Microsoft

Microsoft Azure Fundamentals

THIRD EDITION

Exam Ref AZ-900

Jim Cheshire

Microsoft

Exam Ref AZ-900 Microsoft Azure Fundamentals

Third Edition

Jim Cheshire

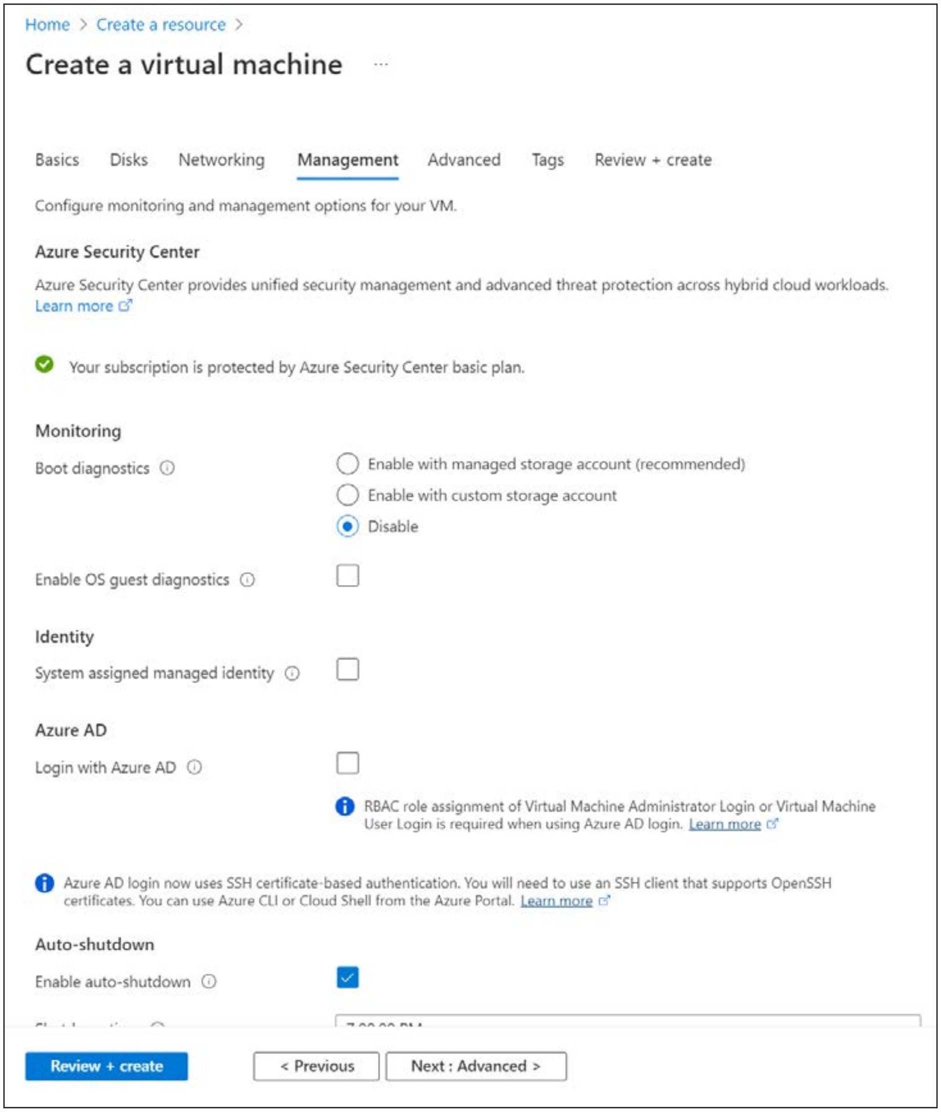

Home > Create a resource >

Create a virtual machine ...

Configure monitoring and management options for your VM.

Azure Security Center

Azure Security Center provides unified security management and advanced threat protection across hybrid cloud workloads. Learn more

✓ Your subscription is protected by Azure Security Center basic plan.

