

Contents

Azure VMware Solution

Overview

About Azure VMware Solution

What's new

Get started

1 - Plan the deployment

2 - Deploy Azure VMware Solution

3 - Connect to on-premises environment

4 - Install and activate VMware HCX in Azure VMware Solution

5 - Configure on-premises HCX Connector

6 - Uninstall VMware HCX in Azure VMware Solution

Tutorials

1 - Network planning checklist

2 - Create a private cloud

3 - Configure networking

4 - Access a private cloud

5 - Create an NSX-T Data Center network segment

6 - Peer on-premises to private cloud

7 - Scale in a private cloud

8 - Delete a private cloud

Concepts

Access and identity

API Management

Hub and spoke

Internet connectivity design considerations

Network design considerations

Networking and interconnectivity

Private clouds and clusters

Run command

Security recommendations

Storage

How-to guides

Azure native integration

Protect, monitor, and manage VMs

Integrate Microsoft Defender for Cloud

Protect web apps with Azure Application Gateway

Deploy Traffic Manager to balance workloads

Configure alerts and work with metrics

Configure customer-managed key encryption at rest

Manage and protect VMs on Azure NetApp Files

Attach Azure NetApp Files datastores to Azure VMware Solution hosts

Attach Azure NetApp Files to Azure VMware Solution VMs

Backup with Azure Backup Server

Set up Backup Server

Back up private cloud VMs with Backup Server

Disaster recovery using Azure Site Recovery

Prepare Azure

Prepare Azure VMware Solution

Set up replication

Run a disaster recovery drill

Fail over to Azure

Reprotect VMs

Fail back from Azure

Create an Azure VMware Solution assessment

Create a placement policy

Configure GitHub Enterprise Server

Configure Identity

Configure external identity source for vCenter Server

Configure external identity source for NSX-T Data Center

Configure Internet connectivity

Enable Managed SNAT for Azure VMware Solution Workloads

Enable Public IP to the NSX-T Data Center Edge for Azure VMware Solution

Disable Internet access or enable a default route

Configure networking

Configure DHCP server or relay

Configure DHCP on L2 stretched networks

Configure DNS forwarder

Configure HCX network extension

HCX Mobility Optimized Networking (MON) guidance

Configure NSX-T Data Center network components

Configure port mirroring

Configure a site-to-site VPN in vWAN

Configure storage policies

Configure VMware syslogs

Configure Windows Server Failover Cluster

Connect multiple private clouds in same region

Deploy Arc

Deploy disaster recovery

Deploy disaster recovery using JetStream DR software

Deploy VMware HCX for disaster recovery

Deploy VMware SRM for disaster recovery

Deploy Zerto disaster recovery

Deploy virtual desktops

Deploy Horizon on Azure VMware Solution

Deploy Citrix on Azure VMware Solution

Deploy vSAN stretched clusters (Preview)

Move resources

Move Azure VMware Solution subscriptions

Move Azure VMware Solution resources across regions

Request host quota for Azure VMware Solution

Rotate cloudadmin credentials

Save costs with a reserved instance

Partner solution ecosystem

Operating system support for VMs

Backup solutions for VMs

Disaster recovery solutions for VMs

Migration solutions for VMs

Security solutions for Azure VMware Solution

Application performance monitoring solutions for Azure VMware Solution

Bitnami appliance deployment

VMware solutions

Configure HCX network extension

Configure HCX network extension high availability

Configure vRealize Operations

Deploy VMware Horizon

Deploy VMs from the content library

Enable HCX access over the internet

Enable SQL Azure hybrid benefit for Azure VMware Solution (Preview)

Enable VMware Cloud director service with Azure VMware solution (Preview)

Upgrade HCX on Azure VMware Solution

Resources

Learning paths

Introduction to Azure VMware Solution

Run VMware vSphere workloads on Azure VMware Solution

Prepare to migrate VMware vSphere workloads to Azure by deploying Azure VMware Solution

Migrate VMware vSphere workloads from on-premises to Azure VMware Solution

Deploy disaster recovery using VMware Site Recovery Manager and Azure VMware Solution

Hands-on Labs

Azure VMware Solution Private Cloud Deployment and Connectivity

Azure VMware Solution Workload Migration with VMware HCX

Disaster Protection with Azure VMware Solution and VMware Site Recovery

Cloud Adoption Framework

Azure VMware Solution landing zone accelerator

Enterprise-Scale for Azure VMware Solution repository

[Azure VMware Solution Landing Zone Accelerator | Video](#)

Videos

[Extend to the Cloud with Azure VMware Solution](#)

[Run your VMware workloads natively on Azure with Azure VMware Solution | Azure Friday](#)

[Azure VMware Solution Technical Overview Series | VMware](#)

[Azure VMware Solution Deployment Deep Dive Series | VMware](#)

Regional availability

Pricing

[Azure VMware Solution pricing](#)

[Azure pricing calculator](#)

SLA

[Azure VMware Solution roadmap](#)

Troubleshooting

[Open a support request for deployment failures](#)

[VMware tools vix error code = 21009](#)

FAQ

Videos

[Playlist](#)

[Integration with Azure Service](#)

[Integration with Azure Application Gateway](#)

[Migration and capacity planning](#)

Reference

[REST API](#)

[Azure CLI](#)

[Azure PowerShell](#)

[Resource Manager template](#)

[Terraform Azure provider](#)

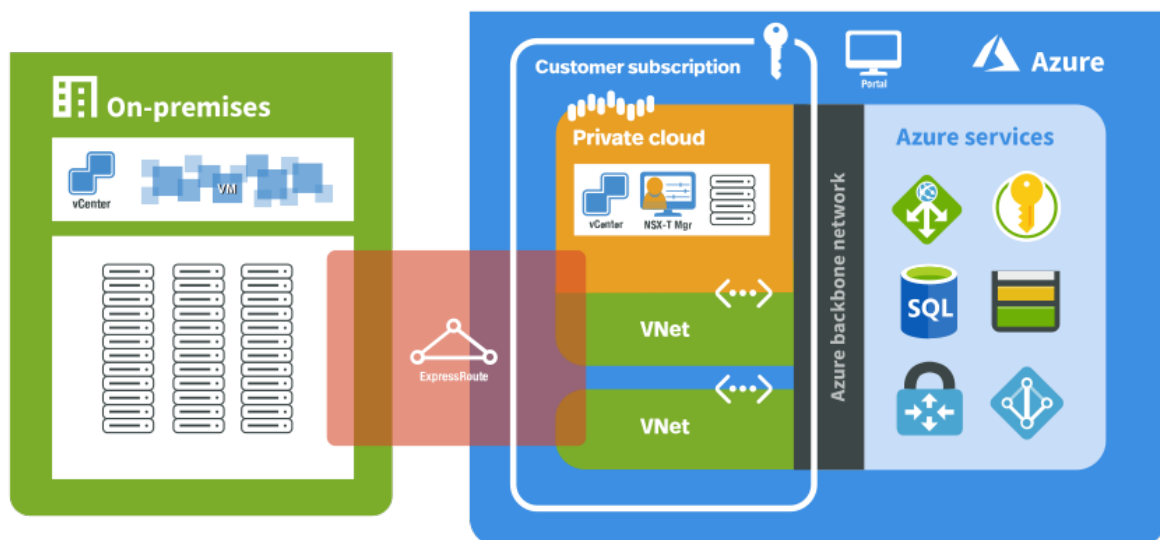
What is Azure VMware Solution?

12/16/2022 • 9 minutes to read • [Edit Online](#)

Azure VMware Solution provides you with private clouds that contain VMware vSphere clusters built from dedicated bare-metal Azure infrastructure. The minimum initial deployment is three hosts, but more hosts can be added one at a time, up to a maximum of 16 hosts per cluster. All provisioned private clouds have VMware vCenter Server, VMware vSAN, VMware vSphere, and VMware NSX-T Data Center. As a result, you can migrate workloads from your on-premises environments, deploy new virtual machines (VMs), and consume Azure services from your private clouds. For information about the SLA, see the [Azure service-level agreements](#) page.

Azure VMware Solution is a VMware validated solution with ongoing validation and testing of enhancements and upgrades. Microsoft manages and maintains the private cloud infrastructure and software. It allows you to focus on developing and running workloads in your private clouds to deliver business value.

The diagram shows the adjacency between private clouds and VNets in Azure, Azure services, and on-premises environments. Network access from private clouds to Azure services or VNets provides SLA-driven integration of Azure service endpoints. ExpressRoute Global Reach connects your on-premises environment to your Azure VMware Solution private cloud.



AV36P and AV52 node sizes available in Azure VMware Solution

The new node sizes increase memory and storage options to optimize your workloads. The gains in performance enable you to do more per server, break storage bottlenecks, and lower transaction costs of latency-sensitive workloads. The availability of the new nodes allow for large latency-sensitive services to be hosted efficiently on the Azure VMware Solution infrastructure.

AV36P key highlights for Memory and Storage optimized Workloads:

- Runs on Intel® Xeon® Gold 6240 Processor with 36 Cores and a Base Frequency of 2.6Ghz and Turbo of 3.9Ghz.
- 768 GB of DRAM Memory
- 19.2 TB Storage Capacity with all NVMe based SSDs (With Random Read of 636500 IOPS and Random Write

of 223300 IOPS)

- 1.5TB of NVMe Cache

AV52 key highlights for Memory and Storage optimized Workloads:

- Runs on Intel® Xeon® Platinum 8270 with 52 Cores and a Base Frequency of 2.7Ghz and Turbo of 4.0Ghz.
- 1.5 TB of DRAM Memory
- 38.4TB storage capacity with all NVMe based SSDs (With Random Read of 636500 IOPS and Random Write of 223300 IOPS)
- 1.5TB of NVMe Cache

For pricing and region availability, see the [Azure VMware Solution pricing page](#) and see the [Products available by region page](#).

Hosts, clusters, and private clouds

Azure VMware Solution clusters are based upon hyper-converged infrastructure. The following table shows the CPU, memory, disk and network specifications of the host.

HOST TYPE	CPU (GHZ)	RAM (GB)	VSAN CACHE TIER (TB, RAW)	VSAN CAPACITY TIER (TB, RAW)	NETWORK INTERFACE CARDS	REGIONAL AVAILABILITY
AV36	Dual Intel Xeon Gold 6140 CPUs with 18 cores/CPU @ 2.3 GHz, Total 36 physical cores (72 logical cores with hyperthreading)	576	3.2 (NVMe)	15.20 (SSD)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	All product regions
AV36P	Dual Intel Xeon Gold 6240 CPUs with 18 cores/CPU @ 2.6 GHz / 3.9 GHz Turbo, Total 36 physical cores (72 logical cores with hyperthreading)	768	1.5 (Intel Cache)	19.20 (NVMe)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	Selected regions (*)

HOST TYPE	CPU (GHZ)	RAM (GB)	VSAN CACHE TIER (TB, RAW)	VSAN CAPACITY TIER (TB, RAW)	NETWORK INTERFACE CARDS	REGIONAL AVAILABILITY
AV52	Dual Intel Xeon Platinum 8270 CPUs with 26 cores/CPU @ 2.7 GHz / 4.0 GHz Turbo, Total 52 physical cores (104 logical cores with hyperthreading)	1,536	1.5 (Intel Cache)	38.40 (NVMe)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	Selected regions (*)

An Azure VMware Solution cluster requires a minimum number of three hosts. You can only use hosts of the same type in a single Azure VMware Solution private cloud. Hosts used to build or scale clusters come from an isolated pool of hosts. Those hosts have passed hardware tests and have had all data securely deleted before being added to a cluster.

(*) details available via the Azure pricing calculator.

You can deploy new or scale existing private clouds through the Azure portal or Azure CLI.

Networking

Azure VMware Solution offers a private cloud environment accessible from on-premises sites and Azure-based resources. Services such as Azure ExpressRoute, VPN connections, or Azure Virtual WAN deliver the connectivity. However, these services require specific network address ranges and firewall ports for enabling the services.

When you deploy a private cloud; private networks for management, provisioning, and vMotion get created. You'll use these private networks to access VMware vCenter Server and VMware NSX-T Data Center NSX-T Manager and virtual machine vMotion or deployment.

[ExpressRoute Global Reach](#) is used to connect private clouds to on-premises environments. It connects circuits directly at the Microsoft Enterprise Edge (MSEE) level. The connection requires a virtual network (vNet) with an ExpressRoute circuit to on-premises in your subscription. The reason is that vNet gateways (ExpressRoute Gateways) can't transit traffic, which means you can attach two circuits to the same gateway, but it won't send the traffic from one circuit to the other.

Each Azure VMware Solution environment is its own ExpressRoute region (its own virtual MSEE device), which lets you connect Global Reach to the 'local' peering location. It allows you to connect multiple Azure VMware Solution instances in one region to the same peering location.

NOTE

For locations where ExpressRoute Global Reach isn't enabled, for example, because of local regulations, you have to build a routing solution using Azure IaaS VMs. For some examples, see [Azure Cloud Adoption Framework - Network topology and connectivity for Azure VMware Solution](#).

Virtual machines deployed on the private cloud are accessible to the internet through the [Azure Virtual WAN public IP](#) functionality. For new private clouds, internet access is disabled by default.

For more information, see [Networking concepts](#).

Access and security

Azure VMware Solution private clouds use vSphere role-based access control for enhanced security. You can integrate vSphere SSO LDAP capabilities with Azure Active Directory. For more information, see the [Access and Identity concepts](#) page.

vSAN data-at-rest encryption, by default, is enabled and is used to provide vSAN datastore security. For more information, see [Storage concepts](#).

Data Residency and Customer Data

Azure VMware Solution doesn't store customer data.

VMware software versions

The VMware solution software versions used in new deployments of Azure VMware Solution private cloud clusters are:

SOFTWARE	VERSION
VMware vCenter Server	7.0 U3c
ESXi	7.0 U3c
vSAN	7.0 U3c
vSAN on-disk format	10
HCX	4.4.2
VMware NSX-T Data Center NOTE: VMware NSX-T Data Center is the only supported version of NSX Data Center.	3.1.2

The current running software version is applied to new clusters added to an existing private cloud. For more information, see the [VMware software version requirements for HCX](#) and [Understanding vSAN on-disk format versions and compatibility](#).

Host and software lifecycle maintenance

Regular upgrades of the Azure VMware Solution private cloud and VMware software ensure the latest security, stability, and feature sets are running in your private clouds. For more information, see [Host maintenance and lifecycle management](#).

Monitoring your private cloud

Once you've deployed Azure VMware Solution into your subscription, [Azure Monitor logs](#) are generated automatically.

In your private cloud, you can:

- Collect logs on each of your VMs.
- [Download and install the MMA agent](#) on Linux and Windows VMs.

- Enable the [Azure diagnostics extension](#).
- [Create and run new queries](#).
- Run the same queries you usually run on your VMs.

Monitoring patterns inside the Azure VMware Solution are similar to Azure VMs within the IaaS platform. For more information and how-tos, see [Monitoring Azure VMs with Azure Monitor](#).

Customer communication

You can find service issues, planned maintenance, health advisories, and security advisories notifications published through **Service Health** in the Azure portal. You can take timely actions when you set up activity log alerts for these notifications. For more information, see [Create Service Health alerts using the Azure portal](#).

The screenshot shows the Azure Service Health interface. At the top, there's a search bar and filters for Subscription (4 selected), Region (3 selected), and Service (3 selected). The main content area is titled 'Service Health | Health advisories (2)'. Under 'ACTIVE EVENTS', 'Health advisories (2)' is selected, showing a table with two entries:

Issue Name	Tracking ID	Service(s)	Region(s)	Start Time
Azure Disaster Recovery Drill for East ...	9SX9-LT8	Network Infrastructure	East US 2 EUAP	2020-10-21T23:50:16Z
Action Required: Review your comput...	DT38-7C0	Azure VMware Solution	West US	2020-11-06T00:00:00Z

Below the table, it says 'See 1 advisory incident(s) outside of your filter.' On the left sidebar, there are sections for 'ACTIVE EVENTS', 'HISTORY', 'RESOURCE HEALTH', and 'ALERTS'. At the bottom, there's an alert message: 'We have important information for your Azure VMware Solution service in the West US region. A private cloud instance is running low on compute and storage capacity, which can affect its performance or your Service Level Agreements. To avoid disruption to any' with a 'Create a support request' link.

Azure VMware Solution Responsibility Matrix - Microsoft vs Customer

Azure VMware Solution implements a shared responsibility model that defines distinct roles and responsibilities of the two parties involved in the offering: Customer and Microsoft. The shared role responsibilities are illustrated in more detail in following two tables.

The shared responsibility matrix table shows the high-level responsibilities between a customer and Microsoft for different aspects of the deployment/management of the private cloud and the customer application workloads.

Azure VMware Solution – Shared responsibility Matrix

Control boundaries

	Physical Infrastructure	Physical Security	Azure VMware Solution Portal	Hardware Failures	ESXi Host\Patching	VMware NSX-T Data Center	VMware vCenter Server	VSAN	HCK/SRM	Portal/Platform Identity Management	Connectivity to VNET/Internet	Virtual Machines	Guest OS	Applications	3 rd party Solution
Deployment/Lifecycle	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Grey	Grey	Grey	Grey
Provider Configuration	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Black	Black	Black	Black
Tenant Configuration	Black	Black	Black	Black	Black	Black	Black	Black	Black	Black	Black	Black	Black	Black	Black
Support	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Blue	Grey	Grey	Grey	Cyan

■ - Microsoft Responsibility
 ■ - Customer/Tenant Responsibility
 ■ - Not Applicable
 ■ - 3rd Party Vendor

The following table provides a detailed list of roles and responsibilities between the customer and Microsoft, which encompasses the most frequent tasks and definitions. For further questions, contact Microsoft.

ROLE	TASK/DETAILS
------	--------------

ROLE	TASK/DETAILS
Microsoft - Azure VMware Solution	<p>Physical infrastructure</p> <ul style="list-style-type: none"> • Azure regions • Azure availability zones • Express Route/Global reach <p>Compute/Network/Storage</p> <ul style="list-style-type: none"> • Rack and power Bare Metal hosts • Rack and power network equipment <p>Software defined Data Center (SDDC) deploy/lifecycle</p> <ul style="list-style-type: none"> • VMware ESXi deploy, patch, and upgrade • VMware vCenter Servers deploy, patch, and upgrade • VMware NSX-T Data Centers deploy, patch, and upgrade • vSAN deploy, patch, and upgrade <p>SDDC Networking - VMware NSX-T Data Center provider config</p> <ul style="list-style-type: none"> • Microsoft Edge node/cluster, VMware NSX-T Data Center host preparation • Provider Tier-0 and Tenant Tier-1 Gateway • Connectivity from Tier-0 (using BGP) to Azure Network via Express Route <p>SDDC Compute - VMware vCenter Server provider config</p> <ul style="list-style-type: none"> • Create default cluster • Configure virtual networking for vMotion, Management, vSAN, and others <p>SDDC backup/restore</p> <ul style="list-style-type: none"> • Backup and restore VMware vCenter Server • Backup and restore VMware NSX-T Data Center NSX-T Manager <p>SDDC health monitoring and corrective actions, for example: replace failed hosts</p> <p>(optional) HCX deploys with fully configured compute profile on cloud side as add-on</p> <p>(optional) SRM deploys, upgrade, and scale up/down</p> <p>Support - SDDC platforms and HCX</p>

ROLE	TASK/DETAILS
Customer	<p>Request Azure VMware Solution host quote with Microsoft Plan and create a request for SDDCs on Azure portal with:</p> <ul style="list-style-type: none"> • Host count • Management network range • Other information <p>Configure SDDC network and security (VMware NSX-T Data Center)</p> <ul style="list-style-type: none"> • Network segments to host applications • Additional Tier -1 routers • Firewall • VMware NSX-T Data Center LB • IPsec VPN • NAT • Public IP addresses • Distributed firewall/gateway firewall • Network extension using HCX or VMware NSX-T Data Center • AD/LDAP config for RBAC <p>Configure SDDC - VMware vCenter Server</p> <ul style="list-style-type: none"> • AD/LDAP config for RBAC • Deploy and lifecycle management of Virtual Machines (VMs) and application <ul style="list-style-type: none"> ◦ Install operating systems ◦ Patch operating systems ◦ Install antivirus software ◦ Install backup software ◦ Install configuration management software ◦ Install application components ◦ VM networking using VMware NSX-T Data Center segments • Migrate Virtual Machines (VMs) <ul style="list-style-type: none"> ◦ HCX configuration ◦ Live vMotion ◦ Cold migration ◦ Content library sync <p>Configure SDDC - vSAN</p> <ul style="list-style-type: none"> • Define and maintain vSAN VM policies • Add hosts to maintain adequate 'slack space' <p>Configure HCX</p> <ul style="list-style-type: none"> • Download and deploy HCA connector OVA in on-premises • Pairing on-premises HCX connector • Configure the network profile, compute profile, and service mesh • Configure HCX network extension/MON • Upgrade/updates <p>Network configuration to connect to on-premises, VNET, or internet</p> <p>Add or delete hosts requests to cluster from Portal</p> <p>Deploy/lifecycle management of partner (third party) solutions</p>

ROLE	TASK/DETAILS
Partner ecosystem	<p>Support for their product/solution. For reference, the following are some of the supported Azure VMware Solution partner solution/product:</p> <ul style="list-style-type: none">• BCDR - SRM, JetStream, RiverMeadow, and others• Backup - Veeam, Commvault, Rubrik, and others• VDI - Horizon/Citrix• Security solutions - BitDefender, TrendMicro, Checkpoint• Other VMware products - vRA, vROps, AVI

Next steps

The next step is to learn key [private cloud and cluster concepts](#).

What's new in Azure VMware Solution

12/16/2022 • 8 minutes to read • [Edit Online](#)

Microsoft will regularly apply important updates to the Azure VMware Solution for new features and software lifecycle management. You'll receive a notification through Azure Service Health that includes the timeline of the maintenance. For more information, see [Host maintenance and lifecycle management](#).

November 2022

AV36P and AV52 node sizes available in Azure VMware Solution. The new node sizes increase memory and storage options to optimize your workloads. The gains in performance enable you to do more per server, break storage bottlenecks, and lower transaction costs of latency-sensitive workloads. The availability of the new nodes allows for large latency-sensitive services to be hosted efficiently on the Azure VMware Solution infrastructure.

For pricing and region availability, see the [Azure VMware Solution pricing page](#) and see the [Products available by region page](#).

July 2022

- HCX cloud manager in Azure VMware Solution can now be accessible over a public IP address. You can pair HCX sites and create a service mesh from on-premises to Azure VMware Solution private cloud using Public IP.

HCX with public IP is especially useful in cases where On-premises sites are not connected to Azure via Express Route or VPN. HCX service mesh appliances can be configured with public IPs to avoid lower tunnel MTUs due to double encapsulation if a VPN is used for on-premises to cloud connections. For more information, please see [Enable HCX over the internet](#)

- All new Azure VMware Solution private clouds are now deployed with VMware vCenter Server version 7.0 Update 3c and ESXi version 7.0 Update 3c. Any existing private clouds will be upgraded to those versions. For more information, please see [VMware ESXi 7.0 Update 3c Release Notes](#) and [VMware vCenter Server 7.0 Update 3c Release Notes](#). You'll receive a notification through Azure Service Health that includes the timeline of the upgrade. You can reschedule an upgrade as needed. This notification also provides details on the upgraded component, its effect on workloads, private cloud access, and other Azure services.

June 2022

All new Azure VMware Solution private clouds in regions (East US2, Canada Central, North Europe, and Japan East), are now deployed in with VMware vCenter Server version 7.0 Update 3c and ESXi version 7.0 Update 3c.

Any existing private clouds in the above mentioned regions will also be upgraded to these versions. For more information, please see [VMware ESXi 7.0 Update 3c Release Notes](#) and [VMware vCenter Server 7.0 Update 3c Release Notes](#).

May 2022

- All new Azure VMware Solution private clouds in regions (Germany West Central, Australia East, Central US and UK West), are now deployed with VMware vCenter Server version 7.0 Update 3c and ESXi version 7.0 Update 3c. Any existing private clouds in the previously mentioned regions will be upgraded to those versions. For

more information, please see [VMware ESXi 7.0 Update 3c Release Notes](#) and [VMware vCenter Server 7.0 Update 3c Release Notes](#). You'll receive a notification through Azure Service Health that includes the timeline of the upgrade. You can reschedule an upgrade as needed. This notification also provides details on the upgraded component, its effect on workloads, private cloud access, and other Azure services.

- All new Azure VMware Solution private clouds in regions (France Central, Brazil South, Japan West, Australia Southeast, Canada East, East Asia, and Southeast Asia), are now deployed with VMware vCenter Server version 7.0 Update 3c and ESXi version 7.0 Update 3c. Any existing private clouds in the previously mentioned regions will be upgraded to those versions. For more information, please see [VMware ESXi 7.0 Update 3c Release Notes](#) and [VMware vCenter Server 7.0 Update 3c Release Notes](#). You'll receive a notification through Azure Service Health that includes the timeline of the upgrade. You can reschedule an upgrade as needed. This notification also provides details on the upgraded component, its effect on workloads, private cloud access, and other Azure services.

February 2022

Per VMware security advisory [VMSA-2022-0004](#), multiple vulnerabilities in VMware ESXi have been reported to VMware.

