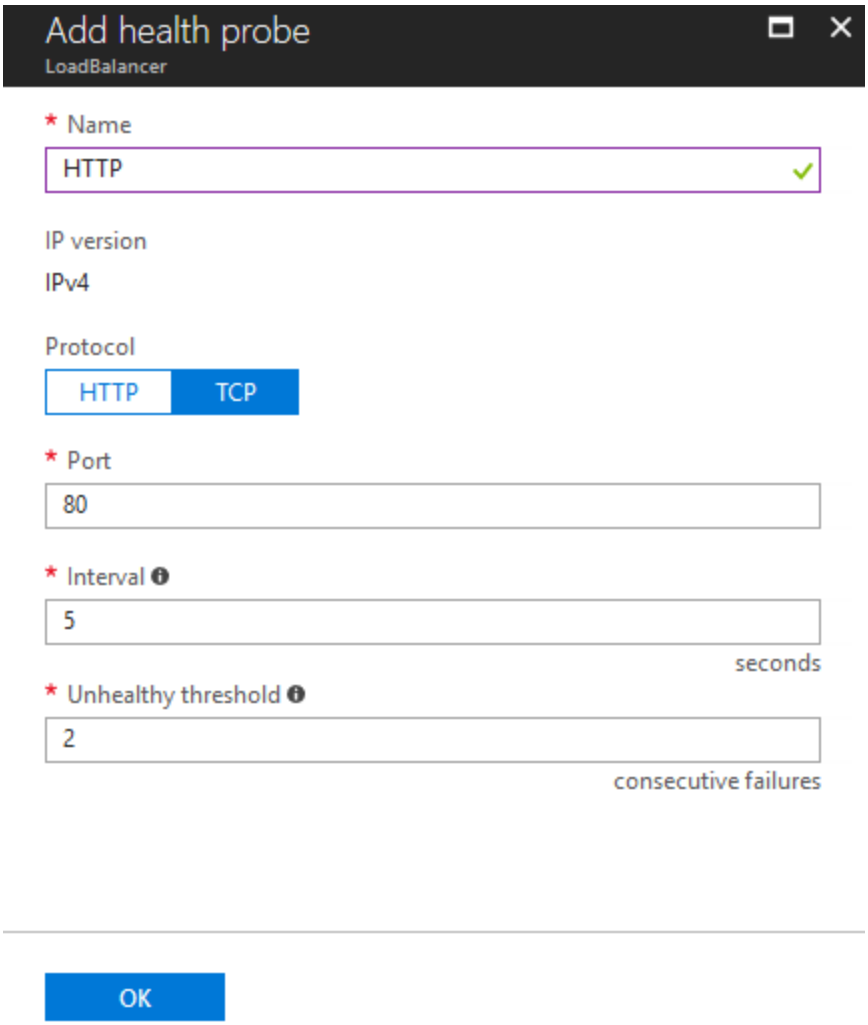

Add a Backend pool

1. Navigate to your Load balancer settings -> **Backend pools**.
2. Click **+Add**.
3. Fill in the Backend pool settings:
    1. Give it a unique name.
    2. Associate the Load balancer to the SecureTransport Edge Availability Set.
    3. Select both SecureTransport Edge Instances as a **Target virtual machine**.
    4. Select each SecureTransport Edge **Network IP Configuration**.
4. Click **OK**.

Add Load Balancing rules

1. Navigate to your Load balancer settings -> **Load Balancing rules**.
2. Click **+Add**.
3. Fill in the Load Balancing rule settings:
    a. Provide a unique name.
    b. Select an **IP version**.
    c. Select **Frontend IP address** from the drop-down menu.
    d. Select a **Protocol**.
    e. Select a **Port** on which the Load Balancer will be accessed.
    f. Select a **Backend Port** to route traffic.
    g. Select a **Backend pool** - target of the load balancer traffic.
    h. Select a **Health probe** for the Load balancing rule.
    i. Select **Session persistence:** *Client IP and protocol*.
    j. Configure **Idle timeout (minutes)**.
    k. Select **Floating IP**.