Monitoring

Boot diagnostics ⓘ
○ Enable with managed storage account (recommended)
○ Enable with custom storage account
● Disable

Enable OS guest diagnostics ⓘ □

Identity

System assigned managed identity ⓘ □

Azure AD

Login with Azure AD ⓘ □

ⓘ RBAC role assignment of Virtual Machine Administrator Login or Virtual Machine User Login is required when using Azure AD login. Learn more

ⓘ Azure AD login now uses SSH certificate-based authentication. You will need to use an SSH client that supports OpenSSH certificates. You can use Azure CLI or Cloud Shell from the Azure Portal. Learn more

Auto-shutdown

Enable auto-shutdown ⓘ ☑

FIGURE 2-14 Virtual machine management settings

As your VM is being deployed, you'll see the status displayed in the Azure portal, as shown in Figure 2-15. You can see the Azure resources that are created to support your VM. You can see the resource name, the resource type, and the status of each resource.

Once all the resources required for your VM are created, your VM will be considered fully deployed. You'll then be able to click the **Go To Resource** button to see the management interface for your VM in the Azure portal, as shown in Figure 2-16.

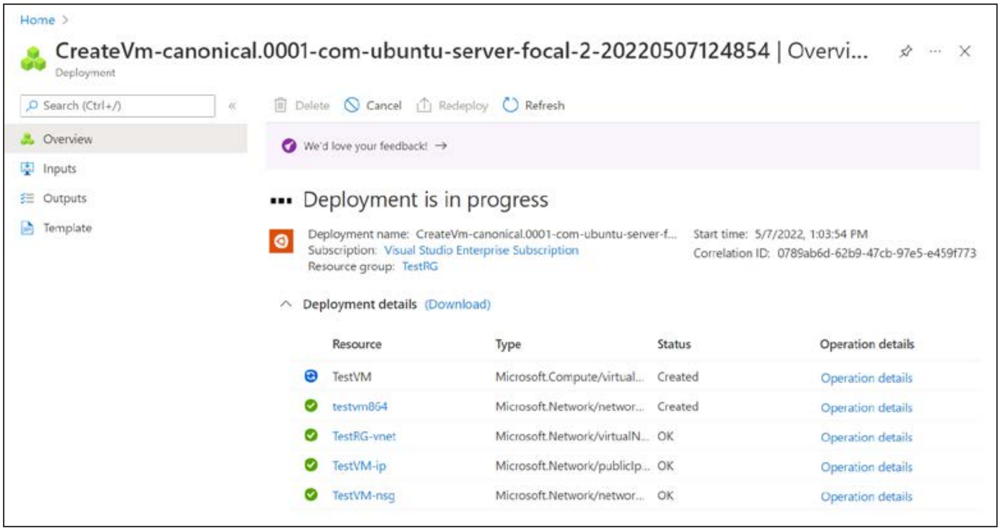

CreateVm-canonical.0001-com-ubuntu-server-focal-2-20220507124854 | Overvi...

Deployment is in progress

Deployment name: CreateVm-canonical.0001-com-ubuntu-server-f...
Subscription: Visual Studio Enterprise Subscription
Resource group: TestRG

Start time: 5/7/2022, 1:03:54 PM
Correlation ID: 0789ab6d-62b9-47cb-97e5-e459f773

Deployment details (Download)

Resource	Type	Status	Operation details
TestVM	Microsoft.Compute/virtual...	Created	Operation details
testvm864	Microsoft.Network/networ...	Created	Operation details
TestRG-vnet	Microsoft.Network/virtualN...	OK	Operation details
TestVM-ip	Microsoft.Network/publicIp...	OK	Operation details
TestVM-nsg	Microsoft.Network/networ...	OK	Operation details