To address the vulnerabilities (CVE-2021-22040 and CVE-2021-22041) reported in this VMware security advisory, ESXi hosts have been patched in all Azure VMware Solution private clouds to ESXi 6.7, Patch Release ESXi670-202111001. All new Azure VMware Solution private clouds are deployed with the same version.

For more information on this ESXi version, see [VMware ESXi 6.7, Patch Release ESXi670-202111001](#).

No further action is required.

December 2021

- Azure VMware Solution (AVS) has completed maintenance activities to address critical vulnerabilities in Apache Log4j. The fixes documented in the VMware security advisory [VMSA-2021-0028.6](#) to address CVE-2021-44228 and CVE-2021-45046 have been applied to these AVS managed VMware products: vCenter Server, NSX-T Data Center, SRM and HCX. We strongly encourage customers to apply the fixes to on-premises HCX connector appliances. We also recommend customers to review the security advisory and apply the fixes for other affected VMware products or workloads. If you need any assistance or have questions, please [contact us](#).

- VMware has announced a security advisory [VMSA-2021-0028](#), addressing a critical vulnerability in Apache Log4j identified by CVE-2021-44228. Azure VMware Solution is actively monitoring this issue. We are addressing this issue by applying VMware recommended workarounds or patches for AVS managed VMware components as they become available. Please note that you may experience intermittent connectivity to these components when we apply a fix. We strongly recommend that you read the advisory and patch or apply the recommended workarounds for any additional VMware products that you may have deployed in Azure VMware Solution. If you need any assistance or have questions, please [contact us](#).

November 2021

Per VMware security advisory [VMSA-2021-0027](#), multiple vulnerabilities in VMware vCenter Server have been reported to VMware.

To address the vulnerabilities (CVE-2021-21980 and CVE-2021-22049) reported in VMware security advisory, vCenter Server has been updated to 6.7 Update 3p release in all Azure VMware Solution private clouds.

For more information, see [VMware vCenter Server 6.7 Update 3p Release Notes](#)

No further action is required.

September 2021

- Per VMware security advisory [VMSA-2021-0020](#), multiple vulnerabilities in the VMware vCenter Server have been reported to VMware. To address the vulnerabilities (CVE-2021-21991, CVE-2021-21992, CVE-2021-21993, CVE-2021-22005, CVE-2021-22006, CVE-2021-22007, CVE-2021-22008, CVE-2021-22009, CVE-2021-22010, CVE-2021-22011, CVE-2021-22012, CVE-2021-22013, CVE-2021-22014, CVE-2021-22015, CVE-2021-22016, CVE-2021-22017, CVE-2021-22018, CVE-2021-22019, CVE-2021-22020) reported in VMware security advisory [VMSA-2021-0020](#), vCenter Server has been updated to 6.7 Update 3o in all Azure VMware Solution private clouds. All new Azure VMware Solution private clouds are deployed with vCenter Server version 6.7 Update 3o. For more information, see [VMware vCenter Server 6.7 Update 3o Release Notes](#). No further action is required.
- All new Azure VMware Solution private clouds are now deployed with ESXi version ESXi670-202103001 (Build number: 17700523). ESXi hosts in existing private clouds have been patched to this version. For more information on this ESXi version, see [VMware ESXi 6.7, Patch Release ESXi670-202103001](#).

July 2021

All new Azure VMware Solution private clouds are now deployed with NSX-T Data Center version 3.1.2. NSX-T Data Center version in existing private clouds will be upgraded through September, 2021 to NSX-T Data Center 3.1.2 release.

You'll receive an email with the planned maintenance date and time. You can reschedule an upgrade. The email also provides details on the upgraded component, its effect on workloads, private cloud access, and other Azure services.

For more information on this NSX-T Data Center version, see [VMware NSX-T Data Center 3.1.2 Release Notes](#).

May 2021

- Per VMware security advisory [VMSA-2021-0010](#), multiple vulnerabilities in VMware ESXi and vSphere Client (HTML5) have been reported to VMware. To address the vulnerabilities ([CVE-2021-21985](#) and [CVE-2021-21986](#)) reported in VMware security advisory [VMSA-2021-0010](#), vCenter Server has been updated in all Azure VMware Solution private clouds. No further action is required.
- Azure VMware Solution service will do maintenance work through May 23, 2021, to apply important updates to the vCenter Server in your private cloud. You'll receive a notification through Azure Service Health that includes the timeline of the maintenance for your private cloud. During this time, VMware vCenter Server will be unavailable and you won't be able to manage VMs (stop, start, create, or delete). It's recommended that, during this time, you don't plan any other activities like scaling up private cloud, creating new networks, and so on, in your private cloud. There is no impact to workloads running in your private cloud.

April 2021

All new Azure VMware Solution private clouds are now deployed with VMware vCenter Server version 6.7U31 and NSX-T Data Center version 2.5.2. We're not using NSX-T Data Center 3.1.1 for new private clouds because of an identified issue in NSX-T Data Center 3.1.1 that impacts customer VM connectivity.

The VMware recommended mitigation was applied to all existing private clouds currently running NSX-T Data Center 3.1.1 on Azure VMware Solution. The workaround has been confirmed that there's no impact to customer VM connectivity.

March 2021

- All new Azure VMware Solution private clouds are deployed with VMware vCenter Server version 6.7U3I and NSX-T Data Center version 3.1.1. Any existing private clouds will be updated and upgraded **through June 2021** to the releases mentioned above. You'll receive an email with the planned maintenance date and time. You can reschedule an upgrade. The email also provides details on the upgraded component, its effect on workloads, private cloud access, and other Azure services. An hour before the upgrade, you'll receive a notification and then again when it finishes.
- Azure VMware Solution service will do maintenance work **through March 19, 2021**, to update the vCenter Server in your private cloud to vCenter Server 6.7 Update 3I version. VMware vCenter Server will be unavailable during this time, so you can't manage your VMs (stop, start, create, delete) or private cloud scaling (adding/removing servers and clusters). However, VMware High Availability (HA) will continue to operate to protect existing VMs.
For more information on this vCenter version, see [VMware vCenter Server 6.7 Update 3I Release Notes](#).
 - Azure VMware Solution will apply the [VMware ESXi 6.7, Patch Release ESXi670-202011002](#) to existing privates **through March 15, 2021**.
 - Documented workarounds for the vSphere stack, as per [VMSA-2021-0002](#), will also be applied **through March 15, 2021**.

NOTE

This is non-disruptive and should not impact Azure VMware Services or workloads. During maintenance, various VMware alerts, such as *Lost network connectivity on DVPorts* and *Lost uplink redundancy on DVPorts*, appear in vCenter Server and clear automatically as the maintenance progresses.

Post update

Once complete, newer versions of VMware solution components will appear. If you notice any issues or have any questions, contact our support team by opening a support ticket.

Plan the Azure VMware Solution deployment

12/16/2022 • 11 minutes to read • [Edit Online](#)

Planning your Azure VMware Solution deployment is critical for a successful production-ready environment for creating virtual machines (VMs) and migration. During the planning process, you'll identify and gather what's needed for your deployment. As you plan, make sure to document the information you gather for easy reference during the deployment. A successful deployment results in a production-ready environment for creating virtual machines (VMs) and migration.

In this how-to article, you'll do the following tasks:

- Identify the Azure subscription, resource group, region, and resource name
- Identify the size hosts and determine the number of clusters and hosts
- Request a host quota for eligible Azure plan
- Identify the /22 CIDR IP segment for private cloud management
- Identify a single network segment
- Define the virtual network gateway
- Define VMware HCX network segments

After you're finished, follow the recommended [Next steps](#) at the end of this article to continue with this getting started guide.

Identify the subscription

Identify the subscription you plan to use to deploy Azure VMware Solution. You can create a new subscription or use an existing one.

NOTE

The subscription must be associated with a Microsoft Enterprise Agreement (EA), a Cloud Solution Provider (CSP) Azure plan or an Microsoft Customer Agreement (MCA). For more information, see [Eligibility criteria](#).

Identify the resource group

Identify the resource group you want to use for your Azure VMware Solution. Generally, a resource group is created specifically for Azure VMware Solution, but you can use an existing resource group.

Identify the region or location

Identify the [region](#) you want Azure VMware Solution deployed.

Define the resource name

The resource name is a friendly and descriptive name in which you title your Azure VMware Solution private cloud, for example, **MyPrivateCloud**.

IMPORTANT

The name must not exceed 40 characters. If the name exceeds this limit, you won't be able to create public IP addresses for use with the private cloud.

Identify the size hosts

Identify the size hosts that you want to use when deploying Azure VMware Solution.

Azure VMware Solution clusters are based upon hyper-converged infrastructure. The following table shows the CPU, memory, disk and network specifications of the host.

HOST TYPE	CPU (GHZ)	RAM (GB)	VSAN CACHE TIER (TB, RAW)	VSAN CAPACITY TIER (TB, RAW)	NETWORK INTERFACE CARDS	REGIONAL AVAILABILITY
AV36	Dual Intel Xeon Gold 6140 CPUs with 18 cores/CPU @ 2.3 GHz, Total 36 physical cores (72 logical cores with hyperthreading)	576	3.2 (NVMe)	15.20 (SSD)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	All product regions
AV36P	Dual Intel Xeon Gold 6240 CPUs with 18 cores/CPU @ 2.6 GHz / 3.9 GHz Turbo, Total 36 physical cores (72 logical cores with hyperthreading)	768	1.5 (Intel Cache)	19.20 (NVMe)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	Selected regions (*)
AV52	Dual Intel Xeon Platinum 8270 CPUs with 26 cores/CPU @ 2.7 GHz / 4.0 GHz Turbo, Total 52 physical cores (104 logical cores with hyperthreading)	1,536	1.5 (Intel Cache)	38.40 (NVMe)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	Selected regions (*)

An Azure VMware Solution cluster requires a minimum number of three hosts. You can only use hosts of the

AREA	DESCRIPTION	PROVISIONED VCPUS	PROVISIONED VRAM (GB)	PROVISIONED VDISK (GB)	TYPICAL CPU USAGE (GHZ)	TYPICAL VRAM USAGE (GB)	TYPICAL RAW VSAN DATASTORE USAGE (GB)
VMware NSX-T Data Center	NSX-T Unified Appliance Node 1	6	24	300	5.5	8.5	613
VMware NSX-T Data Center	NSX-T Unified Appliance Node 2	6	24	300	5.5	8.5	613
VMware NSX-T Data Center	NSX-T Unified Appliance Node 3	6	24	300	5.5	8.5	613
VMware NSX-T Data Center	NSX-T Edge VM 1	8	32	200	1.3	0.6	409
VMware NSX-T Data Center	NSX-T Edge VM 2	8	32	200	1.3	0.6	409
VMware HCX (Optional Add-On)	HCX Manager	4	12	65	1	3.2	152
VMware Site Recovery Manager (Optional Add-On)	SRM Appliance	4	12	33	1	1	93
VMware vSphere (Optional Add-On)	vSphere Replication Manager Appliance	4	8	33	4.3	2.2	84
VMware vSphere (Optional Add-On)	vSphere Replication Server Appliance	2	1	33	1	0.1	84
	Total	59 vCPUs	197.3 GB	2,394 GB	56 GHz	38.3 GB	11,575 GB (9,646 GB with expected 1.2x Data Reduction ratio)

These resource requirements only apply to the first cluster deployed in an Azure VMware Solution private cloud. Subsequent clusters only need to account for the vSphere Cluster Service, ESXi resource requirements and

vSAN System Usage in solution sizing.

The virtual appliance **Typical Raw vSAN Datastore Usage** values account for the space occupied by virtual machine files, including configuration and log files, snapshots, virtual disks and swap files.

The VMware ESXi nodes have compute usage values that account for the vSphere VMkernel hypervisor overhead, vSAN overhead and NSX-T distributed router, firewall and bridging overhead. These are estimates for a standard three cluster configuration. The storage requirements are listed as not applicable (N/A) since a boot volume separate from the vSAN Datastore is used.

The VMware vSAN System Usage storage overhead accounts for vSAN performance management objects, vSAN file system overhead, vSAN checksum overhead and vSAN deduplication and compression overhead. To view this consumption, select the Monitor, vSAN Capacity object for the vSphere Cluster in the vSphere Client.

The VMware HCX and VMware Site Recovery Manager resource requirements are optional Add-Ons to the Azure VMware Solution service. Discount these requirements in the solution sizing if they are not being used.

The VMware Site Recovery Manager Add-On has the option of configuring multiple VMware vSphere Replication Server Appliances. The table above assumes one vSphere Replication Server appliance is used.

Sizing an Azure VMware Solution is an estimate; the sizing calculations from the design phase should be validated during the testing phase of a project to ensure the Azure VMware Solution has been sized correctly for the application workload.

TIP

You can always extend the cluster and add additional clusters later if you need to go beyond the initial deployment number.

NOTE

To learn about the limits for the number of hosts per cluster, the number of clusters per private cloud, and the number of hosts per private cloud, check [Azure subscription and service limits, quotas, and constraints](#).

Request a host quota

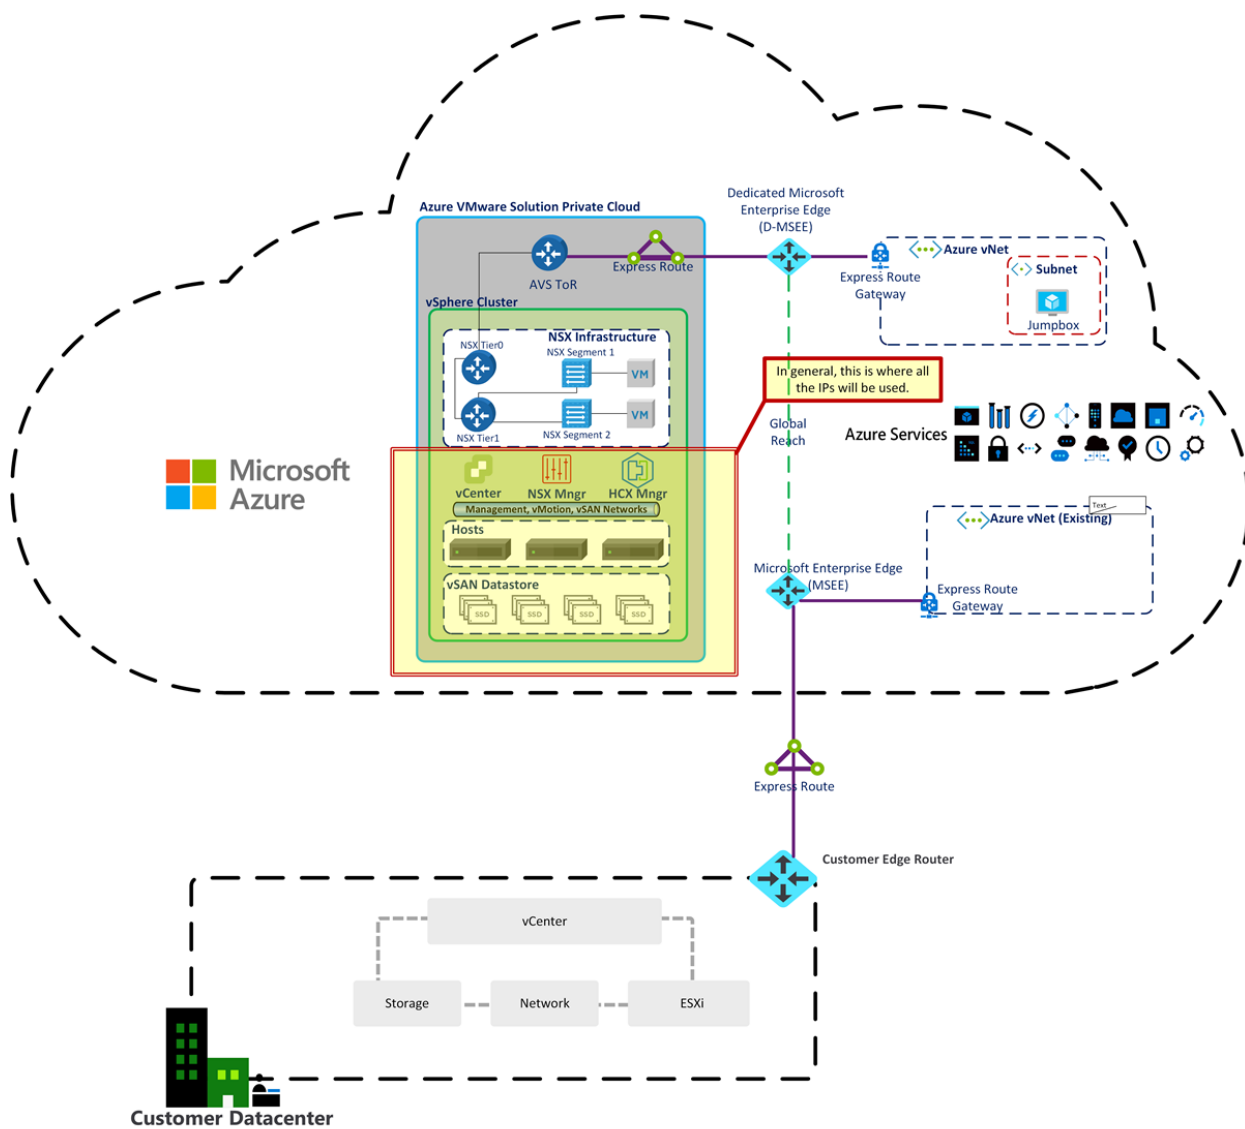
It's crucial to request a host quota early, so after you've finished the planning process, you're ready to deploy your Azure VMware Solution private cloud. Before requesting a host quota, make sure you've identified the Azure subscription, resource group, and region. Also, make sure you've identified the size hosts and determine the number of clusters and hosts you'll need.

After the support team receives your request for a host quota, it takes up to five business days to confirm your request and allocate your hosts.

- [EA customers](#)
- [CSP customers](#)

Define the IP address segment for private cloud management

Azure VMware Solution requires a /22 CIDR network, for example, `10.0.0.0/22`. This address space is carved into smaller network segments (subnets) and used for Azure VMware Solution management segments, including: vCenter Server, VMware HCX, NSX-T Data Center, and vMotion functionality. The diagram highlights Azure VMware Solution management IP address segments.



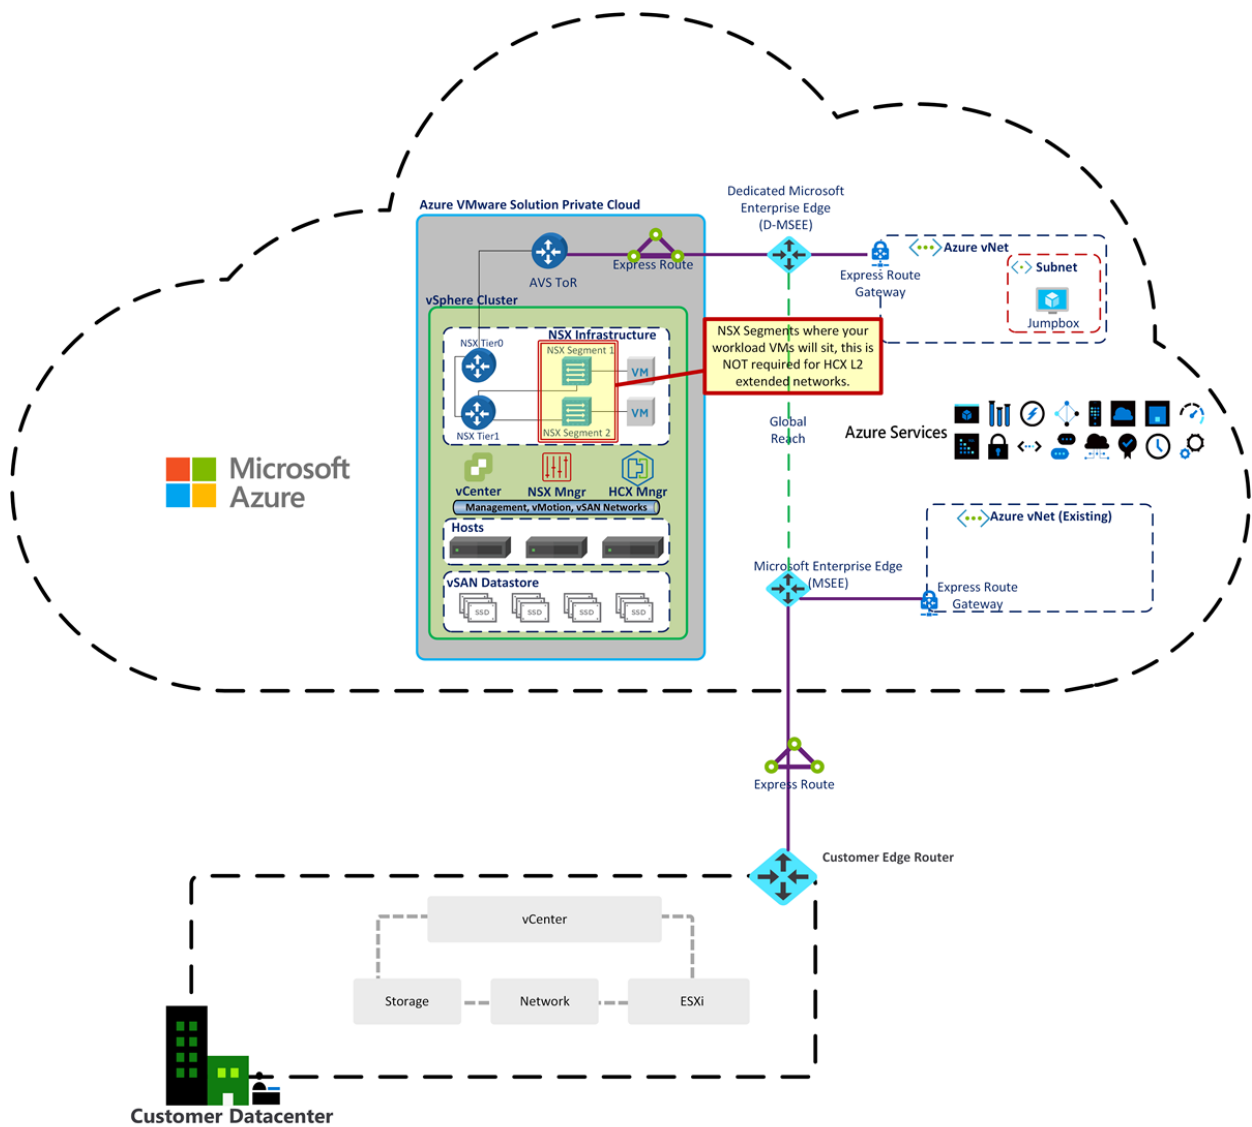
IMPORTANT

The /22 CIDR network address block shouldn't overlap with any existing network segment you already have on-premises or in Azure. For details of how the /22 CIDR network is broken down per private cloud, see [Routing and subnet considerations](#).

Define the IP address segment for VM workloads

Like with any VMware vSphere environment, the VMs must connect to a network segment. As the production deployment of Azure VMware Solution expands, there's often a combination of L2 extended segments from on-premises and local NSX-T network segments.

For the initial deployment, identify a single network segment (IP network), for example, `10.0.4.0/24`. This network segment is used primarily for testing purposes during the initial deployment. The address block shouldn't overlap with any network segments on-premises or within Azure and shouldn't be within the /22 network segment already defined.



Define the virtual network gateway

Azure VMware Solution requires an Azure Virtual Network and an ExpressRoute circuit. Define whether you want to use an *existing* OR *new* ExpressRoute virtual network gateway. If you decide to use a *new* virtual network gateway, you'll create it after creating your private cloud. It's acceptable to use an existing ExpressRoute virtual network gateway. For planning purposes, make a note of which ExpressRoute virtual network gateway you'll use.

NOTE

Preparing for large environments, instead of using the management network used for the on-premises VMware vSphere cluster, create a new /26 network and present that network as a port group to your on-premises VMware vSphere cluster. You can then create up to 10 service meshes and 60 network extenders (-1 per service mesh). You can stretch **eight** networks per network extender by using Azure VMware Solution private clouds.

- **Uplink network:** When deploying VMware HCX on-premises, you'll need to identify an Uplink network for VMware HCX. Use the same network you plan to use for the Management network.
- **vMotion network:** When deploying VMware HCX on-premises, you'll need to identify a vMotion network for VMware HCX. Typically, it's the same network used for vMotion by your on-premises VMware vSphere cluster. At a minimum, identify **two** IPs on this network segment for VMware HCX. You might need larger numbers, depending on the scale of your deployment beyond the pilot or small use case.

You must expose the vMotion network on a distributed virtual switch or vSwitch0. If it's not, modify the environment to accommodate.

NOTE

Many VMware vSphere environments use non-routed network segments for vMotion, which poses no problems.

- **Replication network:** When deploying VMware HCX on-premises, you'll need to define a replication network. Use the same network you're using for your Management and Uplink networks. If the on-premises cluster hosts use a dedicated Replication VMkernel network, reserve **two** IP addresses in this network segment and use the Replication VMkernel network for the replication network.

Determine whether to extend your networks

Optionally, you can extend network segments from on-premises to Azure VMware Solution. If you do extend network segments, identify those networks now following these guidelines:

- Networks must connect to a [vSphere Distributed Switch \(vDS\)](#) in your on-premises VMware environment.
- Networks that are on a [vSphere Standard Switch](#) can't be extended.