Note: For each Load balancing rule select the Health probe that corresponds to the rules protocol and port.

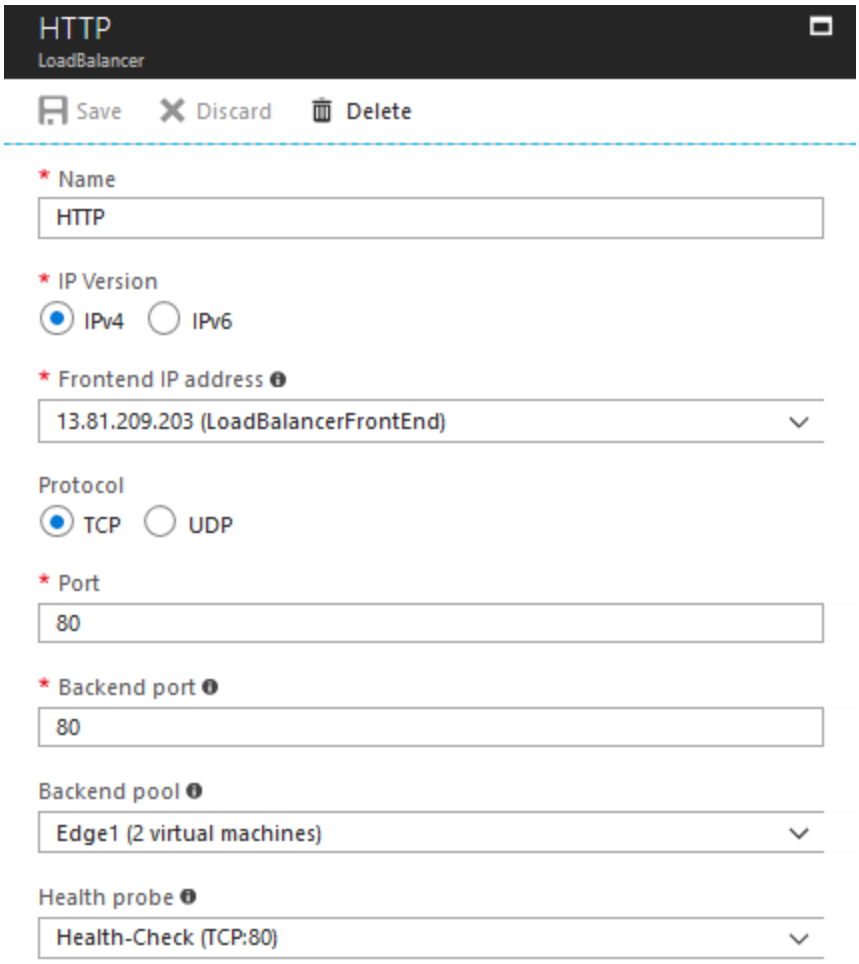

Note The Load balancing rules in the screenshot are configured with default ports or ports specific for the test setup. Please change/add rules according to your specific setup.

Note Check the **FTP does not work through the firewall** section in the *SecureTransport 5.4 Administrator's guide* if you want to configure FTP listeners.

Now, your load balancer is ready to fetch requests and distribute them among your SecureTransport Edge instances in the SecureTransport Edge Availability Set.

If you would like to use a friendly DNS name to access your load balancer, instead of the default DNS name automatically assigned to your load balancer, you can create a custom domain name.

Criteria for a successful setup

12

Make sure that all the listed criteria meet the requirements for a successful SecureTransport:

1. Both SecureTransport Server and Edge Clusters are synchronized and work as expected.
2. Successful end-user login over the desired protocols via the Load Balancer or SecureTransport Edge.
3. Successful file operations and transfers over the desired protocols.