FIGURE 2-15 Virtual machine deployment

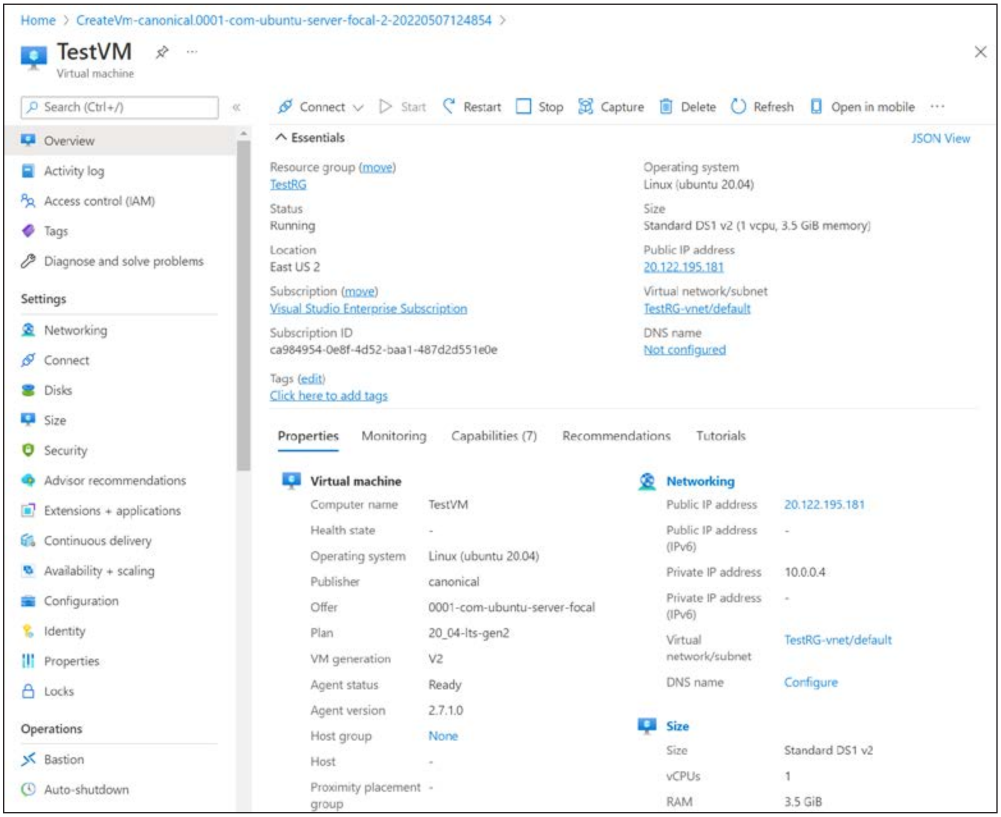

TestVM

Essentials

Resource group (move)
TestRG

Status
Running

Location
East US 2

Subscription (move)
Visual Studio Enterprise Subscription

Subscription ID
ca984954-0e8f-4d52-baa1-487d2d551e0e

Tags (edit)
Click here to add tags

Operating system
Linux (ubuntu 20.04)

Size
Standard DS1 v2 (1 vcpu, 3.5 GiB memory)

Public IP address
20.122.195.181

Virtual network/subnet
TestRG-vnet/default

DNS name
Not configured

Virtual machine

Computer name	TestVM
Health state	-
Operating system	Linux (ubuntu 20.04)
Publisher	canonical
Offer	0001-com-ubuntu-server-focal
Plan	20_04-lts-gen2
VM generation	V2
Agent status	Ready
Agent version	2.7.1.0
Host group	None
Host	-
Proximity placement group	-

Networking

Public IP address	20.122.195.181
Public IP address (IPv6)	-
Private IP address	10.0.0.4
Private IP address (IPv6)	-
Virtual network/subnet	TestRG-vnet/default
DNS name	Configure

Size

Size	Standard DS1 v2
vCPUs	1
RAM	3.5 GiB

FIGURE 2-16 Viewing a virtual machine

The new VM is a guest on a physical computer in an Azure datacenter. In that datacenter is a physical rack of computer servers, and our VM is hosted on one of those servers. The host computer is managed by Microsoft, but you manage the VM because this is an IaaS offering in Azure.