IMPORTANT

These networks are extended as a final step of the configuration, not during deployment.

Next steps

Now that you've gathered and documented the information needed, continue to the next tutorial to create your Azure VMware Solution private cloud.

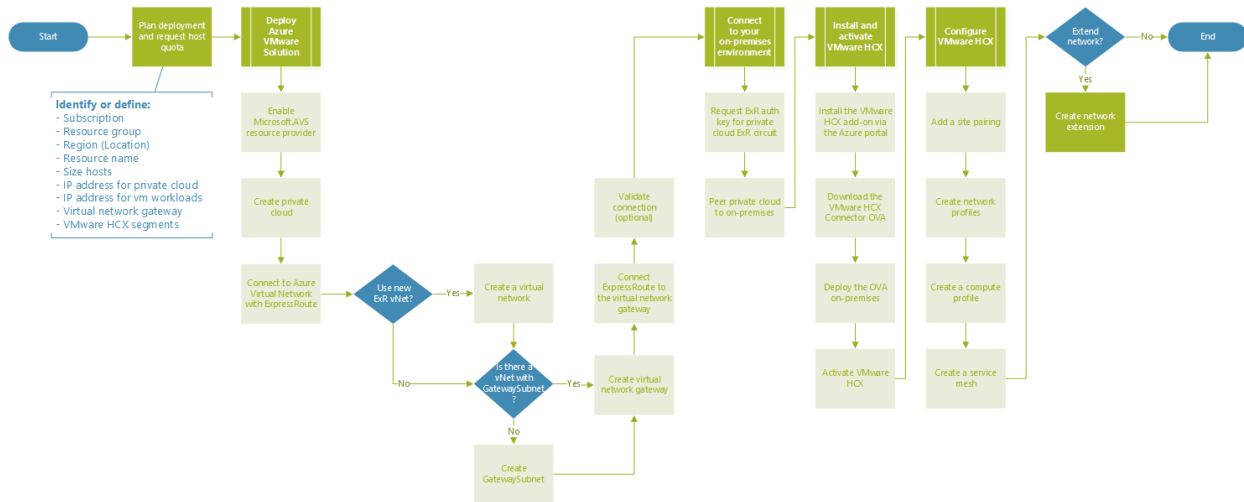
[Deploy Azure VMware Solution](#)

Deploy and configure Azure VMware Solution

12/16/2022 • 8 minutes to read • [Edit Online](#)

Once you've [planned your deployment](#), you'll deploy and configure your Azure VMware Solution private cloud.

The diagram shows the deployment workflow of Azure VMware Solution.



In this how-to, you'll:

- Register the resource provider and create a private cloud
- Connect to a new or existing ExpressRoute virtual network gateway
- Validate the network connect

After you're finished, follow the recommended next steps at the end to continue with the steps of this getting started guide.

Register the Microsoft.AVS resource provider

To use Azure VMware Solution, you must first register the resource provider with your subscription. For more information about resource providers, see [Azure resource providers and types](#).

- [Portal](#)
- [Azure CLI](#)

1. Sign in to the [Azure portal](#).
2. On the Azure portal menu, select **All services**.
3. In the **All services** box, enter **subscription**, and then select **Subscriptions**.
4. Select the subscription from the subscription list to view.
5. Select **Resource providers** and enter **Microsoft.AVS** into the search.
6. If the resource provider is not registered, select **Register**.

Create an Azure VMware Solution private cloud

You can create an Azure VMware Solution private cloud using the Azure portal or the Azure CLI.

- [Portal](#)
- [Azure CLI](#)

1. Sign in to the [Azure portal](#).
2. Select **Create a resource**.
3. In the **Search services and marketplace** text box, type `Azure VMware Solution` and select it from the search results.
4. On the **Azure VMware Solution** window, select **Create**.
5. If you need more hosts, [request a host quota increase](#).
6. On the **Basics** tab, enter values for the fields and then select **Review + Create**.

TIP

You gathered this information during the [planning phase](#) of this quick start.

FIELD	VALUE
Subscription	Select the subscription you plan to use for the deployment. All resources in an Azure subscription are billed together.
Resource group	Select the resource group for your private cloud. An Azure resource group is a logical container into which Azure resources are deployed and managed. Alternatively, you can create a new resource group for your private cloud.
Resource name	Provide the name of your Azure VMware Solution private cloud.
Location	Select a location, such as east us . It's the <i>region</i> you defined during the planning phase.
Size of host	Select the AV36 , AV36P or AV52 SKU.
Number of hosts	Number of hosts allocated for the private cloud cluster. The default value is 3, which you can increase or decrease after deployment. If these nodes are not listed as available, please contact support to request a quota increase . You can also click the link labeled If you need more hosts, request a quota increase in the Azure portal.
Address block for private cloud	Provide an IP address block for the private cloud. The CIDR represents the private cloud management network and is used for the cluster management services, such as vCenter Server and NSX-T Manager. Use /22 address space, for example, 10.175.0.0/22. The address should be unique and not overlap with other Azure Virtual Networks and with on-premises networks.

Microsoft Azure Search resources, services, and docs (G+)

Home > Create a resource > Marketplace > Azure VMware Solution >

Create a private cloud

Prerequisites ***Basics** Tags Review and Create

Project details

Subscription * ⓘ AnyBuild-InternalProdClusters

Resource group * ⓘ [Create new](#)

Private cloud details

Resource name * ⓘ

Location * ⓘ East US

Size of host * ⓘ

Number of hosts ⓘ 3

[Find out how many hosts you need](#)
If you need more hosts, request a quota increase

[Review and Create](#) [Previous](#) [Next : Tags >](#)

7. Verify the information entered, and if correct, select **Create**.

NOTE

This step takes roughly 3-4 hours. Adding a single host in an existing or the same cluster takes between 30 - 45 minutes.

8. Verify that the deployment was successful. Navigate to the resource group you created and select your private cloud. You'll see the status of **Succeeded** when the deployment has finished.

Microsoft Azure Search resources, services, and docs (G+)

Home >

Contoso-westus-sddc AVS Private cloud

Search (Ctrl+/) Delete

Overview

- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems

Settings

- Locks

Essentials

Resource group (change) : Contoso-westus-rg

Status : Succeeded

Location : West US

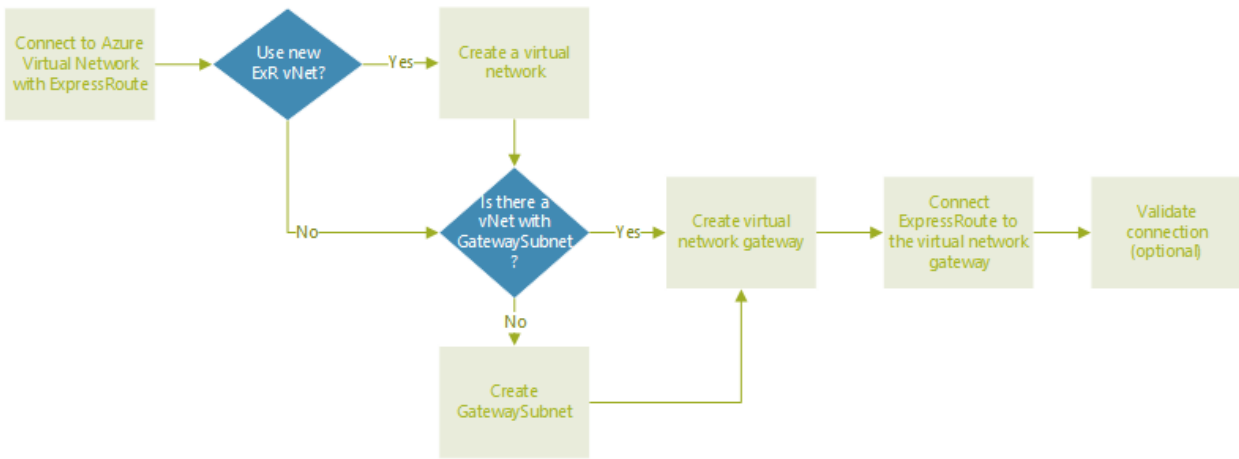
Subscription (change) : Contoso

Subscription ID : 1234abc-d567-8910-abdc-2e2bb12345e6

Tags (change) : [Click here to add tags](#)

Connect to Azure Virtual Network with ExpressRoute

In the planning phase, you defined whether to use an *existing* or *new* ExpressRoute virtual network gateway.



IMPORTANT

If you plan to scale your Azure VMware Solution hosts using [Azure NetApp Files datastores](#), deploying the vNet close to your hosts with an ExpressRoute virtual network gateway is crucial. The closer the storage is to your hosts, the better the performance.

Use a new ExpressRoute virtual network gateway

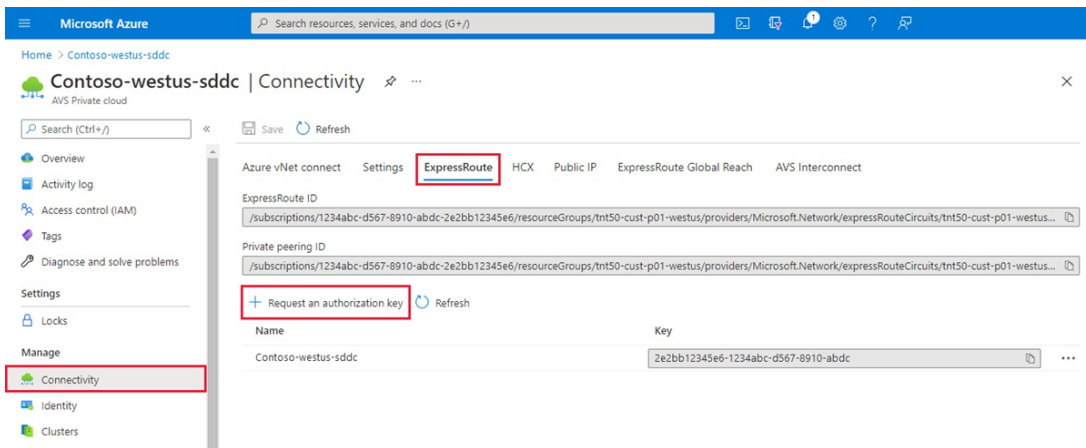
IMPORTANT

You must have a virtual network with a GatewaySubnet that **does not** already have a virtual network gateway.

IF	THEN
You don't already have a virtual network...	Create the following: 1. Virtual network 2. GatewaySubnet 3. Virtual network gateway 4. Connect ExpressRoute to the gateway
You already have a virtual network without a GatewaySubnet...	Create the following: 1. GatewaySubnet 2. Virtual network gateway 3. Connect ExpressRoute to the gateway
You already have a virtual network with a GatewaySubnet...	Create the following: 1. Virtual network gateway 2. Connect ExpressRoute to the gateway

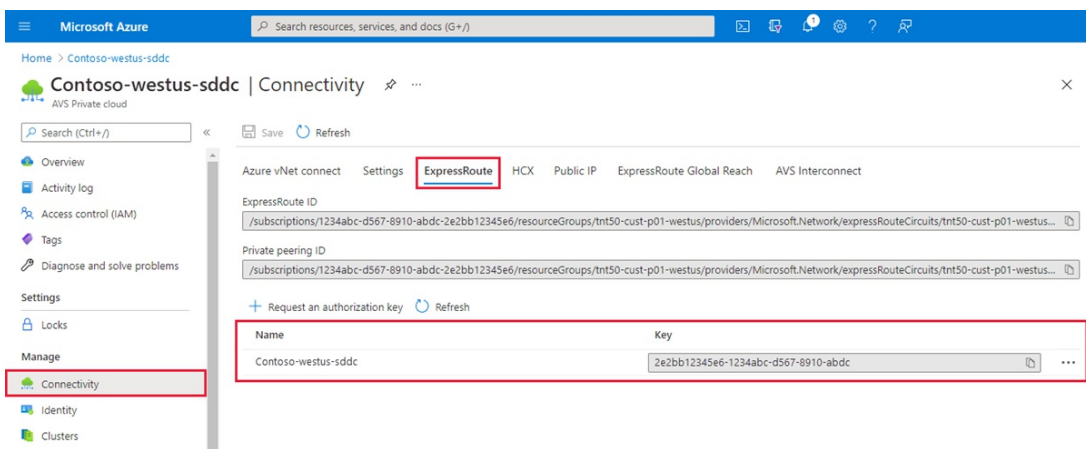
Use an existing virtual network gateway

1. Request an ExpressRoute authorization key:
 - a. In the Azure portal, navigate to the Azure VMware Solution private cloud. Select **Manage** > **Connectivity** > **ExpressRoute** and then select + **Request an authorization key**.



b. Provide a name for it and select **Create**.

It may take about 30 seconds to create the key. Once created, the new key appears in the list of authorization keys for the private cloud.



c. Copy the authorization key and ExpressRoute ID. You'll need them to complete the peering. The authorization key disappears after some time, so copy it as soon as it appears.

2. Navigate to the virtual network gateway you plan to use and select **Connections > + Add**.

3. On the **Add connection** page, provide values for the fields, and select **OK**.

FIELD	VALUE
Name	Enter a name for the connection.
Connection type	Select ExpressRoute .
Redeem authorization	Ensure this box is selected.
Virtual network gateway	The virtual network gateway you intend to use.
Authorization key	Paste the authorization key you copied earlier.
Peer circuit URI	Paste the ExpressRoute ID you copied earlier.



Add connection

PrivateCloudGateway

Directory: Microsoft

i Ensure that the ExpressRoute associated with this authorization is provisioned by the provider before redeeming the authorization.

Name *

privatecloud-connection ✓

Connection type ⓘ

ExpressRoute ▼

Redeem authorization ⓘ

*Virtual network gateway ⓘ

PrivateCloudGateway 🔒

Authorization key *

442cb5d8 ... ✓

Peer circuit URI *

/subscriptions/750a6f9e ... ✓

Subscription ⓘ

▼

Resource group ⓘ

ContosoResourceGroup 🔒

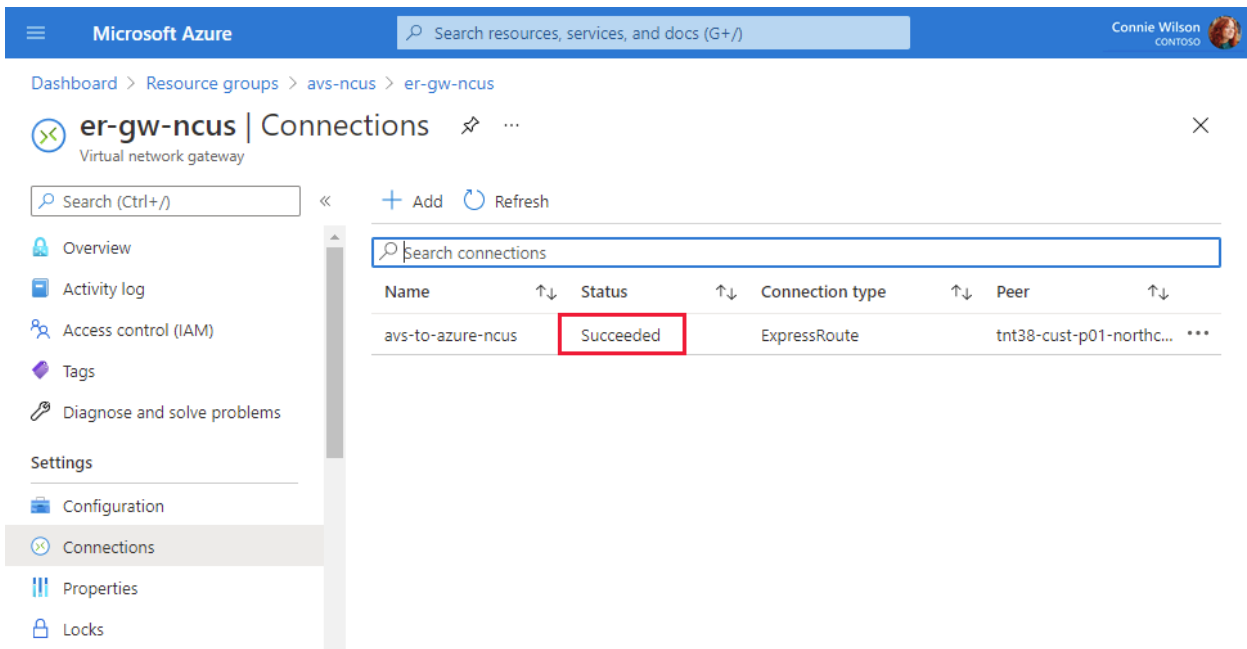
Create new

Location ⓘ

East US ▼

OK

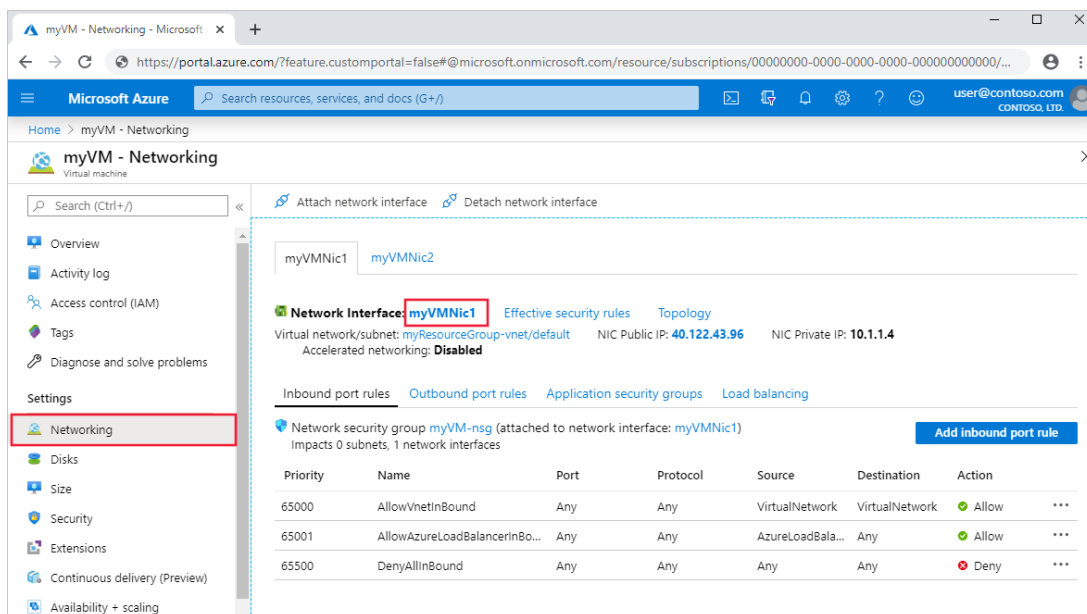
The connection between your ExpressRoute circuit and your Virtual Network is created.



Validate the connection

You should have connectivity between the Azure Virtual Network where the ExpressRoute terminates and the Azure VMware Solution private cloud.

1. Use a [virtual machine](#) within the Azure Virtual Network where the Azure VMware Solution ExpressRoute terminates. For more information, see [Connect to Azure Virtual Network with ExpressRoute](#).
 - a. Log into the [Azure portal](#).
 - b. Navigate to a VM that is in the running state, and under **Settings**, select **Networking** and select the network interface resource.



- c. On the left, select **Effective routes**. You'll see a list of address prefixes that are contained within the `/22` CIDR block you entered during the deployment phase.
2. If you want to log into both vCenter Server and NSX-T Manager, open a web browser and log into the same virtual machine used for network route validation.

You can identify the vCenter Server and NSX-T Manager console's IP addresses and credentials in the Azure portal. Select your private cloud and then **Manage > Identity**.

Microsoft Azure Search resources, services, and docs (G+)

Dashboard > avs-pc-ncus

avs-pc-ncus | Identity AVS Private cloud

Search (Ctrl+/)

- Overview
- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems

Settings

- Locks

Manage

- Connectivity
- Identity**
- Clusters

Login credentials

vCenter credentials

Web client URL	https://10.1.0.2/
Admin username	cloudadmin@vsphere.local
Admin password	[Redacted]
Certificate thumbprint	[Redacted]

NSX-T Manager credentials

Web client URL	https://10.1.0.3/
Admin username	admin
Admin password	[Redacted]
Certificate thumbprint	[Redacted]

Next steps

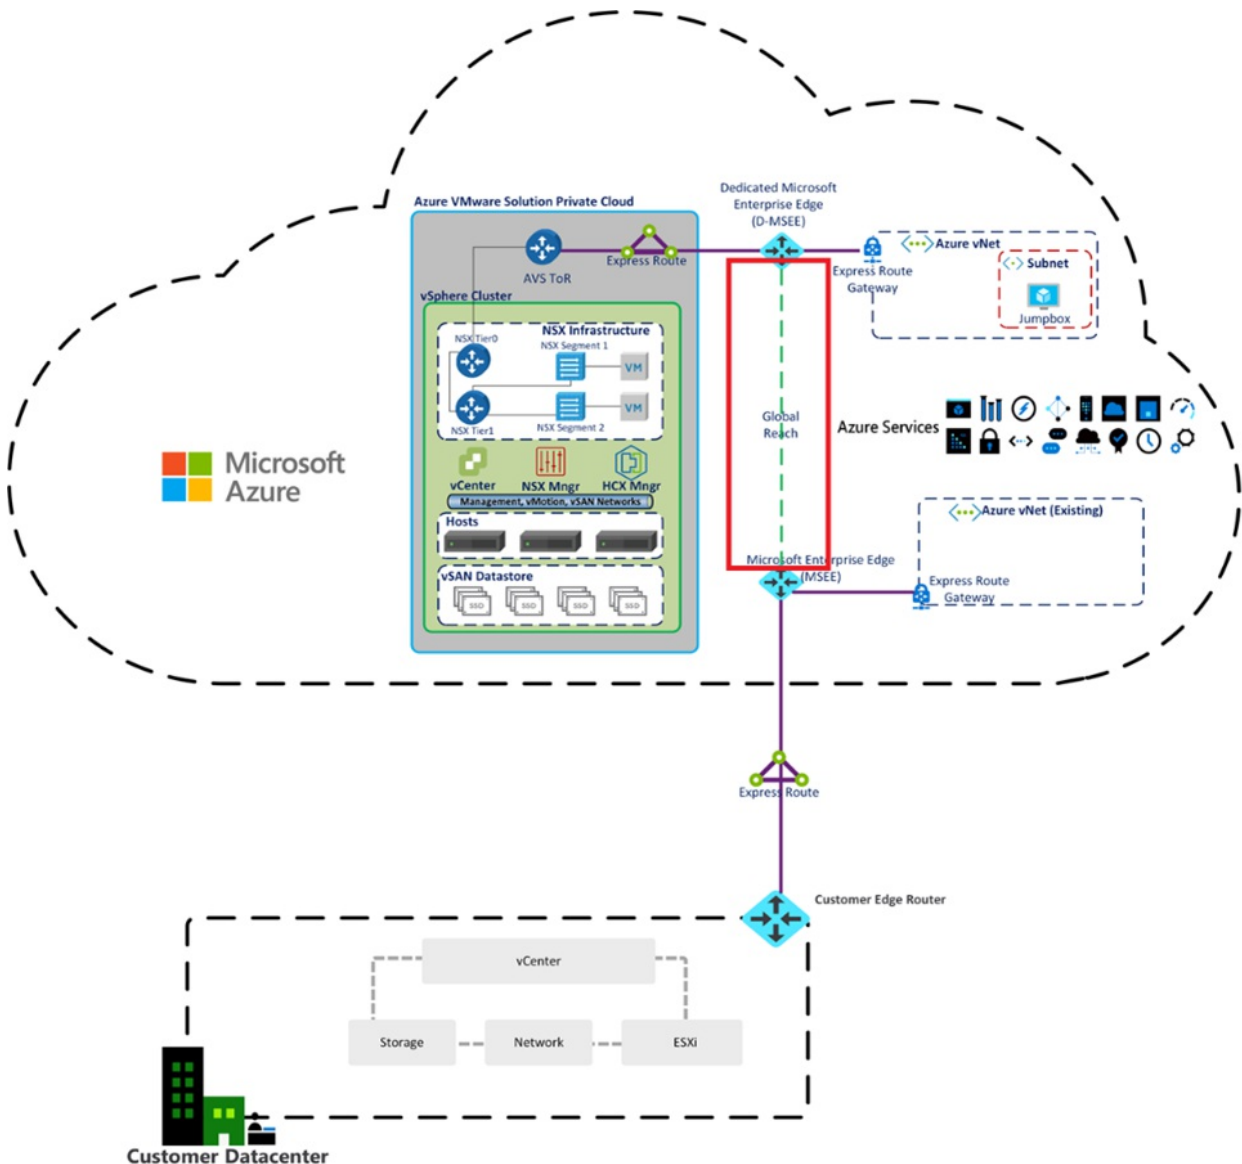
In the next tutorial, you'll connect Azure VMware Solution to your on-premises network through ExpressRoute.

[Connect to your on-premises environment](#)

Tutorial: Peer on-premises environments to Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

After you deploy your Azure VMware Solution private cloud, you'll connect it to your on-premises environment. ExpressRoute Global Reach connects your on-premises environment to your Azure VMware Solution private cloud. The ExpressRoute Global Reach connection is established between the private cloud ExpressRoute circuit and an existing ExpressRoute connection to your on-premises environments.