NOTE **VMs AND BILLING**

You are charged for Azure VMs as long as they are running, and using the default settings as we have here led to a few expensive options. To stop billing for this VM, click the Stop button at the top of the screen shown in Figure 2-16. Azure will save the current state of the VM, and billing will stop. You won't be able to use the VM while it's in a stopped state, but you will also avoid the billing of that VM. Keep in mind that unless you reserve the IP address for your VM, your IP address will likely change the next time you start it.

You can also stop a VM from within the guest operating system on the VM, but when you do that, you will still be charged for the resources the VM uses because it's still allocated to you. That means you'll still incur charges for managed disks and other resources. Deleting the `TestRG` **resource group will ensure you aren't charged for the VM.**

As of right now, this VM is susceptible to downtime due to three types of events: *planned maintenance, unplanned maintenance*, and *unexpected downtime*.

Planned maintenance refers to planned updates that Microsoft makes to the host computer. This includes things like operating system updates, driver updates, and so on. In many cases, updates won't affect your VM, but if Microsoft installs an update that requires a reboot of the host computer, your VM will be down during that reboot.

Azure has underlying systems that constantly monitor the health of computer components. If one of these underlying systems detects that a component within the host computer might fail soon, Azure will flag the computer for unplanned maintenance. In an unplanned maintenance event, Azure will attempt to move your VM to a healthy host computer. When it does this, it preserves the state of the VM, including what's in memory and any files that are open. It only takes Azure a short time to move the VM, during which time it's in a paused state. In a case where the move operation fails, the VM will experience unexpected downtime.

To ensure reliability when a failure occurs in a rack within the Azure datacenter, you can (and you should) take advantage of a feature called *availability sets*. *Availability sets* protect you from maintenance events and downtime caused by hardware failures. To do that, Azure creates some underlying entities in an availability set called *update domains* and *fault domains*. (In order to protect yourself in the event of maintenance events or downtime, you must deploy at least two VMs into your availability set.)

Fault domains are a logical representation of the physical rack in which a host computer is installed. By default, Azure assigns two fault domains to an availability set. If a problem occurs in one fault domain (one computer rack), the VMs in that fault domain will be affected, but VMs in the second fault domain will not. This protects you from unplanned maintenance events and unexpected downtime.

Update domains are designed to protect you from a situation where the host computer is being rebooted. When you create an availability set, Azure creates five update domains by default. These update domains are spread across the fault domains in the availability set. If a reboot is required on computers in the availability set (whether host computers or VMs within the availability set), Azure will only reboot computers in one update domain at a time and it will wait 30 minutes for computers to recover from the reboot before it moves on to the next update domain. Update domains protect you from planned maintenance events.

Figure 2-17 shows a representation of an availability set containing five VMs. There are two fault domains and three update domains. When VMs were created in this availability set, they were assigned as follows:

- The first VM is assigned Fault Domain 0 and Update Domain 0.
- The second VM is assigned Fault Domain 1 and Update Domain 1.
- The third VM is assigned Fault Domain 0 and Update Domain 2.
- The fourth VM is assigned Fault Domain 1 and Update Domain 0.
- The fifth VM is assigned Fault Domain 0 and Update Domain 1.