NOTE

You can connect through VPN, but that's out of scope for this quick start guide.

In this article, you'll:

- Create an ExpressRoute auth key in the on-premises ExpressRoute circuit
- Peer the private cloud with your on-premises ExpressRoute circuit

- Verify on-premises network connectivity

After you're finished, follow the recommended next steps at the end to continue with the steps of this getting started guide.

Prerequisites

- Review the documentation on how to [enable connectivity in different Azure subscriptions](#).
- A separate, functioning ExpressRoute circuit for connecting on-premises environments to Azure, which is *circuit 1* for peering.
- Ensure that all gateways, including the ExpressRoute provider's service, supports 4-byte Autonomous System Number (ASN). Azure VMware Solution uses 4-byte public ASNs for advertising routes.

NOTE

If advertising a default route to Azure (0.0.0.0/0), ensure a more specific route containing your on-premises networks is advertised in addition to the default route to enable management access to Azure VMware Solution. A single 0.0.0.0/0 route will be discarded by Azure VMware Solution's management network to ensure successful operation of the service.

Create an ExpressRoute auth key in the on-premises ExpressRoute circuit

The circuit owner creates an authorization, which creates an authorization key to be used by a circuit user to connect their virtual network gateways to the ExpressRoute circuit. An authorization is valid for only one connection.

NOTE

Each connection requires a separate authorization.

1. From **ExpressRoute circuits** in the left navigation, under **Settings**, select **Authorizations**.
2. Enter the name for the authorization key and select **Save**.

Home > ExpressRoute circuits > Contoso-westus-sddc

Contoso-westus-sddc | Authorizations

Save Discard Refresh

Search (Ctrl+F)

Overview
Activity log
Access control (IAM)
Tags
Diagnose and solve problems

Settings

Configuration
Connections
Authorizations
Peerings
Properties
Locks

You can create authorizations that can be redeemed by other circuit users. Circuit users are owners of virtual network gateways (that are not within the same subscription as the ExpressRoute circuit). Each authorization can be redeemed with one virtual network.

To redeem authorizations, circuit users will need the resource ID of the ExpressRoute and an unused authorization key.
[Learn more](#)

Resource ID
/subscriptions/ ... /resourceGroups/ ...

Name	Provisioning state	Use status	Authorization key
js-er-avs-auth	Succeeded	Used	b2cbd ...
js-er-az-auth	Succeeded	Available	36d91 ...

Enter name


Once created, the new key appears in the list of authorization keys for the circuit.

3. Copy the authorization key and the ExpressRoute ID. You'll use them in the next step to complete the peering.

Peer private cloud to on-premises

Now that you've created an authorization key for the private cloud ExpressRoute circuit, you can peer it with your on-premises ExpressRoute circuit. The peering is done from the on-premises ExpressRoute circuit in the **Azure portal**. You'll use the resource ID (ExpressRoute circuit ID) and authorization key of your private cloud ExpressRoute circuit to finish the peering.

1. From the private cloud, under Manage, select **Connectivity > ExpressRoute Global Reach > Add**.



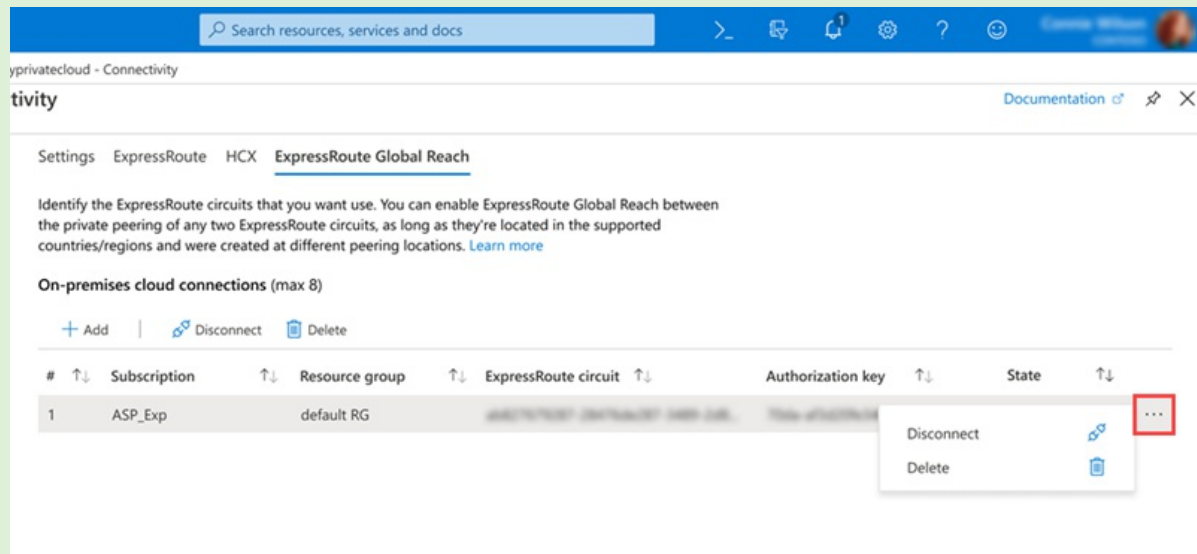
2. Enter the ExpressRoute ID and the authorization key created in the previous section.

The screenshot shows the 'On-prem cloud connections' form. It includes several input fields: 'Subscription' (set to 'Contoso'), 'Resource group' (set to 'contoso-westus-rg'), and 'ExpressRoute circuit' (empty). Below these fields is a red box containing the word 'or'. Underneath, there is a section 'If you have a circuit ID, copy/paste below' with a text input field containing 'Enter an ExpressRoute circuit ID'. Below that is an 'Authorization key' text input field. At the bottom of the form, there are two buttons: 'Create' (highlighted with a red box) and 'Cancel'.

3. Select **Create**. The new connection shows in the on-premises cloud connections list.

TIP

You can delete or disconnect a connection from the list by selecting **More**.



The screenshot shows the Azure portal interface for 'ExpressRoute Global Reach'. At the top, there's a search bar and navigation icons. Below that, the page title is 'tivity' with a 'Documentation' link. The main content area has tabs for 'Settings', 'ExpressRoute', 'HCX', and 'ExpressRoute Global Reach'. A descriptive paragraph explains that you can enable ExpressRoute Global Reach between private peering of two ExpressRoute circuits. Below this, there's a section for 'On-premises cloud connections (max 8)' with '+ Add', 'Disconnect', and 'Delete' buttons. A table lists the connections with columns: '#', 'Subscription', 'Resource group', 'ExpressRoute circuit', 'Authorization key', and 'State'. The first row shows a connection with ID '1', subscription 'ASP_Exp', and resource group 'default RG'. A red box highlights the 'More' menu icon (three dots) in the 'State' column, which has opened a dropdown menu with 'Disconnect' and 'Delete' options.

Verify on-premises network connectivity

In your **on-premises edge router**, you should now see where the ExpressRoute connects the NSX-T network segments and the Azure VMware Solution management segments.

IMPORTANT

Everyone has a different environment, and some will need to allow these routes to propagate back into the on-premises network.

Next steps

Continue to the next tutorial to install VMware HCX add-on in your Azure VMware Solution private cloud.

[Install VMware HCX](#)

Install and activate VMware HCX in Azure VMware Solution

12/16/2022 • 4 minutes to read • [Edit Online](#)

VMware HCX is an application mobility platform designed for simplifying application migration, rebalancing workloads, and optimizing disaster recovery across data centers and clouds.

VMware HCX has two component services: **HCX Cloud Manager** and **HCX Connector**. These components work together for HCX operations.

In this article, you'll learn how to install and activate the VMware HCX Cloud Manager and VMware HCX Connector components.

HCX Cloud manager is typically deployed as the destination (cloud side), but it can also be used as the source in cloud-to-cloud deployments. HCX Connector is deployed at the source (on-premises environment). A download link is provided for deploying HCX Connector appliance from within the HCX Cloud Manager.

In this how-to, you'll:

- Install VMware HCX Cloud through the Azure portal.
- Download and deploy the VMware HCX Connector in on-premises.
- Activate VMware HCX with a license key.

After HCX is deployed, follow the recommended [Next steps](#).

Prerequisite

- See [Prepare for HCX installations](#)

Install VMware HCX Cloud

1. In your Azure VMware Solution private cloud, select **Manage** > **Add-ons**.
2. Select **Get started** for **HCX Workload Mobility**.

Microsoft Azure

Search resources, services, and docs (G+)

Home > Contoso-westus-sddc

Contoso-westus-sddc | Add-ons

AVS Private cloud

Search (Ctrl+)

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Settings

Locks

Manage

Connectivity

Identity

Clusters

+ Add-ons

Overview Disaster recovery Migration using HCX

Enhance your private cloud with these optional features

Azure VMware Solutions offers various add-ons that can be enabled to provide additional functionality. Each service has its own requirements please select one to get started

Disaster Recovery

Install SRM in your private cloud to enable disaster recovery in your private cloud. SRM requires a "BYOL" license key as part of the install.

[Get Started](#)

HCX Workload Mobility

Install VMware HCX in your private cloud for workload migration. HCX advanced will be installed by default with HCX Enterprise as an optional upgrade via support.

[Get Started](#)

3. Select the **I agree with terms and conditions** checkbox and then select **Install**.

Once installed, you'll see the HCX Cloud Manager URL and the HCX keys required for the HCX on-premises connector site pairing on the **Migration using HCX** tab.

IMPORTANT

If you don't see the HCX key after installing, click the **ADD** button to generate the key which you can then use for site pairing.

Microsoft Azure

Home > Contoso-westus-sddc

Contoso-westus-sddc | Add-ons

Search (Ctrl+/)

Overview | Disaster recovery | **Migration using HCX**

HCX is an application mobility platform that is designed for simplifying application migration, workload rebalancing, and business continuity across data centers and clouds. [Learn more.](#)

HCX plan: HCX Advanced [Upgrade to HCX Enterprise \(Preview\)](#)

1. Configure HCX appliance
Using the IP address below launch the HCX portal. Download HCX appliance (OVA file) from Administration page and deploy on the site where source vCenter environment is running. [Learn more.](#)

HCX Cloud Manager IP:

2. Connect with on-premise using HCX keys
After you deploy the VMware HCX Connector appliance on-premises and start the appliance, you're ready to activate using below license keys. [Learn more.](#)

+ Add Refresh Delete

HCX key name	Activation key	Status
am	D7D9CF6583B440C7BF2B8...	Available

or

Uninstall HCX Advanced
To permanently remove all HCX components from your private cloud click the uninstall button. To downgrade to HCX Advanced edition but keep HCX please contact [support](#)

[Uninstall](#)

HCX license edition

HCX offers various [services](#) based on the type of license installed with the system. Advanced delivers basic connectivity and mobility services to enable hybrid interconnect and migration services. HCX Enterprise offers more services than what standard licenses provide, such as Mobility Groups, Replication assisted vMotion (RAV), Mobility Optimized Networking, Network Extension High availability, OS assisted Migration, etc.

NOTE

VMware HCX Enterprise is available for Azure VMware Solution customers at no additional cost.

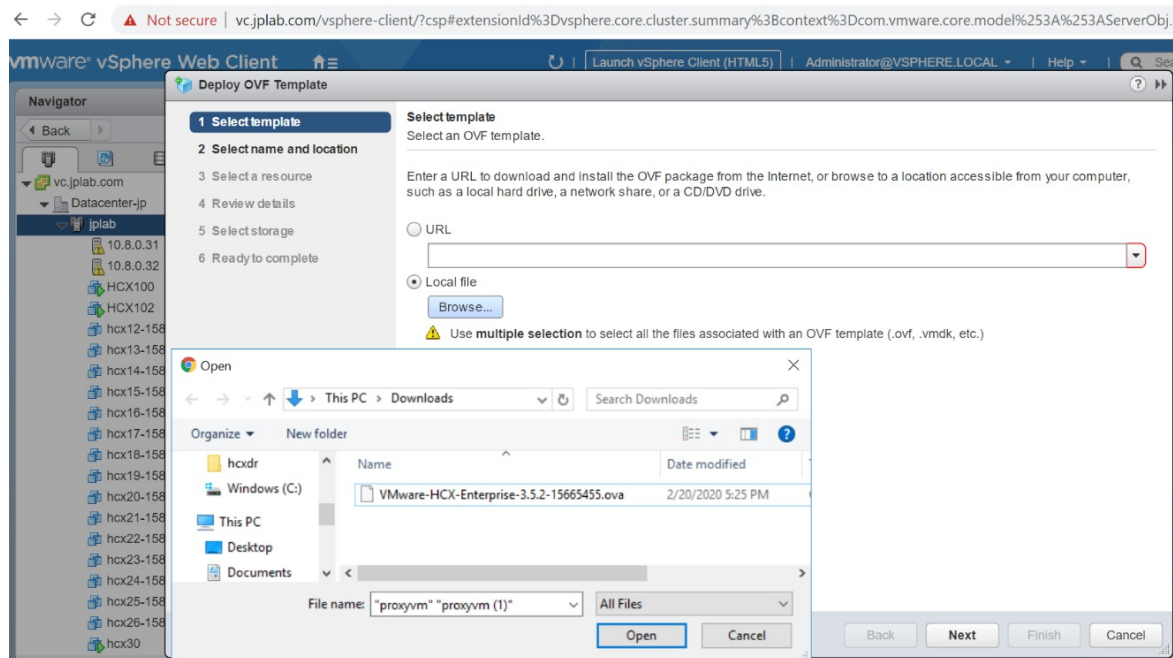
- After HCX is deployed, you can upgrade the license from Advanced to Enterprise using a [support request](#) to have HCX Enterprise Edition enabled.
- Downgrading from HCX Enterprise Edition to HCX Advanced is possible without redeploying. First, ensure you've reverted to an HCX Advanced configuration state and you aren't using the Enterprise features. If you plan to downgrade, ensure that no scheduled migrations, [Enterprise services](#) like RAV and HCX MON, etc. aren't in use. Open a [support request](#) to request downgrade.

Download and deploy the VMware HCX Connector in on-premises

In this step, you'll download the VMware HCX Connector OVA file, and then you'll deploy the VMware HCX Connector to your on-premises vCenter Server.

1. Open a browser window, sign in to the Azure VMware Solution HCX Manager on port 443 with the `cloudadmin@vsphere.local` user credentials

2. Under **Administration > System Updates**, select **Request Download Link**. If the box is greyed, wait a few seconds for it to generate a link.
3. Either download or receive a link for the VMware HCX Connector OVA file you deploy on your local vCenter Server.
4. In your on-premises vCenter Server, select an **OVF template** to deploy the VMware HCX Connector to your on-premises vSphere cluster.
5. Navigate to and select the OVA file that you downloaded and then select **Open**.



6. Select a name and location, and select a resource or cluster where you're deploying the VMware HCX Connector. Then review the details and required resources and select **Next**.
7. Review license terms, select the required storage and network, and then select **Next**.
8. Select the **VMware HCX management network segment** that you defined during the planning state. Then select **Next**.
9. In **Customize template**, enter all required information and then select **Next**.

Deploy OVF Template

✓ 1 Select an OVF template
✓ 2 Select a name and folder
✓ 3 Select a compute resource
✓ 4 Review details
✓ 5 License agreements
✓ 6 Select storage
✓ 7 Select networks
8 Customize template
9 Ready to complete

Customize template
Customize the deployment properties of this software solution.

✓ All properties have valid values

▼ Passwords	2 settings
CLI "admin" User Password	The password for default CLI user for this VM.
	Password <input type="text"/>
	Confirm Password <input type="text"/>
root Password	The password for root user.
	Password <input type="text"/>
	Confirm Password <input type="text"/>
▼ Network properties	4 settings
Hostname	The hostname for this VM.
	<input type="text"/>
Network 1 IPv4 Address	The IPv4 Address for this interface. Leave this empty for DHCP base IP assignment.
	<input type="text"/>

CANCEL BACK **NEXT**

10. Verify and then select **Finish** to deploy the VMware HCX Connector OVA.

IMPORTANT

You will need to turn on the virtual appliance manually. After powering on, wait 10-15 minutes before proceeding to the next step.

Activate VMware HCX

After deploying the VMware HCX Connector OVA on-premises and starting the appliance, you're ready to activate it. First, you'll need to get a license key from the Azure VMware Solution portal, and then you'll activate it in VMware HCX Manager. Finally, you'll need a key for each on-premises HCX connector deployed.

1. In your Azure VMware Solution private cloud, select **Manage > Add-ons > Migration using HCX**. Then copy the **Activation key**.

HCX key name	Activation key	Status
am	D7D9CF6583B440C7BF2B8...	✓ Available

2. Sign in to the on-premises VMware HCX Manager at `https://HCXManagerIP:9443` with the `admin` credentials. Make sure to include the `9443` port number with the VMware HCX Manager IP address.

TIP

You defined the **admin** user password during the VMware HCX Manager OVA file deployment.

3. In **Licensing**, enter your key for **HCX Advanced Key** and select **Activate**.

IMPORTANT

VMware HCX Manager must have open internet access or a proxy configured.

4. In **Datacenter Location**, provide the nearest location for installing the VMware HCX Manager on-premises. Then select **Continue**.
5. In **System Name**, modify the name or accept the default and select **Continue**.
6. Select **Yes, Continue**.
7. In **Connect your vCenter**, provide the FQDN or IP address of your vCenter server and the appropriate credentials, and then select **Continue**.

TIP

The vCenter Server is where you deployed the VMware HCX Connector in your datacenter.

8. In **Configure SSO/PSC**, provide your Platform Services Controller's FQDN or IP address, and select **Continue**.

NOTE

Typically, it's the same as your vCenter Server FQDN or IP address.

9. Verify that the information entered is correct and select **Restart**.

NOTE

You'll experience a delay after restarting before being prompted for the next step.

After the services restart, you'll see vCenter Server displayed as green on the screen that appears. Both vCenter Server and SSO must have the appropriate configuration parameters, which should be the same as the previous screen.

The screenshot shows the VMware HCX Manager dashboard. The top navigation bar includes 'vm HCX Manager', 'Dashboard', 'Appliance Summary', 'Configuration', and 'Administration'. The top right corner displays '10.254.11.3', 'Version: 3.5.3 Build 16639459', 'Type: Connector', and 'admin'. The main content area is divided into several sections:

- System Information:** Shows 'hcxmanager' with IP Address: 10.254.11.9, Version: 3.5.3 Build 16639459, Uptime: 17 minutes, and Current Time: Friday, 04 September 2020 08:40:17 PM UTC.
- Resource Usage:** Three progress bars show CPU (Used 1625 MHz, Free 669 MHz, Capacity 2294 MHz, 71%), Memory (Used 2271 MB, Free 9717 MB, Capacity 11989 MB, 19%), and Storage (Used 4.5G, Free 74G, Capacity 79G, 6%).
- Configuration Table:** A table with three columns: NSX, vCenter, and SSO. The vCenter and SSO rows are highlighted with a red border. The vCenter row shows 'https://10.254.11.5' with a green dot indicating it is active. The SSO row shows 'https://10.254.11.5'. Each row has a 'MANAGE' button below it.

Next steps

Continue to the next tutorial to configure the VMware HCX Connector. After you've configured the VMware HCX Connector, you'll have a production-ready environment for creating virtual machines (VMs) and migration.

[Configure VMware HCX in Azure VMware Solution](#)

[VMware blog series - cloud migration](#)

[Uninstall VMware HCX in Azure VMware Solution](#)

Configure on-premises VMware HCX Connector

12/16/2022 • 6 minutes to read • [Edit Online](#)

Once you've [installed the VMware HCX add-on](#), you're ready to configure the on-premises VMware HCX Connector for your Azure VMware Solution private cloud.

In this how-to, you'll:

- Pair your on-premises VMware HCX Connector with your Azure VMware Solution HCX Cloud Manager
- Configure the network profile, compute profile, and service mesh
- Check the appliance status and validate that migration is possible

After you complete these steps, you'll have a production-ready environment for creating virtual machines (VMs) and migration.

Prerequisites

- [VMware HCX Connector](#) has been installed.
- If you plan to use VMware HCX Enterprise, make sure you've enabled the [VMware HCX Enterprise](#) add-on through a [support request](#). VMware HCX Enterprise edition is available and supported on Azure VMware Solution, at no additional cost.
- If you plan to [enable VMware HCX MON](#), make sure you have:
 - NSX-T Data Center or vSphere Distributed Switch (vDS) on-premises for HCX Network Extension (vSphere Standard Switch not supported)
 - One or more active stretched network segment
- [VMware software version requirements](#) have been met.
- Your on-premises vSphere environment (source environment) meets the [minimum requirements](#).
- [Azure ExpressRoute Global Reach](#) is configured between on-premises and Azure VMware Solution private cloud ExpressRoute circuits.
- [All required ports](#) are open for communication between on-premises components and Azure VMware Solution private.
- [Define VMware HCX network segments](#). The primary use cases for VMware HCX are workload migrations and disaster recovery.
- [Review the VMware HCX Documentation](#) for information on using HCX.

Add a site pairing

In your data center, you can connect or pair the VMware HCX Cloud Manager in Azure VMware Solution with the VMware HCX Connector.

IMPORTANT

As per the [Azure VMware Solution limits](#) the maximum site pairs is 25 and maximum service meshes is 10 in a single HCX manager system, this includes inbound and outbound site pairings.

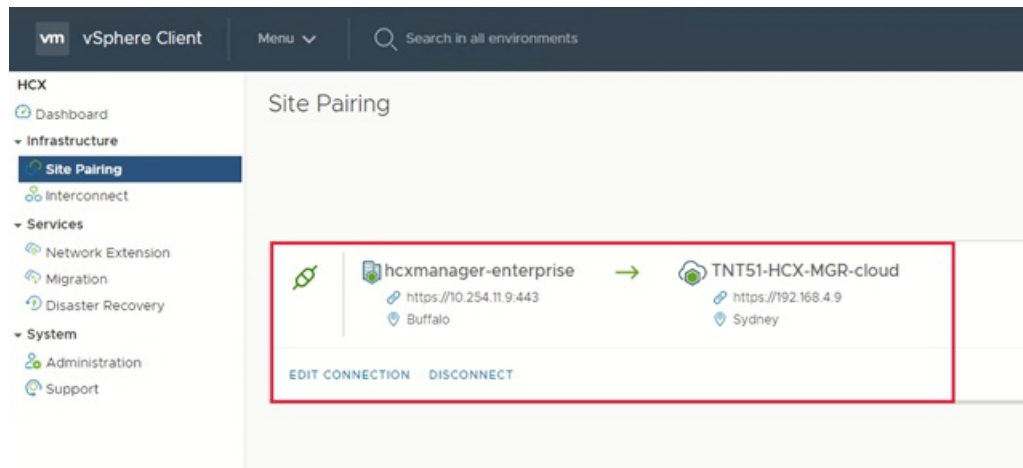
1. Sign in to your on-premises vCenter Server, and under **Home**, select **HCX**.
2. Under **Infrastructure**, select **Site Pairing** and select the **Connect To Remote Site** option (in the middle of the screen).
3. Enter the Azure VMware Solution HCX Cloud Manager URL or IP address that you noted earlier `https://x.x.x.9` and the credentials for a user which holds the CloudAdmin role in your private cloud. Then select **Connect**.

NOTE

To successfully establish a site pair:

- Your VMware HCX Connector must be able to route to your HCX Cloud Manager IP over port 443.
- A service account from your external identity source, such as Active Directory, is recommended for site pairing connections. For more information about setting up separate accounts for connected services, see [Access and Identity Concepts](#).

You'll see a screen showing that your VMware HCX Cloud Manager in Azure VMware Solution and your on-premises VMware HCX Connector are connected (paired).



Create network profiles

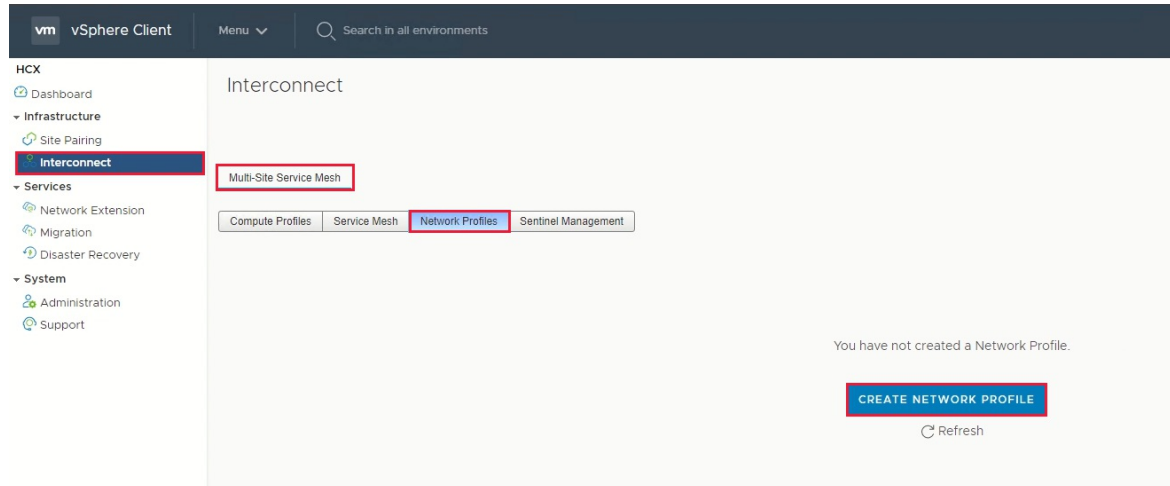
VMware HCX Connector deploys a subset of virtual appliances (automated) that require multiple IP segments. When you create your network profiles, you use the IP segments you identified during the [planning phase](#). You'll create four network profiles:

- Management
- vMotion
- Replication
- Uplink

NOTE

- Azure VMware Solution connected via VPN should set Uplink Network Profile MTU's to 1350 to account for IPSec overhead.
- Azure VMWare Solution defaults to 1500 MTU and is sufficient for most ExpressRoute implementations.
 - If your ExpressRoute provider does not support jumbo frame, MTU may need to be lowered in ExpressRoute setups as well.
 - Changes to MTU should be performed on both HCX Connector (on-premises) and HCX Cloud Manager (Azure VMware Solution) network profiles.

1. Under **Infrastructure**, select **Interconnect** > **Multi-Site Service Mesh** > **Network Profiles** > **Create Network Profile**.



2. For each network profile, select the network and port group, provide a name, and create the segment's IP pool. Then select **Create**.

Create Network Profile ✕

vCenter * 10.254.11.5

Network * Distributed Portgroup Standard Switch Network NSX Logical Switch

PortGroup	Host ID	VLAN
<input type="radio"/> vlan98	host-109	98
<input type="radio"/> VM Network	host-12,host-109	0

2 Networks

Name * _____

IP Pools

IP Pool - 0

IP Ranges	Prefix Length	Gateway
<input type="text" value="HOW MANY FREE IP ADDRESSES DO YOU NEED?"/>	_____	_____

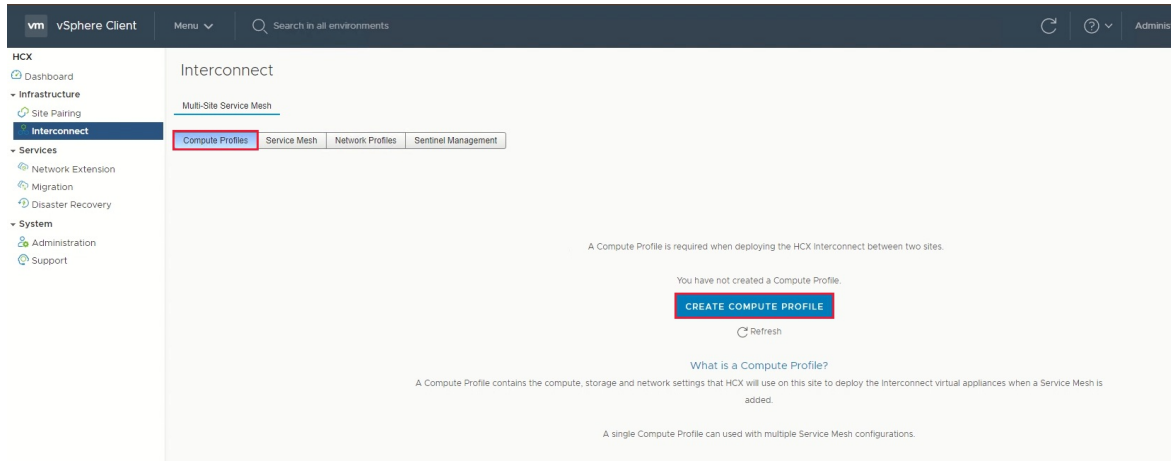
Primary DNS _____ Secondary DNS _____ DNS Suffix _____

MTU * 1500

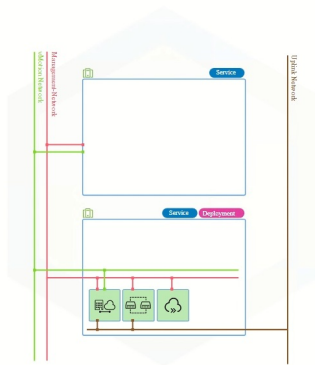

For an end-to-end overview of this procedure, view the [Azure VMware Solution: HCX Network Profile](#) video.

Create a compute profile

1. Under Infrastructure, select Interconnect > Compute Profiles > Create Compute Profile.



2. Enter a name for the profile and select Continue.



Creating a Compute Profile

During the creation of a Compute Profile, each of the following elements will be defined:

Services: The services that should be enabled on this site.

Service Clusters: The list of clusters on which the Services are to be enabled.

Deployment Resource Pool: The resource pool that should be used when deploying the Interconnect Appliances.

Deployment Storage: The storage that should be used when deploying the Interconnect Appliances.

Networks:

Management Network: The network via which the management interface of vCenter and the ESXi hosts can be reached.

Uplink Network: The network via which the Interconnect Appliances on the remote site can be reached.

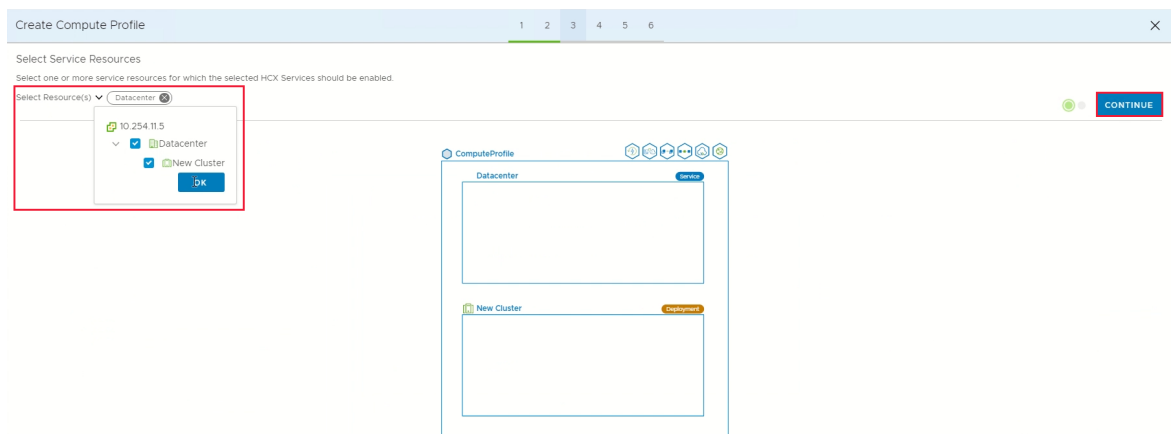
vMotion Network: The network via which the vMotion interface of the ESXi hosts can be reached.

vSphere Replication network: The network via which the vSphere Replication interface of ESXi Hosts can be reached. In most of the cases this is the same as Management Network.

Guest network: The network via which the non vSphere systems talk to OS Assisted Migration Service. In some cases this can be same as Management Network.

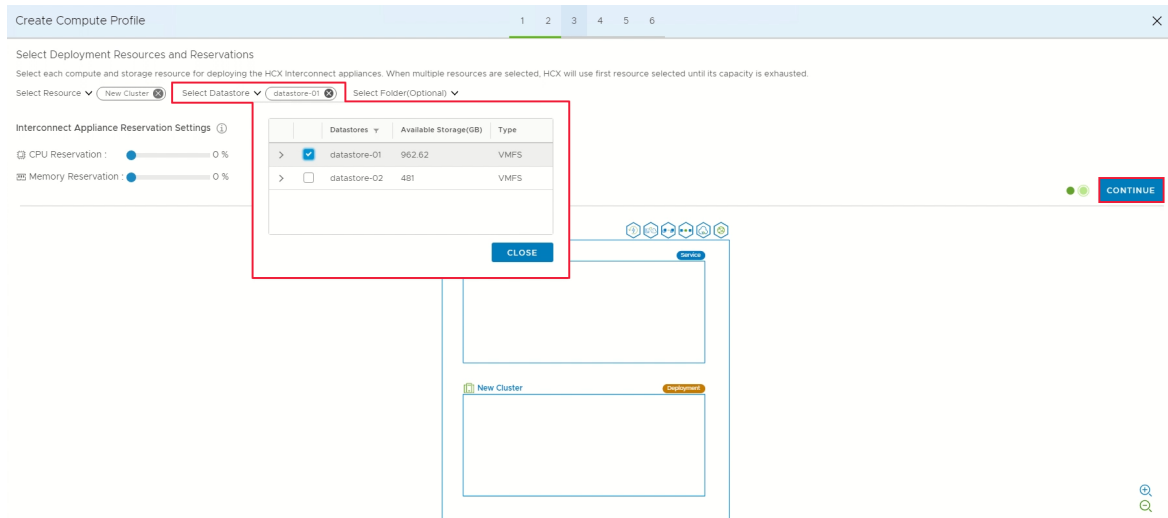
For each of the above networks, you need to have a pool of free IP Addresses, Gateway IP and subnet mask that will be used for deploying the Interconnect Appliances while creating a Service Mesh.

3. Select the services to enable, such as migration, network extension, or disaster recovery, and then select Continue.
4. In Select Service Resources, select one or more service resources (clusters) to enable the selected VMware HCX services.
5. When you see the clusters in your on-premises datacenter, select Continue.

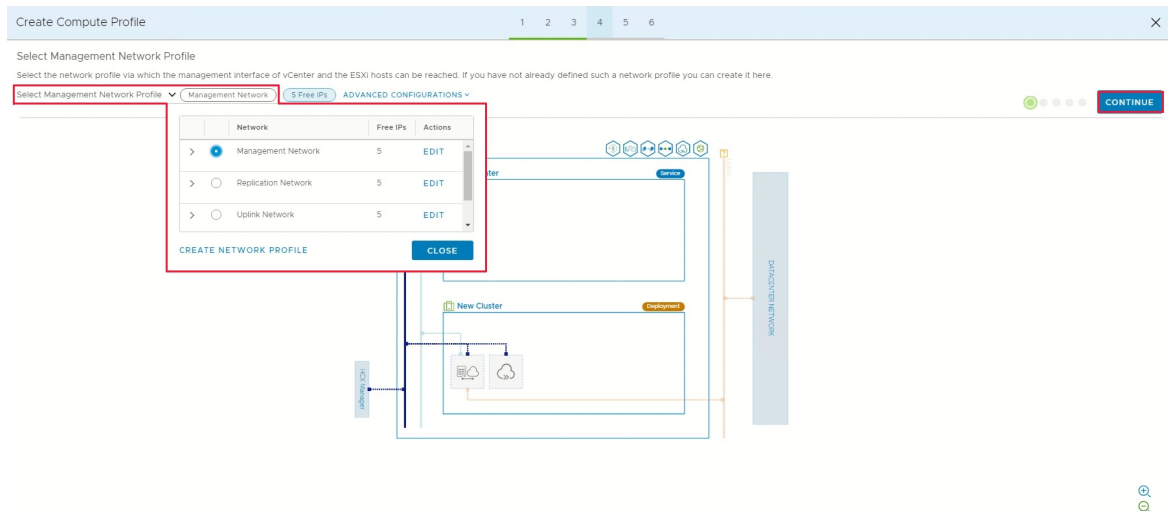


- From **Select Datastore**, select the datastore storage resource for deploying the VMware HCX Interconnect appliances. Then select **Continue**.

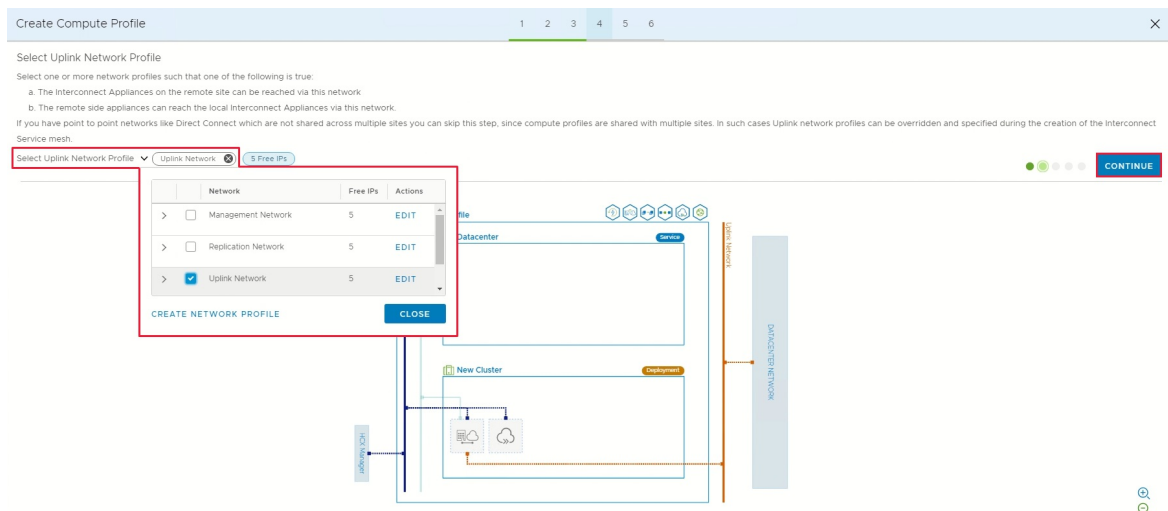
When multiple resources are selected, VMware HCX uses the first resource selected until its capacity is exhausted.



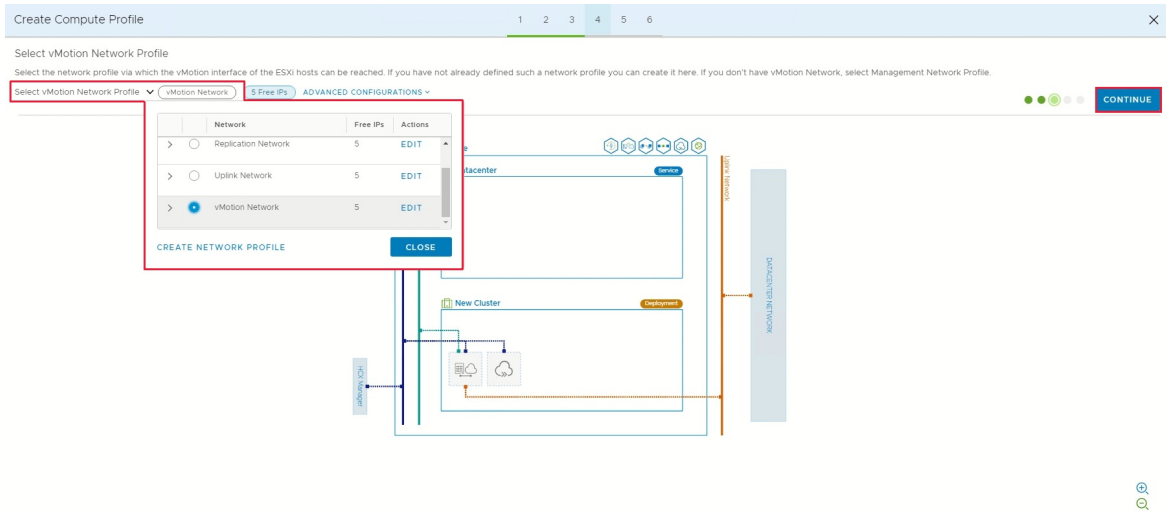
- From **Select Management Network Profile**, select the management network profile that you created in previous steps. Then select **Continue**.



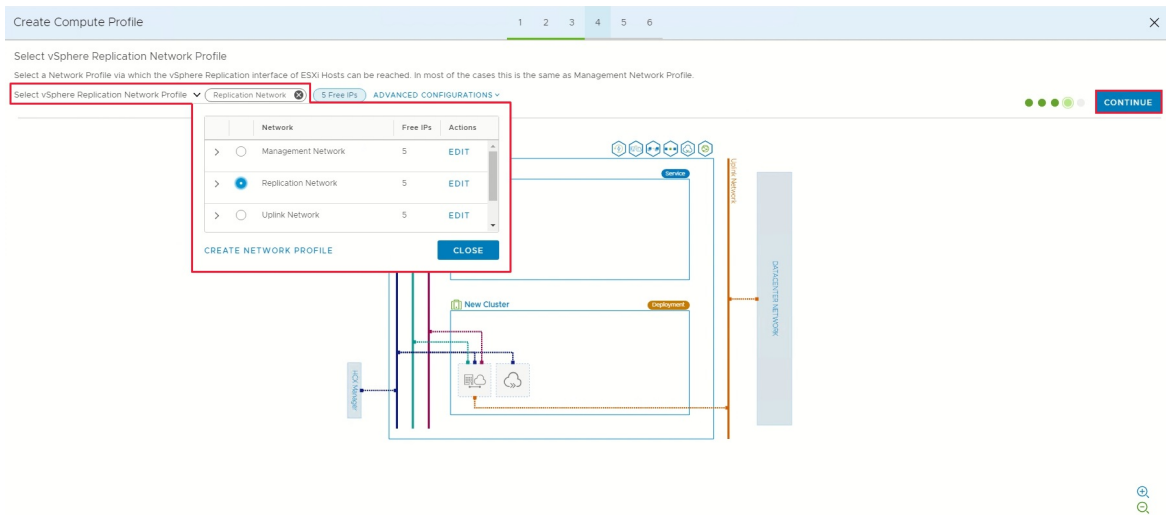
- From **Select Uplink Network Profile**, select the uplink network profile you created in the previous procedure. Then select **Continue**.



9. From **Select vMotion Network Profile**, select the vMotion network profile that you created in previous steps. Then select **Continue**.



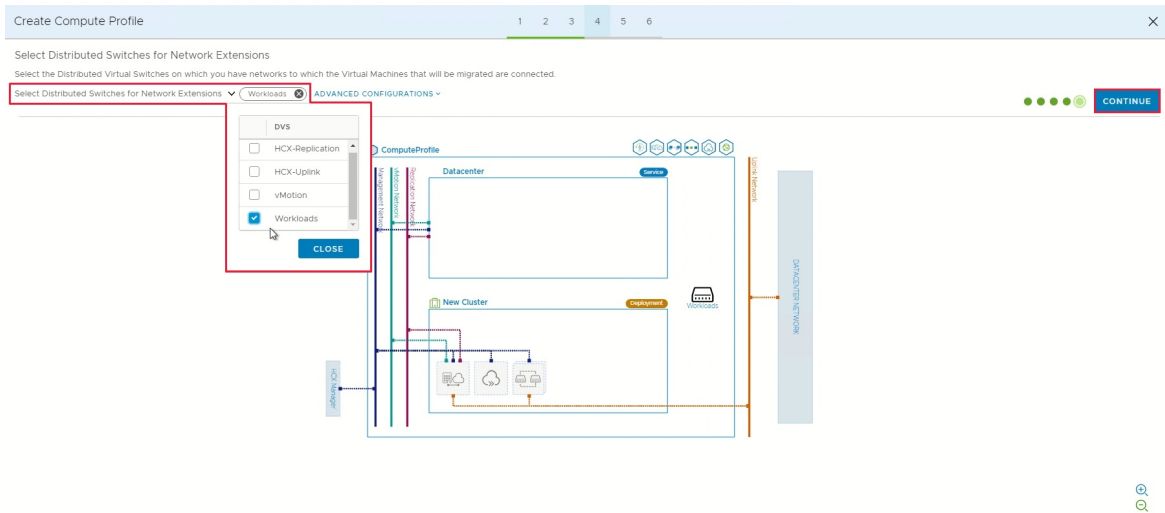
10. From **Select vSphere Replication Network Profile**, select the replication network profile that you created in previous steps. Then select **Continue**.



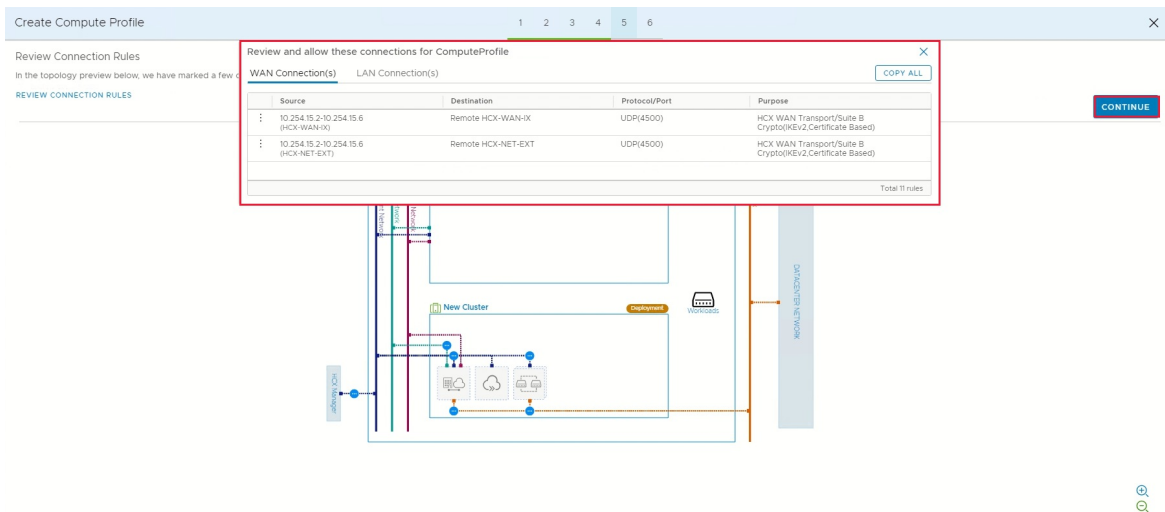
11. From **Select Distributed Switches for Network Extensions**, select the switches containing the virtual machines to be migrated to Azure VMware Solution on a layer-2 extended network. Then select **Continue**.

NOTE

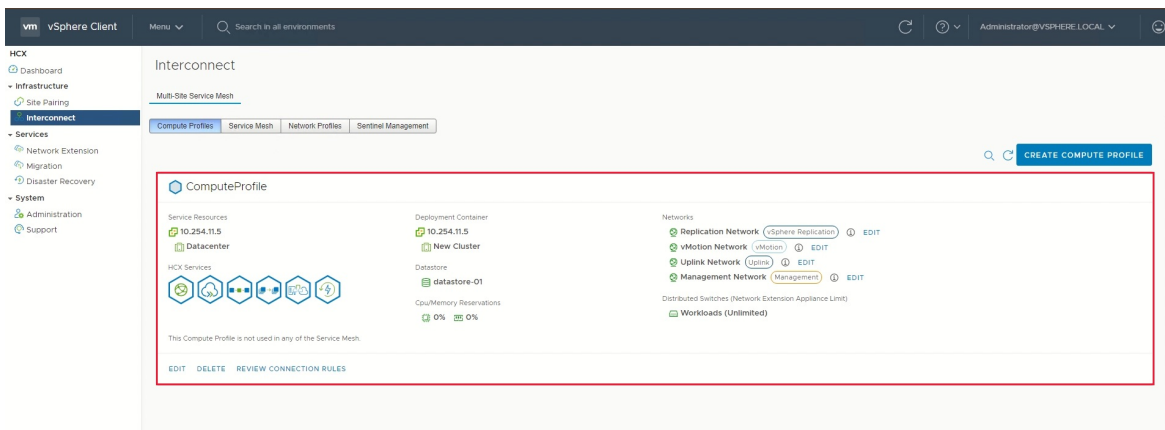
If you are not migrating virtual machines on layer-2 (L2) extended networks, you can skip this step.



12. Review the connection rules and select Continue.



13. Select Finish to create the compute profile.



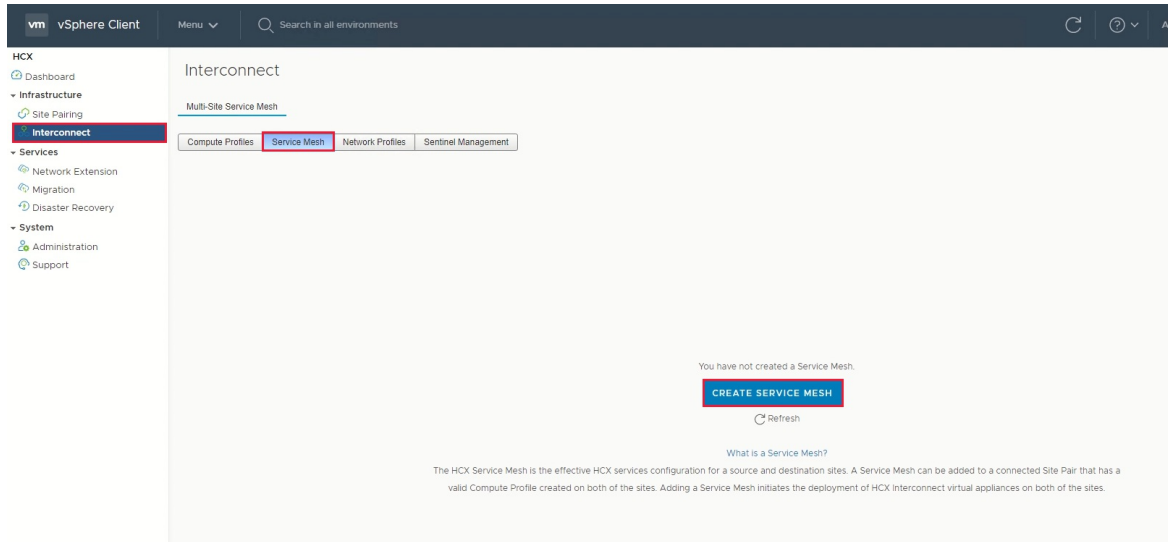
For an end-to-end overview of this procedure, view the [Azure VMware Solution: Compute Profile](#) video.