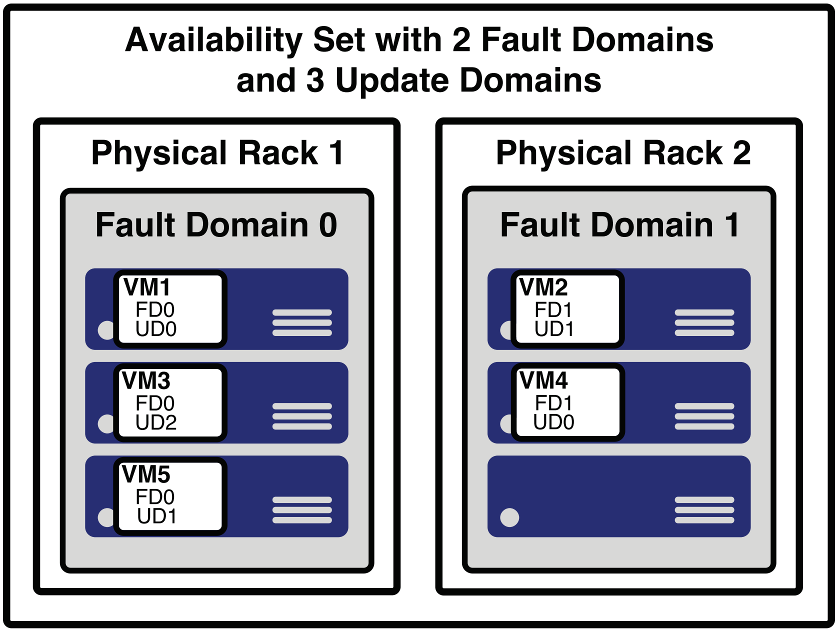

FIGURE 2-17 A representation of an availability set

You can verify the placement of fault domains and update domains by creating five VMs in an availability set with two fault domains and three update domains. If you then look at the availability set created in the Azure portal, as shown in Figure 2-18, you can see the same configuration depicted in Figure 2-17.

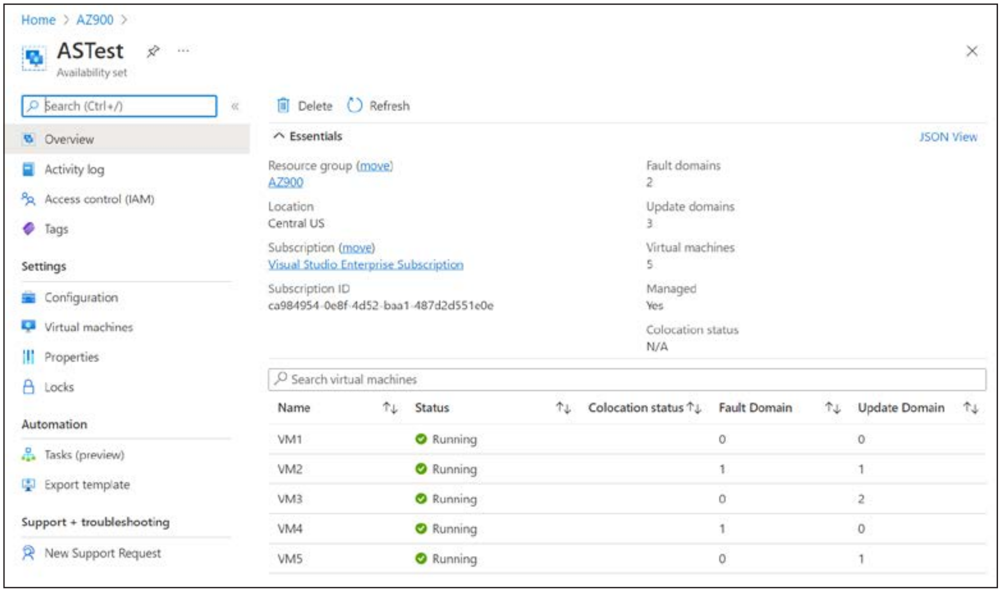

Resource group (move)
AZ900

Location
Central US

Subscription (move)
Visual Studio Enterprise Subscription

Subscription ID
ca984954-0e8f-4d52-baa1-487d2d551e0e

Fault domains
2

Update domains
3

Virtual machines
5

Managed
Yes

Colocation status
N/A

Name	Status	Colocation status	Fault Domain	Update Domain
VM1	✓ Running		0	0
VM2	✓ Running		1	1
VM3	✓ Running		0	2
VM4	✓ Running		1	0
VM5	✓ Running		0	1

FIGURE 2-18 An availability set in the Azure portal showing fault domains and update domains

Notice in Figure 2-18 that the availability set is named `ASTest`. In this availability set, we run five VMs that are all running a web server and host the website for an application. Suppose you need a database for this application, and you want to host that database on VMs as well. In that situation, you would want to separate the database VMs into their own availability set. As a best practice, you should always separate your workloads into separate availability sets.