Create a service mesh

IMPORTANT

Make sure port UDP 4500 is open between your on-premises VMware HCX Connector 'uplink' network profile addresses and the Azure VMware Solution HCX Cloud 'uplink' network profile addresses. (500 UDP was previously required in legacy versions of HCX. See <https://ports.vmware.com> for latest information)

1. Under **Infrastructure**, select **Interconnect > Service Mesh > Create Service Mesh**.



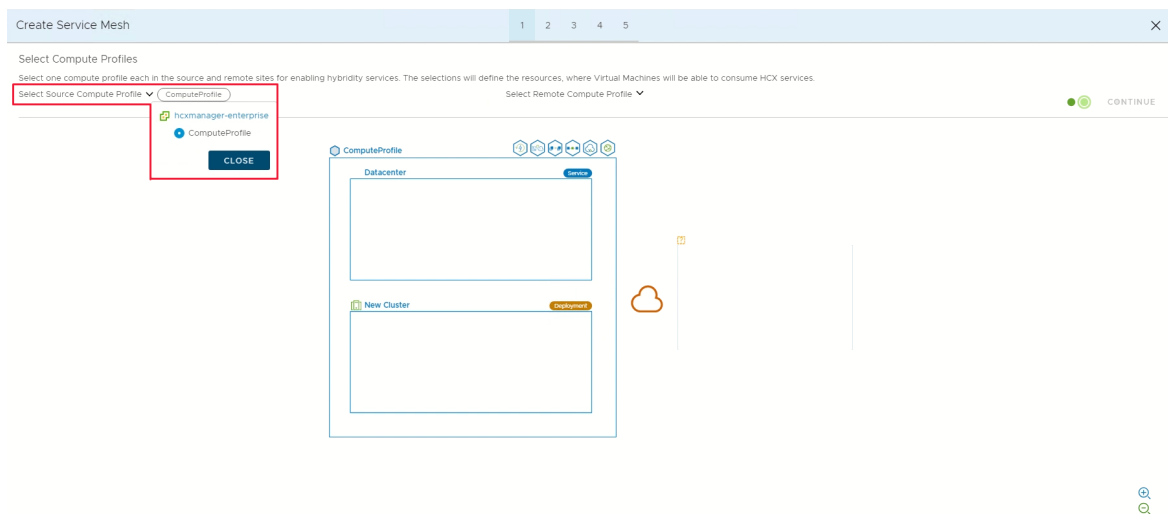
2. Review the pre-populated sites, and then select **Continue**.

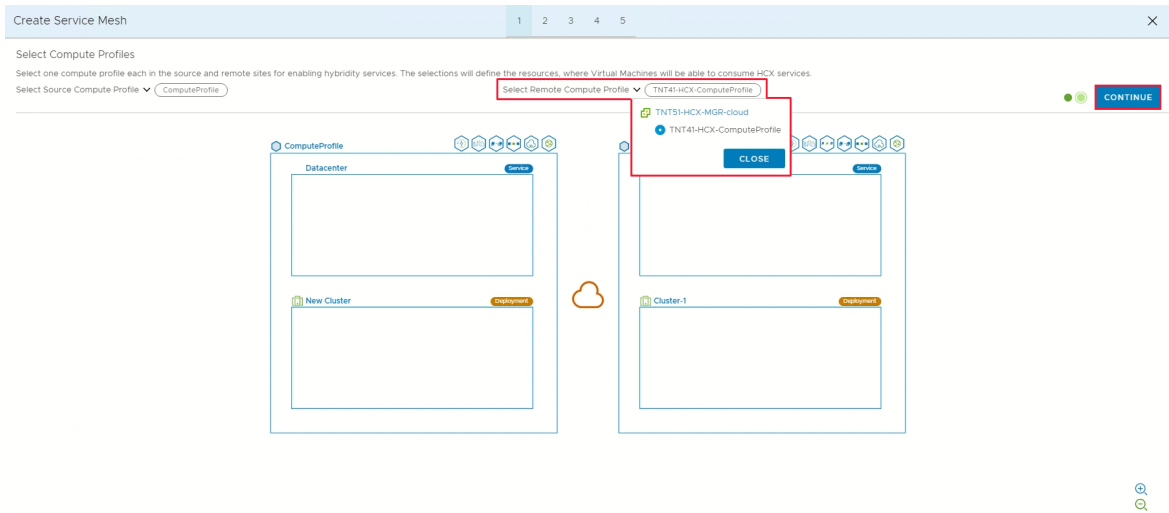
NOTE

If this is your first service mesh configuration, you won't need to modify this screen.

3. Select the source and remote compute profiles from the drop-down lists, and then select **Continue**.

The selections define the resources where VMs can consume VMware HCX services.






4. Review services that will be enabled, and then select **Continue**.

5. In **Advanced Configuration - Override Uplink Network profiles**, select **Continue**.

Uplink network profiles connect to the network through which the remote site's interconnect appliances can be reached.

6. In **Advanced Configuration - Network Extension Appliance Scale Out**, review and select **Continue**.

You can have up to eight VLANs per appliance, but you can deploy another appliance to add another eight VLANs. You must also have IP space to account for the more appliances, and it's one IP per appliance. For more information, see [VMware HCX Configuration Limits](#).

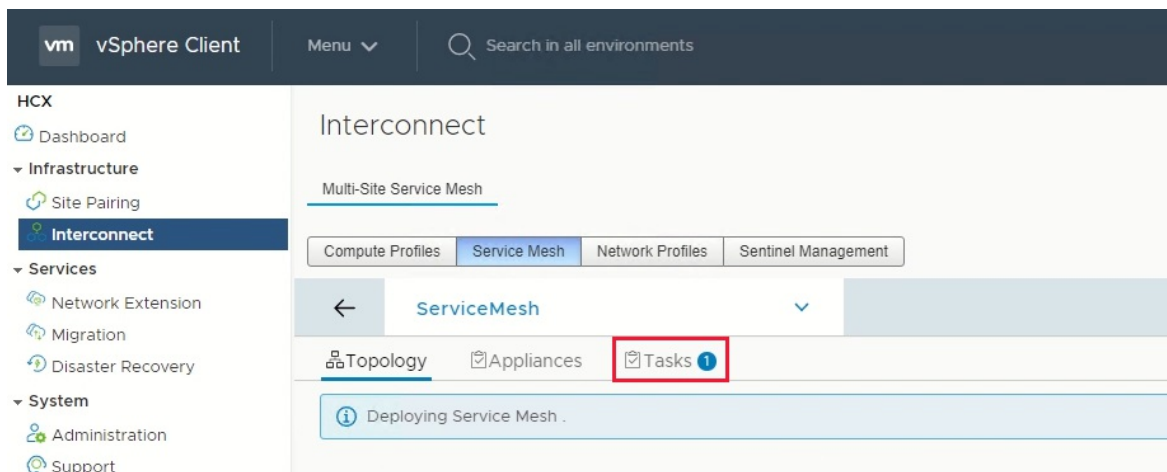


7. In **Advanced Configuration - Traffic Engineering**, review and make any modifications that you feel are necessary, and then select **Continue**.

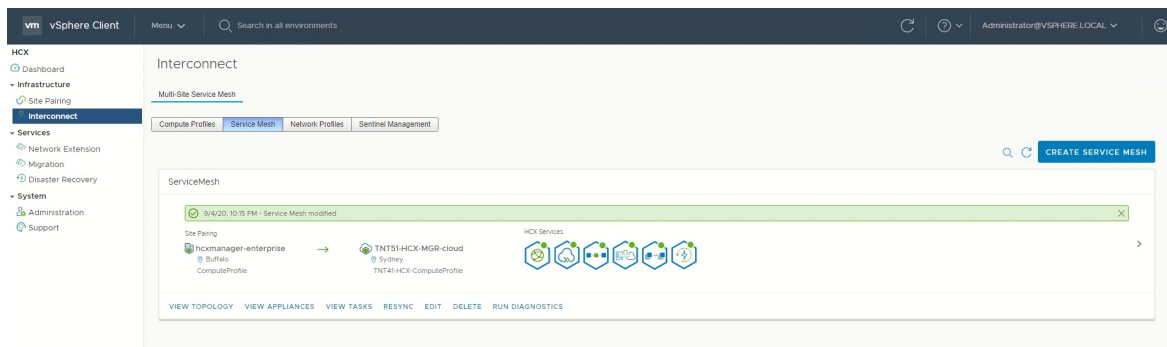
8. Review the topology preview and select **Continue**.

9. Enter a user-friendly name for this service mesh and select **Finish** to complete.

10. Select **View Tasks** to monitor the deployment.

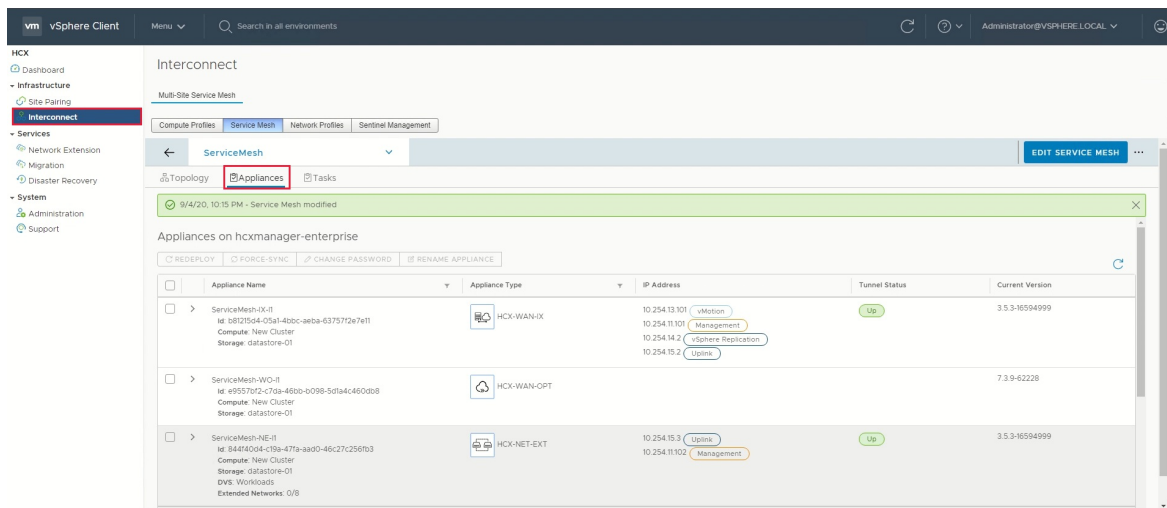


When the service mesh deployment finishes successfully, you'll see the services as green.



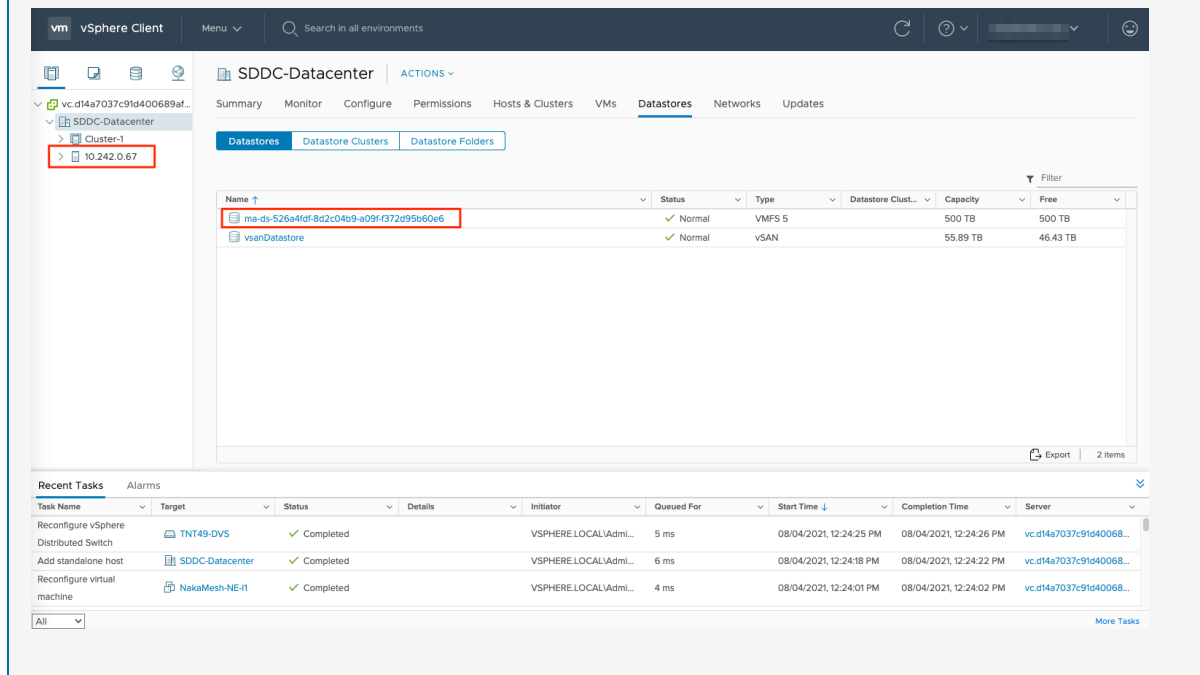
11. Verify the service mesh's health by checking the appliance status.

12. Select Interconnect > Appliances.



NOTE

After establishing the service mesh, you may notice a new datastore and a new host in your private cloud. This is perfectly normal behavior after establishing a service mesh.



The HCX interconnect tunnel status should indicate **UP** and in green. You're ready to migrate and protect Azure VMware Solution VMs using VMware HCX. Azure VMware Solution supports workload migrations (with or without a network extension). So you can still migrate workloads in your vSphere environment, along with on-premises creation of networks and deployment of VMs onto those networks. For more information, see the [VMware HCX Documentation](#).

For an end-to-end overview of this procedure, view the [Azure VMware Solution: Service Mesh](#) video.

Next steps

Now that you've configured the HCX Connector, you can also learn about:

- [Create a HCX network extension](#)
- [VMware HCX Mobility Optimized Networking \(MON\) guidance](#)

Uninstall VMware HCX in Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

In this article, you'll learn how to uninstall HCX in Azure VMware solution. You can uninstall HCX from the cloud side through the portal, which removes the existing pairing and software. Removing HCX returns the resources to your private cloud occupied by the HCX virtual appliances.

Generally, the workflow is to clean-up from the HCX on-premises side first, then clean-up on the HCX Cloud side afterwards.

Prerequisites

- Make sure you don't have any active migrations in progress.
- Ensure that L2 extensions are no longer needed or the networks have been `unstretched` to the destination.
- For workloads using MON, ensure that you've removed the default gateways. Otherwise, it may result in workloads not being able to communicate or function.
- [Uninstall HCX deployment from Connector on-premises.](#)

Uninstall HCX

1. In your Azure VMware Solution private cloud, select **Manage** > **Add-ons**.
2. Select **Get started for HCX Workload Mobility**, then select **Uninstall**.
3. Enter **yes** to confirm the uninstall.

Overview Disaster recovery Migration using HCX

HCX is an application mobility platform that is designed for simplifying application migration, workload rebalancing, and business continuity across data centers and clouds. [Learn more](#)

HCX plan

HCX Enterprise

1. Configure HCX appliance

Using the IP address below launch the HCX portal. Download HCX appliance (OVA file) from Administration page and deploy on the site where source vCenter environment is running. [Learn more](#)

HCX Cloud Manager IP

`https://192.168.192.9/`

2. Connect with on-premise using HCX keys

After you deploy the VMware HCX Connector appliance on-premises and start the appliance, you're ready to activate using below license keys. [Learn more](#)

+ Add Refresh Delete

HCX key name

Activation key

Status

or

Uninstall HCX Advanced

To permanently remove all HCX components from your private cloud click the uninstall button. To downgrade to HCX Advanced edition but keep HCX please contact [support](#).

Uninstall

After uninstalling HCX, it no longer has the vCenter Server plugin. If necessary, you can reinstall it.

[Configure VMware HCX in Azure VMware Solution](#)

[VMware blog series - cloud migration](#)

[Install and activate VMware HCX in Azure VMware Solution](#)

Networking planning checklist for Azure VMware Solution

12/16/2022 • 6 minutes to read • [Edit Online](#)

Azure VMware Solution offers a VMware private cloud environment accessible for users and applications from on-premises and Azure-based environments or resources. The connectivity is delivered through networking services such as Azure ExpressRoute and VPN connections. It requires specific network address ranges and firewall ports to enable the services. This article provides you with the information you need to properly configure your networking to work with Azure VMware Solution.

In this tutorial, you'll learn about:

- Virtual network and ExpressRoute circuit considerations
- Routing and subnet requirements
- Required network ports to communicate with the services
- DHCP and DNS considerations in Azure VMware Solution

Prerequisite

Ensure that all gateways, including the ExpressRoute provider's service, supports 4-byte Autonomous System Number (ASN). Azure VMware Solution uses 4-byte public ASNs for advertising routes.

Virtual network and ExpressRoute circuit considerations

When you create a virtual network connection in your subscription, the ExpressRoute circuit is established through peering, using an authorization key and a peering ID you request in the Azure portal. The peering is a private, one-to-one connection between your private cloud and the virtual network.

NOTE

The ExpressRoute circuit is not part of a private cloud deployment. The on-premises ExpressRoute circuit is beyond the scope of this document. If you require on-premises connectivity to your private cloud, you can use one of your existing ExpressRoute circuits or purchase one in the Azure portal.

When deploying a private cloud, you receive IP addresses for vCenter Server and NSX-T Manager. To access those management interfaces, you'll need to create more resources in your subscription's virtual network. You can find the procedures for creating those resources and establishing [ExpressRoute private peering](#) in the tutorials.

The private cloud logical networking comes with pre-provisioned NSX-T Data Center configuration. A Tier-0 gateway and Tier-1 gateway are pre-provisioned for you. You can create a segment and attach it to the existing Tier-1 gateway or attach it to a new Tier-1 gateway that you define. NSX-T Data Center logical networking components provide East-West connectivity between workloads and North-South connectivity to the internet and Azure services.

IMPORTANT

If you plan to scale your Azure VMware Solution hosts using [Azure NetApp Files datastores](#), deploying the vNet close to your hosts with an ExpressRoute virtual network gateway is crucial. The closer the storage is to your hosts, the better the performance.

Routing and subnet considerations

The Azure VMware Solution private cloud is connected to your Azure virtual network using an Azure ExpressRoute connection. This high bandwidth, low latency connection allows you to access services running in your Azure subscription from your private cloud environment. The routing is Border Gateway Protocol (BGP) based, automatically provisioned, and enabled by default for each private cloud deployment.

Azure VMware Solution private clouds require a minimum of a `/22` CIDR network address block for subnets, shown below. This network complements your on-premises networks. Therefore, the address block shouldn't overlap with address blocks used in other virtual networks in your subscription and on-premises networks. Within this address block, management, provisioning, and vMotion networks get provisioned automatically.

NOTE

Permitted ranges for your address block are the RFC 1918 private address spaces (10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16), except for 172.17.0.0/16.

Example `/22` CIDR network address block: `10.10.0.0/22`

The subnets:

NETWORK USAGE	SUBNET	EXAMPLE
Private cloud management	<code>/26</code>	<code>10.10.0.0/26</code>
HCX Mgmt Migrations	<code>/26</code>	<code>10.10.0.64/26</code>
Global Reach Reserved	<code>/26</code>	<code>10.10.0.128/26</code>
NSX-T Data Center DNS Service	<code>/32</code>	<code>10.10.0.192/32</code>
Reserved	<code>/32</code>	<code>10.10.0.193/32</code>
Reserved	<code>/32</code>	<code>10.10.0.194/32</code>
Reserved	<code>/32</code>	<code>10.10.0.195/32</code>
Reserved	<code>/30</code>	<code>10.10.0.196/30</code>
Reserved	<code>/29</code>	<code>10.10.0.200/29</code>
Reserved	<code>/28</code>	<code>10.10.0.208/28</code>
ExpressRoute peering	<code>/27</code>	<code>10.10.0.224/27</code>

NETWORK USAGE	SUBNET	EXAMPLE
ESXi Management	/25	10.10.1.0/25
vMotion Network	/25	10.10.1.128/25
Replication Network	/25	10.10.2.0/25
vSAN	/25	10.10.2.128/25
HCX Uplink	/26	10.10.3.0/26
Reserved	/26	10.10.3.64/26
Reserved	/26	10.10.3.128/26
Reserved	/26	10.10.3.192/26

Required network ports

SOURCE	DESTINATION	PROTOCOL	PORT	DESCRIPTION
Private Cloud DNS server	On-Premises DNS Server	UDP	53	DNS Client - Forward requests from Private Cloud vCenter Server for any on-premises DNS queries (check DNS section below)
On-premises DNS Server	Private Cloud DNS server	UDP	53	DNS Client - Forward requests from on-premises services to Private Cloud DNS servers (check DNS section below)
On-premises network	Private Cloud vCenter Server	TCP(HTTP)	80	vCenter Server requires port 80 for direct HTTP connections. Port 80 redirects requests to HTTPS port 443. This redirection helps if you use <code>http://server</code> instead of <code>https://server</code> .

SOURCE	DESTINATION	PROTOCOL	PORT	DESCRIPTION
Private Cloud management network	On-premises Active Directory	TCP	389/636	These ports are open to allow communications for Azure VMware Solutions vCenter Server to communicate to any on-premises Active Directory/LDAP server(s). These port(s) are optional - for configuring on-premises AD as an identity source on the Private Cloud vCenter. Port 636 is recommended for security purposes.
Private Cloud management network	On-premises Active Directory Global Catalog	TCP	3268/3269	These ports are open to allow communications for Azure VMware Solutions vCenter Server to communicate to any on-premises Active Directory/LDAP global catalog server(s). These port(s) are optional - for configuring on-premises AD as an identity source on the Private Cloud vCenter Server. Port 3269 is recommended for security purposes.
On-premises network	Private Cloud vCenter Server	TCP(HTTPS)	443	This port allows you to access vCenter Server from an on-premises network. The default port that the vCenter Server system uses to listen for connections from the vSphere Client. To enable the vCenter Server system to receive data from the vSphere Client, open port 443 in the firewall. The vCenter Server system also uses port 443 to monitor data transfer from SDK clients.

SOURCE	DESTINATION	PROTOCOL	PORT	DESCRIPTION
On-premises network	HCX Manager	TCP(HTTPS)	9443	Hybrid Cloud Manager Virtual Appliance Management Interface for Hybrid Cloud Manager system configuration.
Admin Network	Hybrid Cloud Manager	SSH	22	Administrator SSH access to Hybrid Cloud Manager.
HCX Manager	Interconnect (HCX-IX)	TCP(HTTPS)	8123	HCX Bulk Migration Control
HCX Manager	Interconnect (HCX-IX), Network Extension (HCX-NE)	HTTP TCP(HTTPS)	9443	Send management instructions to the local HCX Interconnect using the REST API.
Interconnect (HCX-IX)	L2C	TCP(HTTPS)	443	Send management instructions from Interconnect to L2C when L2C uses the same path as the Interconnect.
HCX Manager, Interconnect (HCX-IX)	ESXi Hosts	TCP	80,902	Management and OVF deployment.
HCX NE, Interconnect (HCX-IX) at Source	HCX NE, Interconnect (HCX-IX) at Destination	UDP	4500	Required for IPSEC Internet key exchange (IKEv2) to encapsulate workloads for the bidirectional tunnel. Network Address Translation-Traversal (NAT-T) is also supported.
Interconnect (HCX-IX) local	Interconnect (HCX-IX) (remote)	UDP	500	Required for IPSEC Internet key exchange (ISAKMP) for the bidirectional tunnel.
On-premises vCenter Server network	Private Cloud management network	TCP	8000	vMotion of VMs from on-premises vCenter Server to Private Cloud vCenter Server

[For a full list of HCX port requirements](#)

DHCP and DNS resolution considerations

Applications and workloads running in a private cloud environment require name resolution and DHCP services for lookup and IP address assignments. A proper DHCP and DNS infrastructure are required to provide these services. You can configure a virtual machine to provide these services in your private cloud environment.

Use the DHCP service built-in to NSX or use a local DHCP server in the private cloud instead of routing broadcast DHCP traffic over the WAN back to on-premises.

Next steps

In this tutorial, you learned about the considerations and requirements for deploying an Azure VMware Solution private cloud. Once you have the proper networking in place, continue to the next tutorial to create your Azure VMware Solution private cloud.

[Create an Azure VMware Solution private cloud](#)

Tutorial: Deploy an Azure VMware Solution private cloud

12/16/2022 • 5 minutes to read • [Edit Online](#)

The Azure VMware Solution private gives you the ability to deploy a vSphere cluster in Azure. For each private cloud created, there's one vSAN cluster by default. You can add, delete, and scale clusters. The minimum number of hosts per cluster is three. More hosts can be added one at a time, up to a maximum of 16 hosts per cluster. The maximum number of clusters per private cloud is 12. The initial deployment of Azure VMware Solution has three hosts.

You use vCenter Server and NSX-T Manager to manage most other aspects of cluster configuration or operation. All local storage of each host in a cluster is under the control of vSAN.

TIP

You can always extend the cluster and add more clusters later if you need to go beyond the initial deployment number.

Because Azure VMware Solution doesn't allow you to manage your private cloud with your cloud vCenter Server at launch, you'll need to do more steps for the configuration. This tutorial covers these steps and related prerequisites.

In this tutorial, you'll learn how to:

- Create an Azure VMware Solution private cloud
- Verify the private cloud deployed

Prerequisites

- Appropriate administrative rights and permission to create a private cloud. You must be at minimum contributor level in the subscription.
- Follow the information you gathered in the [planning](#) tutorial to deploy Azure VMware Solution.
- Ensure you have the appropriate networking configured as described in the [Network planning checklist](#).
- Hosts provisioned and the Microsoft.AVS [resource provider has been registered](#).

Create a private cloud

You can create an Azure VMware Solution private cloud using the Azure portal or the Azure CLI.

- [Portal](#)
- [Azure CLI](#)

1. Sign in to the [Azure portal](#).
2. Select **Create a resource**.
3. In the **Search services and marketplace** text box, type `Azure VMware Solution` and select it from the search results.
4. On the **Azure VMware Solution** window, select **Create**.
5. If you need more hosts, [request a host quota increase](#).

6. On the **Basics** tab, enter values for the fields and then select **Review + Create**.

TIP

You gathered this information during the [planning phase](#) of this quick start.

FIELD	VALUE
Subscription	Select the subscription you plan to use for the deployment. All resources in an Azure subscription are billed together.
Resource group	Select the resource group for your private cloud. An Azure resource group is a logical container into which Azure resources are deployed and managed. Alternatively, you can create a new resource group for your private cloud.
Resource name	Provide the name of your Azure VMware Solution private cloud.
Location	Select a location, such as east us . It's the <i>region</i> you defined during the planning phase.
Size of host	Select the AV36 , AV36P or AV52 SKU.
Number of hosts	Number of hosts allocated for the private cloud cluster. The default value is 3, which you can increase or decrease after deployment. If these nodes are not listed as available, please contact support to request a quota increase . You can also click the link labeled If you need more hosts, request a quota increase in the Azure portal.
Address block for private cloud	Provide an IP address block for the private cloud. The CIDR represents the private cloud management network and is used for the cluster management services, such as vCenter Server and NSX-T Manager. Use /22 address space, for example, 10.175.0.0/22. The address should be unique and not overlap with other Azure Virtual Networks and with on-premises networks.

Microsoft Azure Search resources, services, and docs (G+)

Home > Create a resource > Marketplace > Azure VMware Solution >

Create a private cloud

Prerequisites *Basics Tags Review and Create

Project details

Subscription * ⓘ AnyBuild-InternalProdClusters

Resource group * ⓘ [Create new](#)

Private cloud details

Resource name * ⓘ

Location * ⓘ East US

Size of host * ⓘ

Number of hosts ⓘ 3

[Find out how many hosts you need](#)
If you need more hosts, request a quota increase

[Review and Create](#) [Previous](#) [Next : Tags >](#)

7. Verify the information entered, and if correct, select **Create**.

NOTE

This step takes roughly 3-4 hours. Adding a single host in an existing or the same cluster takes between 30 - 45 minutes.

8. Verify that the deployment was successful. Navigate to the resource group you created and select your private cloud. You'll see the status of **Succeeded** when the deployment has finished.

Microsoft Azure Search resources, services, and docs (G+)

Home >

Contoso-westus-sddc AVS Private cloud

Search (Ctrl+/) Delete

Overview

- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems

Settings

- Locks

Essentials

Resource group (change) : Contoso-westus-rg

Status : Succeeded

Location : West US

Subscription (change) : Contoso

Subscription ID : 1234abc-d567-8910-abdc-2e2bb12345e6

Tags (change) : [Click here to add tags](#)

Next steps

In this tutorial, you've learned how to:

- Create an Azure VMware Solution private cloud

- Verify the private cloud deployed
- Delete an Azure VMware Solution private cloud

Continue to the next tutorial to learn how to create a jump box. You use the jump box to connect to your environment to manage your private cloud locally.

[Access an Azure VMware Solution private cloud](#)

Tutorial: Configure networking for your VMware private cloud in Azure

12/16/2022 • 6 minutes to read • [Edit Online](#)

An Azure VMware Solution private cloud requires an Azure Virtual Network. Because Azure VMware Solution doesn't support your on-premises vCenter Server, you'll need to do additional steps to integrate with your on-premises environment. Setting up an ExpressRoute circuit and a virtual network gateway is also required.

If you plan to scale your Azure VMware Solution hosts using [Azure NetApp Files datastores](#), deploying the vNet close to your hosts with an ExpressRoute virtual network gateway is crucial. The closer the storage is to your hosts, the better the performance.

In this tutorial, you learn how to:

- Create a virtual network
- Create a virtual network gateway
- Connect your ExpressRoute circuit to the gateway

NOTE

Before you create a new vNet, evaluate if you already have an existing vNet in Azure and plan to use it to connect to Azure VMware Solution; or whether to create a new vNet entirely.

- To use an existing vNet in same Azure subscription as Azure VMware Solution, use the [Azure vNet connect](#) tab under **Connectivity**.
- To use an existing vNet in a different Azure subscription than Azure VMware Solution, use the guidance on [Connect to the private cloud manually](#).
- To create a new vNet in same Azure subscription as Azure VMware Solution, use the [Azure vNet connect](#) tab or create one [manually](#).

Connect with the Azure vNet connect feature

You can use the [Azure vNet connect](#) feature to use an existing vNet or create a new vNet to connect to Azure VMware Solution. [Azure vNet connect](#) is a function to configure vNet connectivity, it does not record configuration state; browse the Azure portal to check what settings have been configured.

NOTE

Address space in the vNet cannot overlap with the Azure VMware Solution private cloud CIDR.

Prerequisites

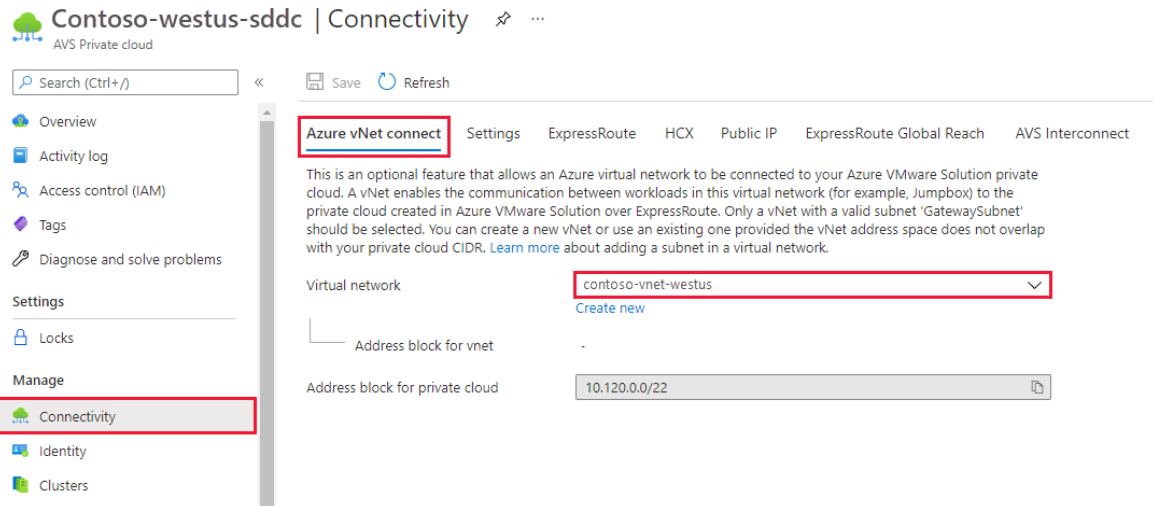
Before selecting an existing vNet, there are specific requirements that must be met.

1. vNet must contain a gateway subnet.
2. In the same region as Azure VMware Solution private cloud.
3. In the same resource group as Azure VMware Solution private cloud.
4. vNet must contain an address space that doesn't overlap with Azure VMware Solution.
5. Validate solution design is within Azure VMware Solution limits (Microsoft technical documentation/[azure/azure-resource-manager/management/azure-subscription-service-limits](#)).

Select an existing vNet

When you select an existing vNet, the Azure Resource Manager (ARM) template that creates the vNet and other resources gets redeployed. The resources, in this case, are the public IP, gateway, gateway connection, and ExpressRoute authorization key. If everything is set up, the deployment won't change anything. However, if anything is missing, it gets created automatically. For example, if the GatewaySubnet is missing, then it gets added during the deployment.

1. In your Azure VMware Solution private cloud, under **Manage**, select **Connectivity**.
2. Select the **Azure vNet connect** tab and then select the existing vNet.



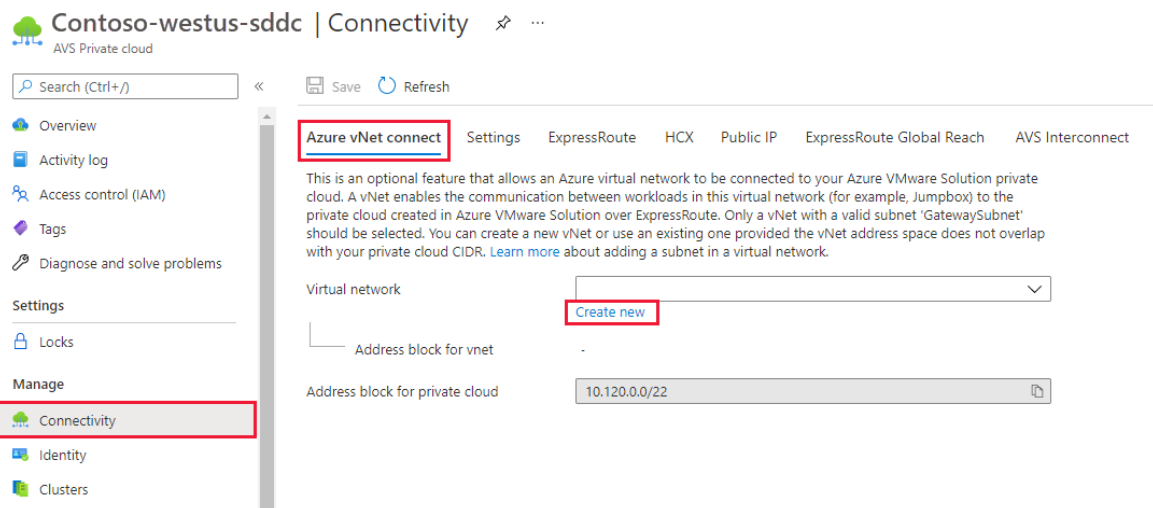
3. Select **Save**.

At this point, the vNet validates if overlapping IP address spaces between Azure VMware Solution and vNet are detected. If detected, change the network address of either the private cloud or the vNet so they don't overlap.

Create a new vNet

When you create a new vNet, the required components to connect to Azure VMware Solution are automatically created.

1. In your Azure VMware Solution private cloud, under **Manage**, select **Connectivity**.
2. Select the **Azure vNet connect** tab and then select **Create new**.



3. Provide or update the information for the new vNet and then select **OK**.

At this point, the vNet validates if overlapping IP address spaces between Azure VMware Solution and

vNet are detected. If detected, change the private cloud or vNet's network address so they don't overlap.

Create virtual network

×

This virtual network enables the communication between workloads in this virtual network (e.g. a JumpHost) to the private cloud created in Azure VMware Solution over an Express route. A default address range and a subnet is selected for this virtual network. For changing the default address range and subnet of this virtual network, follow these steps. Step 1: Change the "Address Range" to desired range (e.g. 172.16.0.0/16). Step 2: Add a subnet under "Subnets" with the name as "GatewaySubnet" and provide subnet's address range in CIDR notation (e.g. 172.16.1.0/24). [Learn more about virtual networks](#)

Name * ✓

Address space

The virtual network's address space specified as one or more address prefixes in CIDR notation (e.g. 10.0.0.0/16).

<input type="checkbox"/>	Address range	Addresses	Overlap	
<input type="checkbox"/>	<input type="text" value="172.24.0.0/16"/>	172.24.0.4 - 172.24.255.254 (65531 addresses)	None	
<input type="checkbox"/>	<input type="text"/>	(0 Addresses)	None	

Subnets

The subnet's address range in CIDR notation (e.g. 10.0.0.0/24). It must be contained by the address space of the virtual network.

<input type="checkbox"/>	Subnet name	Address range	Addresses	
<input type="checkbox"/>	GatewaySubnet	172.24.0.0/24	172.24.0.4 - 172.24.0.254 (251 addresses)	
<input type="checkbox"/>	<input type="text"/>	<input type="text"/>	(0 Addresses)	

OK

Discard

The vNet with the provided address range and GatewaySubnet is created in your subscription and resource group.

Connect to the private cloud manually

Create a vNet manually

1. Sign in to the [Azure portal](#).
2. Navigate to the resource group you created in the [create a private cloud tutorial](#) and select **+ Add** to define a new resource.
3. In the **Search the Marketplace** text box, type **Virtual Network**. Find the Virtual Network resource and select it.
4. On the **Virtual Network** page, select **Create** to set up your virtual network for your private cloud.
5. On the **Create Virtual Network** page, enter the details for your virtual network.
6. On the **Basics** tab, enter a name for the virtual network, select the appropriate region, and select **Next : IP Addresses**.
7. On the **IP Addresses** tab, under **IPv4 address space**, enter the address space you created in the

previous tutorial.

IMPORTANT

You must use an address space that **does not** overlap with the address space you used when you created your private cloud in the preceding tutorial.

8. Select **+ Add subnet**, and on the **Add subnet** page, give the subnet a name and appropriate address range. When complete, select **Add**.
9. Select **Review + create**.

Home > Virtual Network >

Create virtual network ...

Validation passed

Basics IP Addresses Security Tags **Review + create**

Basics

Subscription	Contoso
Resource group	contoso-uswest-rg
Name	contoso-uswest-vnet
Region	West US

IP addresses

Address space	10.0.0.0/16
Subnet	default (10.0.0.0/24)

Tags

None

Security

BastionHost	Disabled
DDoS protection plan	Basic
Firewall	Disabled

Create < Previous Next > [Download a template for automation](#)

10. Verify the information and select **Create**. Once the deployment is complete, you'll see your virtual network in the resource group.

Create a virtual network gateway

Now that you've created a virtual network, you'll create a virtual network gateway.

1. In your resource group, select **+ Add** to add a new resource.
2. In the **Search the Marketplace** text box, type **Virtual network gateway**. Find the Virtual Network resource and select it.
3. On the **Virtual Network gateway** page, select **Create**.
4. On the Basics tab of the **Create virtual network gateway** page, provide values for the fields, and then select **Review + create**.

FIELD	VALUE
Subscription	Pre-populated value with the Subscription to which the resource group belongs.
Resource group	Pre-populated value for the current resource group. Value should be the resource group you created in a previous test.
Name	Enter a unique name for the virtual network gateway.
Region	Select the geographical location of the virtual network gateway.
Gateway type	Select ExpressRoute .
SKU	Leave the default value: standard .
Virtual network	Select the virtual network you created previously. If you don't see the virtual network, make sure the gateway's region matches the region of your virtual network.
Gateway subnet address range	This value is populated when you select the virtual network. Don't change the default value.
Public IP address	Select Create new .

Create virtual network gateway ...

Basics Tags Review + create

Azure has provided a planning and design guide to help you configure the various VPN gateway options. [Learn more.](#)

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription *

Resource group ⓘ Select a virtual network to get resource group

Instance details

Name *

Region *

Gateway type * ⓘ VPN ExpressRoute

SKU * ⓘ

Virtual network * ⓘ

[Create virtual network](#)

ⓘ Only virtual networks in the currently selected subscription and region are listed.

Public IP address

Public IP address * ⓘ Create new Use existing

Public IP address name *

Public IP address SKU Basic

Assignment Dynamic Static

Azure recommends using a validated VPN device with your virtual network gateway. To view a list of validated devices and instructions for configuration, refer to Azure's [documentation](#) regarding validated VPN devices.

[Review + create](#)

[Previous](#)

[Next: Tags >](#)

[Download a template for automation](#)

- Verify that the details are correct, and select **Create** to start your virtual network gateway deployment.
- Once the deployment completes, move to the next section to connect your ExpressRoute connection to the virtual network gateway containing your Azure VMware Solution private cloud.

Connect ExpressRoute to the virtual network gateway

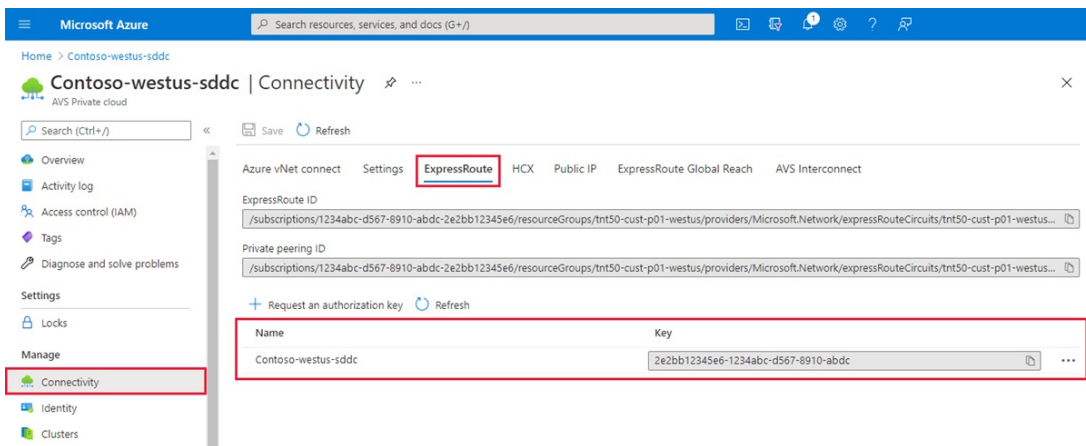
Now that you've deployed a virtual network gateway, you'll add a connection between it and your Azure VMware Solution private cloud.

- Request an ExpressRoute authorization key:
 - In the Azure portal, navigate to the Azure VMware Solution private cloud. Select **Manage > Connectivity > ExpressRoute** and then select **+ Request an authorization key**.

The screenshot shows the Azure portal interface for an Azure VMware Solution private cloud. The breadcrumb navigation is 'Home > Contoso-westus-sddc'. The main heading is 'Contoso-westus-sddc | Connectivity'. The left sidebar shows the 'Connectivity' menu item highlighted. The main content area shows the 'ExpressRoute' configuration page. The 'ExpressRoute' tab is selected, and the 'Request an authorization key' button is highlighted. The page displays the ExpressRoute ID and Private peering ID, both pointing to the same resource. Below this, there is a table with columns for 'Name' and 'Key'. The table contains one entry with the name 'Contoso-westus-sddc' and the key '2e2bb12345e6-1234abc-d567-8910-abdc'.

b. Provide a name for it and select **Create**.

It may take about 30 seconds to create the key. Once created, the new key appears in the list of authorization keys for the private cloud.



c. Copy the authorization key and ExpressRoute ID. You'll need them to complete the peering. The authorization key disappears after some time, so copy it as soon as it appears.

2. Navigate to the virtual network gateway you plan to use and select **Connections > + Add**.

3. On the **Add connection** page, provide values for the fields, and select **OK**.

FIELD	VALUE
Name	Enter a name for the connection.
Connection type	Select ExpressRoute .
Redeem authorization	Ensure this box is selected.
Virtual network gateway	The virtual network gateway you intend to use.
Authorization key	Paste the authorization key you copied earlier.
Peer circuit URI	Paste the ExpressRoute ID you copied earlier.



Add connection

PrivateCloudGateway

Directory: Microsoft

i Ensure that the ExpressRoute associated with this authorization is provisioned by the provider before redeeming the authorization.

Name *

privatecloud-connection ✓

Connection type ⓘ

ExpressRoute ▾

Redeem authorization ⓘ

*Virtual network gateway ⓘ

PrivateCloudGateway 🔒

Authorization key *

442cb5d8 ... ✓

Peer circuit URI *

/subscriptions/750a6f9e ... ✓

Subscription ⓘ

▾

Resource group ⓘ

ContosoResourceGroup 🔒

Create new

Location ⓘ

East US ▾

OK

The connection between your ExpressRoute circuit and your Virtual Network is created.

Microsoft Azure

Search resources, services, and docs (G+)

Connie Wilson
CONTOSO

Dashboard > Resource groups > avs-ncus > er-gw-ncus

er-gw-ncus | Connections

Virtual network gateway

Search (Ctrl+ /)

+ Add Refresh

Search connections

Name	Status	Connection type	Peer
avs-to-azure-ncus	Succeeded	ExpressRoute	tnt38-cust-p01-northc...

Overview

Activity log

Access control (IAM)

Tags

Diagnose and solve problems

Settings

Configuration

Connections

Properties

Locks

Next steps

In this tutorial, you learned how to:

- Create a Virtual Network using the vNet Connect Feature
- Create a Virtual Network Manually
- Create a Virtual Network gateway
- Connect your ExpressRoute circuit to the gateway

Continue to the next tutorial to learn how to create the NSX-T network segments used for VMs in vCenter Server.

[Create an NSX-T network segment](#)

Tutorial: Access an Azure VMware Solution private cloud

12/16/2022 • 2 minutes to read • [Edit Online](#)

Azure VMware Solution doesn't allow you to manage your private cloud with your on-premises vCenter Server. Instead, you'll need to connect to the Azure VMware Solution vCenter Server instance through a jump box.

In this tutorial, you'll create a jump box in the resource group you created in the [previous tutorial](#) and sign into the Azure VMware Solution vCenter Server. This jump box is a Windows virtual machine (VM) on the same virtual network you created. It provides access to both vCenter Server and the NSX Manager.

In this tutorial, you learn how to:

- Create a Windows VM to access the Azure VMware Solution vCenter
- Sign into vCenter Server from this VM

Create a new Windows virtual machine

1. In the resource group, select **Add**, search for **Microsoft Windows 10**, and select it. Then select **Create**.



Microsoft Windows 10
Microsoft

Microsoft Windows 10 [Save for later](#)

Select a software plan
Windows 10 Pro, Version 1909

Overview Plans

This software is provided by Microsoft. Use of this software in Microsoft Azure is not permitted except under a volume licensing agreement with Microsoft. By clicking Create, I acknowledge that I or the company I work for is licensed to use this software under a volume licensing agreement with Microsoft and that the right to use it will be subject to that agreement.

Learn more

[What's new for Business in the Windows 10 May 2019 Update](#)
[What's new for IT pros in the Windows 10 May 2019 Update](#)

2. Enter the required information in the fields, and then select **Review + create**.

For more information on the fields, see the following table.

FIELD	VALUE
Subscription	Value is pre-populated with the Subscription belonging to the Resource Group.
Resource group	Value is pre-populated for the current Resource Group, which you created in the preceding tutorial.
Virtual machine name	Enter a unique name for the VM.

FIELD	VALUE
Region	Select the geographical location of the VM.
Availability options	Leave the default value selected.
Image	Select the VM image.
Size	Leave the default size value.
Authentication type	Select Password .
Username	Enter the user name for logging on to the VM.
Password	Enter the password for logging on to the VM.
Confirm password	Enter the password for logging on to the VM.
Public inbound ports	<p>Select None.</p> <ul style="list-style-type: none"> To control access to the VM only when you want to access it, use JIT access. To securely access the jump box server from the internet without exposing any network port, use an Azure Bastion.

3. Once validation passes, select **Create** to start the virtual machine creation process.

Connect to the local vCenter of your private cloud

1. From the jump box, sign in to vSphere Client with VMware vCenter Server SSO using a cloud admin username and verify that the user interface displays successfully.
2. In the Azure portal, select your private cloud, and then **Manage > Identity**.

The URLs and user credentials for private cloud vCenter Server and NSX-T Manager display.