Availability sets certainly provide a benefit in protecting from downtime in certain situations, but they also have some disadvantages. First, every machine in an availability set must be explicitly created. While you can use an ARM template to deploy multiple virtual machines in one deployment, you still must configure those machines with the software and configuration necessary to support your application.

An availability set also requires that you configure something in front of your VMs that will handle the distribution of traffic to those VMs. For example, if your availability set is servicing a website hosted on the VMs, you'll need to configure a load balancer that will handle the job of routing users of your website to the VMs that are running it.

Another disadvantage of availability sets relates to cost. In a situation where your VM needs to be changed often based on things like load on the application, you might find yourself paying for many more VMs than you need.

Azure offers another feature for VMs called *scale sets* that solves these problems nicely. When you create a scale set, you tell Azure what operating system you want to run and then you tell Azure how many VMs you want in your scale set. You have many other options such as creating a load balancer or gateway and so forth. Azure will create as many VMs as you specify (up to 1,000) in one easy step.

MORE INFO **USING A CUSTOM IMAGE**

The default set of templates for VMs are basic and include only the operating system. However, you can create a VM, install all the necessary components you need (including your own applications), and then create an image that can be used when creating scale sets.

For more information on using custom images, see *https://bit.ly/az900-customvmimages*.

Scale sets are deployed in availability sets automatically, so you automatically benefit from multiple fault domains and update domains. Unlike VMs in an availability set, however, VMs in a scale set are also compatible with availability zones, so you are protected from problems in an Azure datacenter.

As you might imagine, you can also scale a scale set in a situation where you need more or fewer VMs. You might start with only one VM in a scale set, but as the load on that VM increases, you might want to automatically add additional VMs. Scale sets provide that functionality by using Azure's auto-scale feature. You define scaling rules that use metrics like CPU, disk usage, network usage, and so forth. You can configure when Azure should add additional instances and when it should scale back and deallocate instances. This is a great way to ensure availability while reducing costs by taking advantage of the elasticity that auto-scale provides.

MORE INFO **SCALING AND AVAILABILITY SETS**

Before the introduction of scale sets, you had the ability to configure auto-scale rules for an availability set. You'll probably still see third-party documentation and training that talks about scaling availability sets, but that functionality has been replaced with scale sets.

Microsoft guarantees an SLA of 99.95 percent when you use a multi-VM deployment scenario, and for most production scenarios, a multi-VM deployment is preferred. However, if you use a single-instance VM, and you use premium storage, Microsoft guarantees a 99.9 percent SLA. Premium storage uses solid-state drives (SSDs) that are located on the same physical server that is hosting the VM for enhanced performance and uptime.

An Azure VM is a great choice if you need the flexibility to run your own applications in a virtualized environment with the greatest amount of flexibility, but many businesses today are turning to desktop virtualization instead of running multiple VMs.

In a desktop virtualization model, a business installs an operating system and applications on one central server. The desktop virtualization infrastructure makes it possible for employees to access the operating system and applications from virtually any device, provided it has access to the network. The OS and applications aren't downloaded to the employee's device. Instead, the employee uses the applications in a virtualized environment that makes it feel like the applications are running locally.

Most businesses have applications that all their employees need to use. For example, employees might need access to Microsoft Word, Microsoft Excel, Microsoft Outlook, and so on. In many situations, businesses fill this need by purchasing a Microsoft Office license for all employees and installing Office apps on each computer.