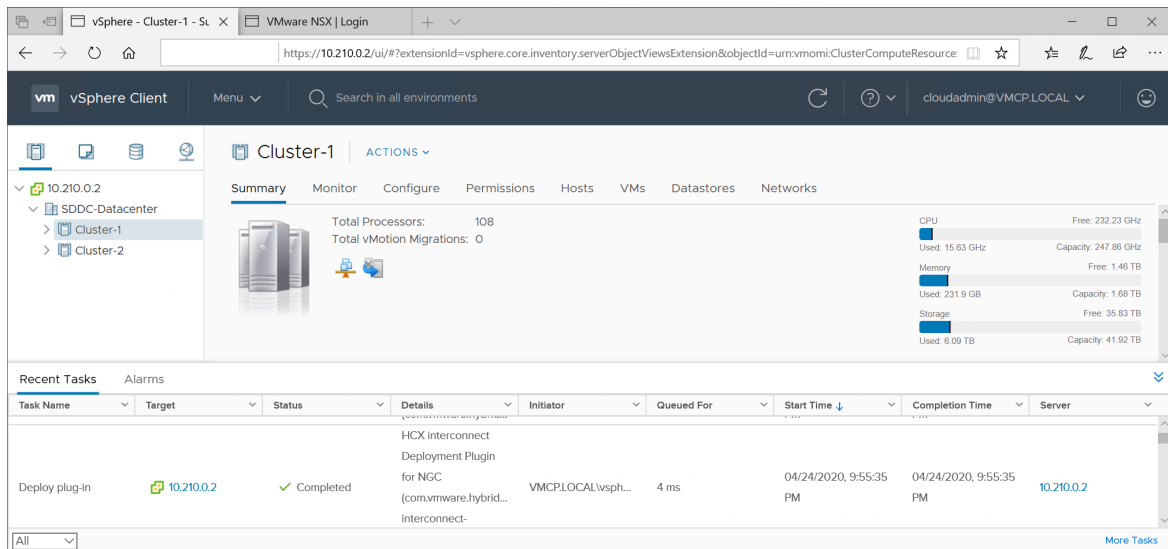
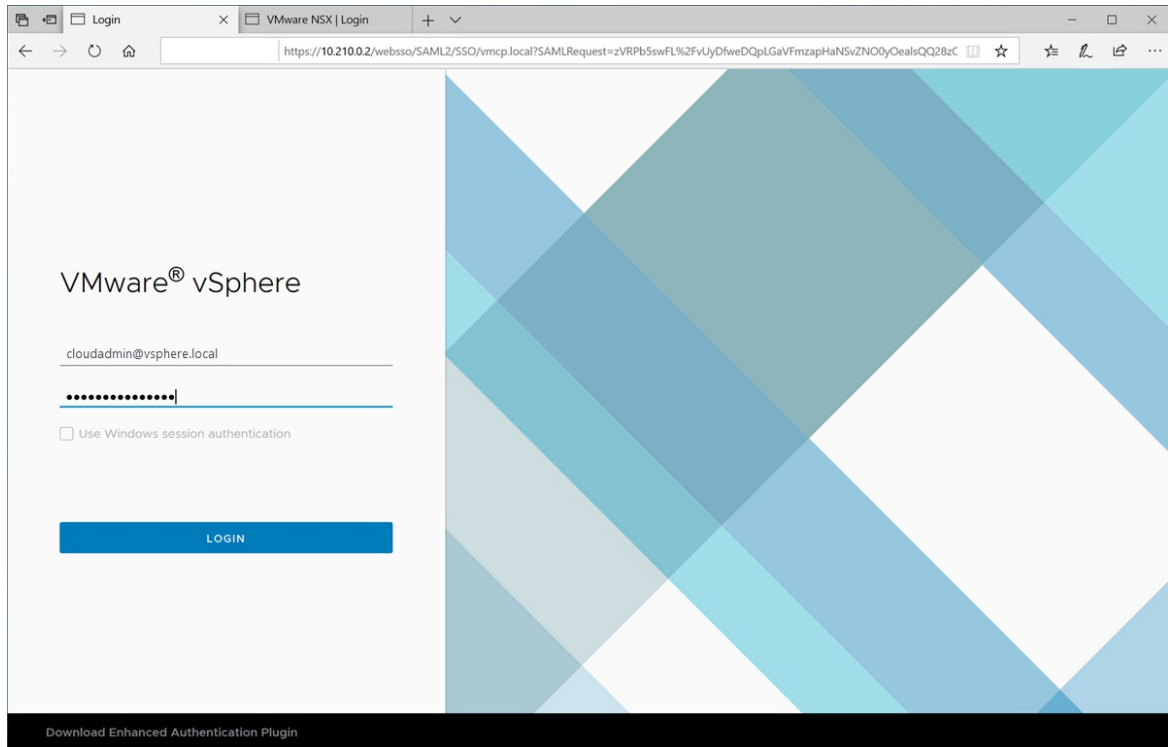
The screenshot shows the Azure portal interface for a private cloud named 'avs-pc-ncus'. The 'Identity' section is expanded, showing a list of login credentials. The 'vCenter credentials' section includes a Web client URL of 'https://10.1.0.2/', an Admin username of 'cloudadmin@vsphere.local', and an Admin password field. The 'NSX-T Manager credentials' section includes a Web client URL of 'https://10.1.0.3/', an Admin username of 'admin', and an Admin password field. Both sections also have a Certificate thumbprint field.

3. Navigate to the VM you created in the preceding step and connect to the virtual machine.

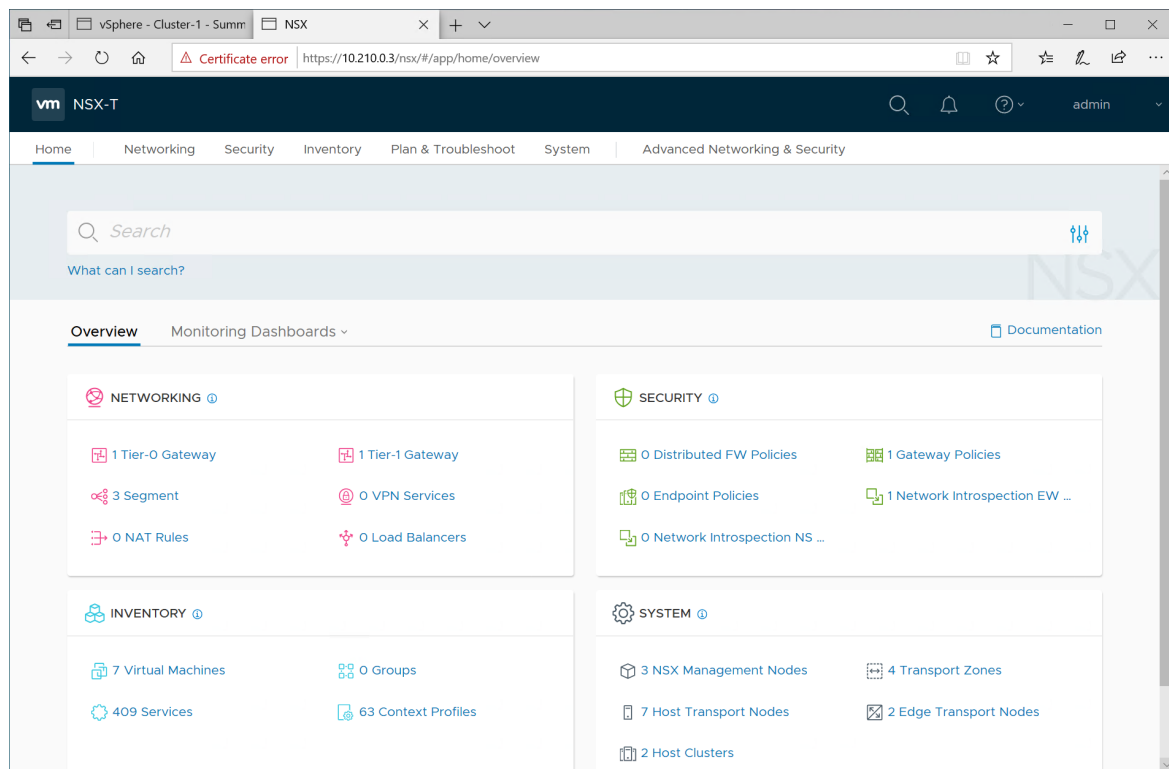
If you need help with connecting to the VM, see [connect to a virtual machine](#) for details.

4. In the Windows VM, open a browser and navigate to the vCenter Server and NSX-T Manager URLs in two tabs.

5. In the vSphere Client tab, enter the `ccloudadmin@vsphere.local` user credentials from the previous step.



6. In the second tab of the browser, sign in to NSX-T Manager.



Next steps

In this tutorial, you learned how to:

- Create a Windows VM to use to connect to vCenter Server
- Login to vCenter Server from your VM

Continue to the next tutorial to learn how to create a virtual network to set up local management for your private cloud clusters.

[Create a Virtual Network](#)

Tutorial: Add an NSX-T Data Center network segment in Azure VMware Solution

12/16/2022 • 3 minutes to read • [Edit Online](#)

After deploying Azure VMware Solution, you can configure an NSX-T Data Center network segment from NSX-T Manager or the Azure portal. Once configured, the segments are visible in Azure VMware Solution, NSX-T Manager, and vCenter Server. NSX-T Data Center comes pre-provisioned by default with an NSX-T Data Center Tier-0 gateway in **Active/Active** mode and a default NSX-T Data Center Tier-1 gateway in **Active/Standby** mode. These gateways let you connect the segments (logical switches) and provide East-West and North-South connectivity.

TIP

The Azure portal presents a simplified view of NSX-T Data Center operations a VMware administrator needs regularly and targeted at users not familiar with NSX-T Manager.

In this tutorial, you learn how to:

- Add network segments using either NSX-T Manager or the Azure portal
- Verify the new network segment

Prerequisites

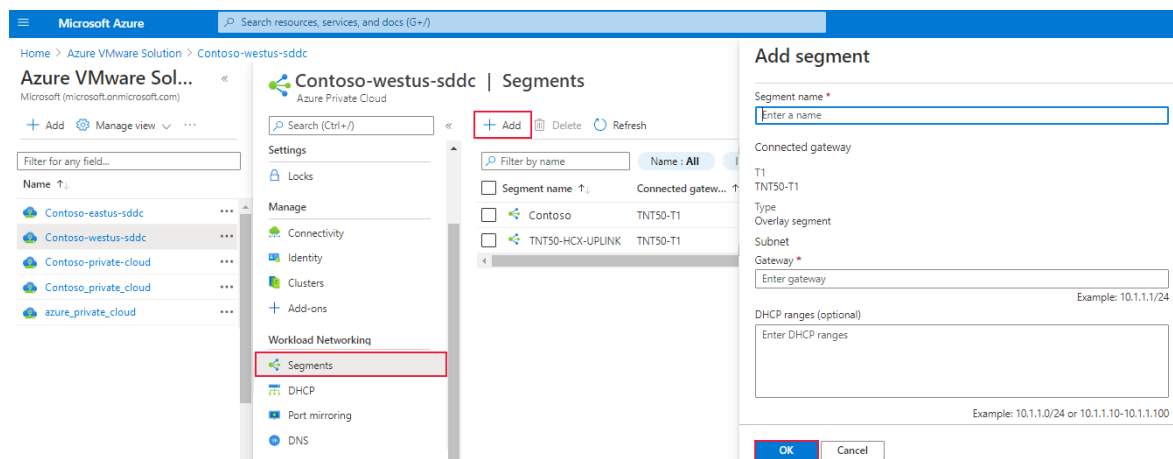
An Azure VMware Solution private cloud with access to the vCenter Server and NSX-T Manager interfaces. For more information, see the [Configure networking](#) tutorial.

Use Azure portal to add an NSX-T Data Center network segment

NOTE

If you plan to use DHCP, you'll need to [configure a DHCP server or DHCP relay](#) before you can configure an NSX-T network segment.

1. In your Azure VMware Solution private cloud, under **Workload Networking**, select **Segments** > **Add**.
2. Provide the details for the new logical segment and select **OK**.



- **Segment name** - Name of the segment that is visible in vCenter.
- **Subnet gateway** - Gateway IP address for the segment's subnet with a subnet mask. VMs are attached to a logical segment, and all VMs connecting to this segment belong to the same subnet. Also, all VMs attached to this logical segment must carry an IP address from the same segment.
- **DHCP (optional)** - DHCP ranges for a logical segment. You must configure a [DHCP server or DHCP relay](#) to consume DHCP on Segments.

NOTE

The **Connected gateway** is selected by default and is read-only. It shows Tier-1 gateway and type of segment information.

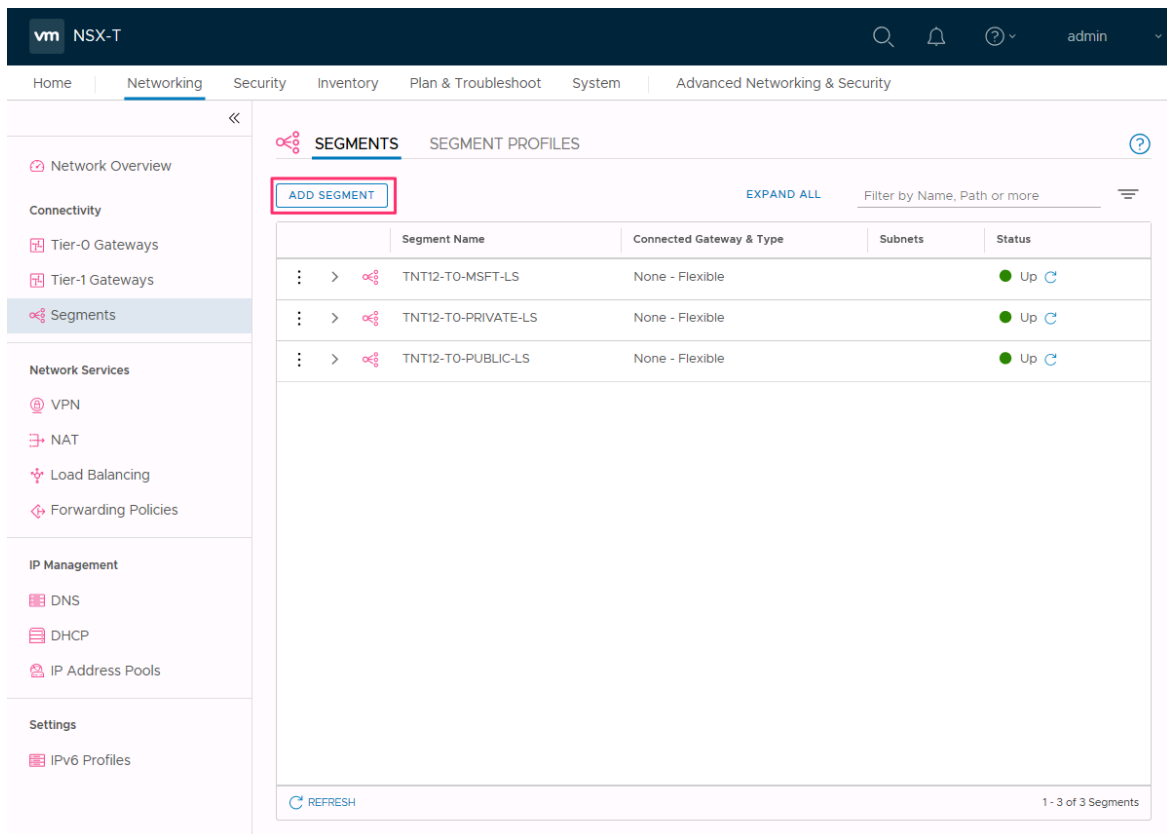
- **T1** - Name of the Tier-1 gateway in NSX-T Manager. A private cloud comes with an NSX-T Tier-0 gateway in Active/Active mode and a default NSX-T Tier-1 gateway in Active/Standby mode. Segments created through the Azure VMware Solution console only connect to the default Tier-1 gateway, and the workloads of these segments get East-West and North-South connectivity. You can only create more Tier-1 gateways through NSX-T Manager. Tier-1 gateways created from the NSX-T Manager console are not visible in the Azure VMware Solution console.
- **Type** - Overlay segment supported by Azure VMware Solution.

The segment is now visible in Azure VMware Solution, NSX-T Manger, and vCenter.

Use NSX-T Manager to add network segment

The virtual machines (VMs) created in vCenter Server are placed onto the network segments created in NSX-T Data Center and are visible in vCenter Server.

1. In NSX-T Manager, select **Networking > Segments**, and then select **Add Segment**.



2. Enter a name for the segment.

- Select the Tier-1 Gateway (TNTxx-T1) as the **Connected Gateway** and leave the **Type** as Flexible.
- Select the pre-configured overlay **Transport Zone** (TNTxx-OVERLAY-TZ) and then select **Set Subnets**.

The screenshot shows the NSX-T Segments configuration interface. The 'ADD SEGMENT' form is active, with the following fields highlighted in red:

- Segment Name: Is01
- Connected Gateway & Type: TNT12-T1
- Subnets: Flex
- Set Subnets: (checked)
- Transport Zone: TNT12-OVERLAY-TZ

Below the form, a table lists existing segments:

Segment Name	Connected Gateway & Type	Status
TNT12-TO-MSFT-LS	None - Flexible	Up
TNT12-TO-PRIVATE-LS	None - Flexible	Up
TNT12-TO-PUBLIC-LS	None - Flexible	Up

- Enter the gateway IP address and then select **Add**.

IMPORTANT

The IP address needs to be on a non-overlapping RFC1918 address block, which ensures connection to the VMs on the new segment.

The 'Set Subnets' dialog box is shown with the following details:

- Segment: #Subnets 1
- Gateway IP/Prefix Length: 10.10.230.1/24 (highlighted in red)
- Format CIDR e.g. 10.12.2.1/24
- DHCP Ranges: Enter DHCP Ranges (empty field)
- Formats, e.g. 10.12.2.64/26, 10.12.2.2-10.12.2.50
- Buttons: ADD (highlighted in red), CANCEL, APPLY

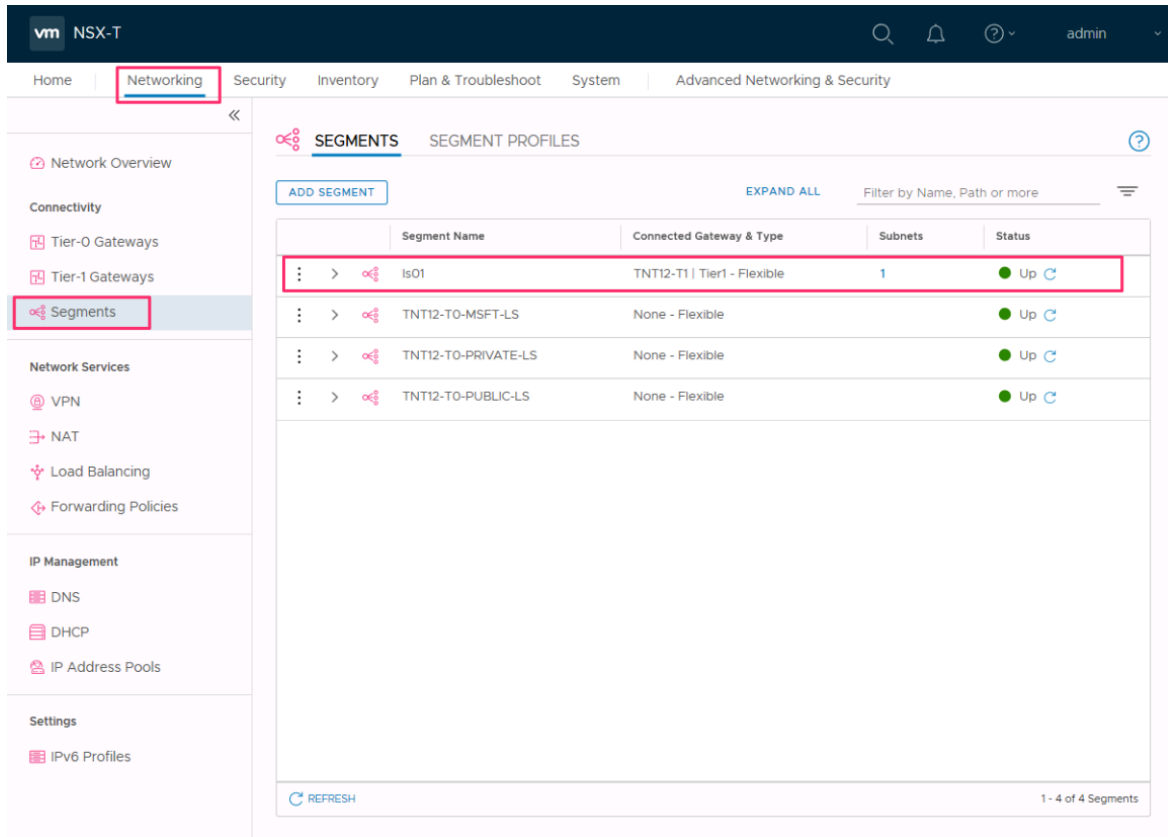
- Select **Apply** and then **Save**.

7. Select **No** to decline the option to continue configuring the segment.

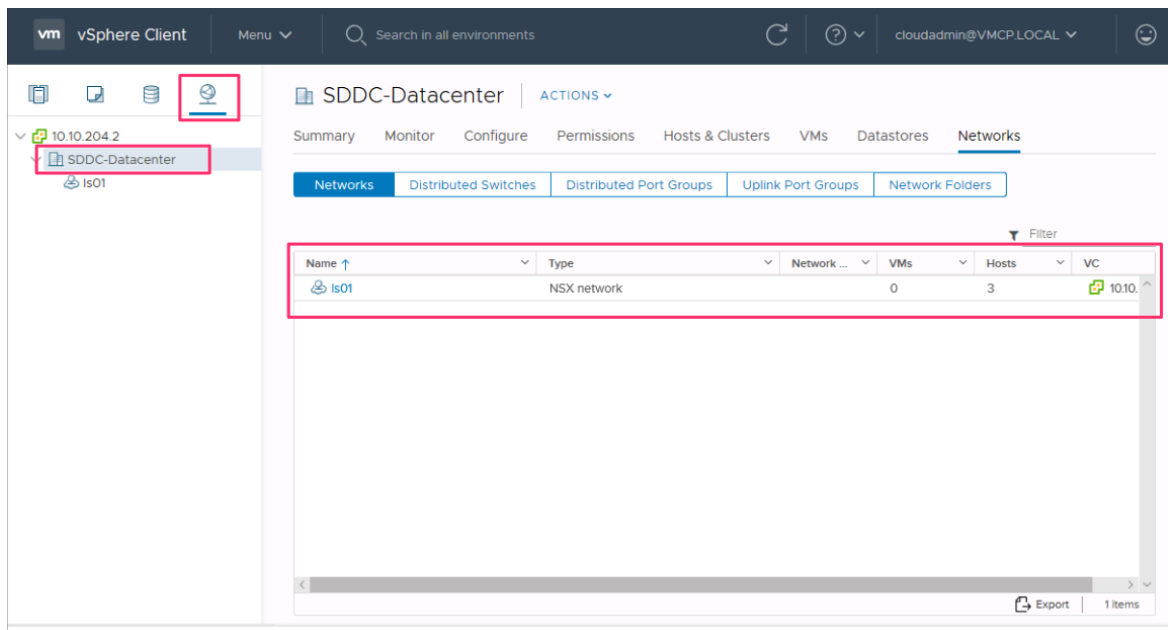
Verify the new network segment

Verify the presence of the new network segment. In this example, **Is01** is the new network segment.

1. In NSX-T Manager, select **Networking > Segments**.



2. In vCenter Server, select **Networking > SDDC-Datacenter**.



Next steps

In this tutorial, you created an NSX-T Data Center network segment to use for VMs in vCenter Server.


You can now:

- [Configure and manage DHCP for Azure VMware Solution](#)
- [Create a Content Library to deploy VMs in Azure VMware Solution](#)
- [Peer on-premises environments to a private cloud](#)

Tutorial: Peer on-premises environments to Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

After you deploy your Azure VMware Solution private cloud, you'll connect it to your on-premises environment. ExpressRoute Global Reach connects your on-premises environment to your Azure VMware Solution private cloud. The ExpressRoute Global Reach connection is established between the private cloud ExpressRoute circuit and an existing ExpressRoute connection to your on-premises environments.



NOTE

You can connect through VPN, but that's out of scope for this quick start guide.

In this article, you'll:

- Create an ExpressRoute auth key in the on-premises ExpressRoute circuit
- Peer the private cloud with your on-premises ExpressRoute circuit

- Verify on-premises network connectivity

After you're finished, follow the recommended next steps at the end to continue with the steps of this getting started guide.

Prerequisites

- Review the documentation on how to [enable connectivity in different Azure subscriptions](#).
- A separate, functioning ExpressRoute circuit for connecting on-premises environments to Azure, which is *circuit 1* for peering.
- Ensure that all gateways, including the ExpressRoute provider's service, supports 4-byte Autonomous System Number (ASN). Azure VMware Solution uses 4-byte public ASNs for advertising routes.

NOTE

If advertising a default route to Azure (0.0.0.0/0), ensure a more specific route containing your on-premises networks is advertised in addition to the default route to enable management access to Azure VMware Solution. A single 0.0.0.0/0 route will be discarded by Azure VMware Solution's management network to ensure successful operation of the service.

Create an ExpressRoute auth key in the on-premises ExpressRoute circuit

The circuit owner creates an authorization, which creates an authorization key to be used by a circuit user to connect their virtual network gateways to the ExpressRoute circuit. An authorization is valid for only one connection.

NOTE

Each connection requires a separate authorization.

1. From **ExpressRoute circuits** in the left navigation, under Settings, select **Authorizations**.
2. Enter the name for the authorization key and select **Save**.

Home > ExpressRoute circuits > Contoso-westus-sddc

Contoso-westus-sddc | Authorizations

Save Discard Refresh

Search (Ctrl+F)

Overview
Activity log
Access control (IAM)
Tags
Diagnose and solve problems

Settings

Configuration
Connections
Authorizations
Peerings
Properties
Locks

You can create authorizations that can be redeemed by other circuit users. Circuit users are owners of virtual network gateways (that are not within the same subscription as the ExpressRoute circuit). Each authorization can be redeemed with one virtual network.

To redeem authorizations, circuit users will need the resource ID of the ExpressRoute and an unused authorization key.
[Learn more](#)

Resource ID
/subscriptions/ ... /resourceGroups/ ...

Name	Provisioning state	Use status	Authorization key
js-er-avs-auth	Succeeded	Used	b2cbd ...
js-er-az-auth	Succeeded	Available	36d91 ...

Enter name

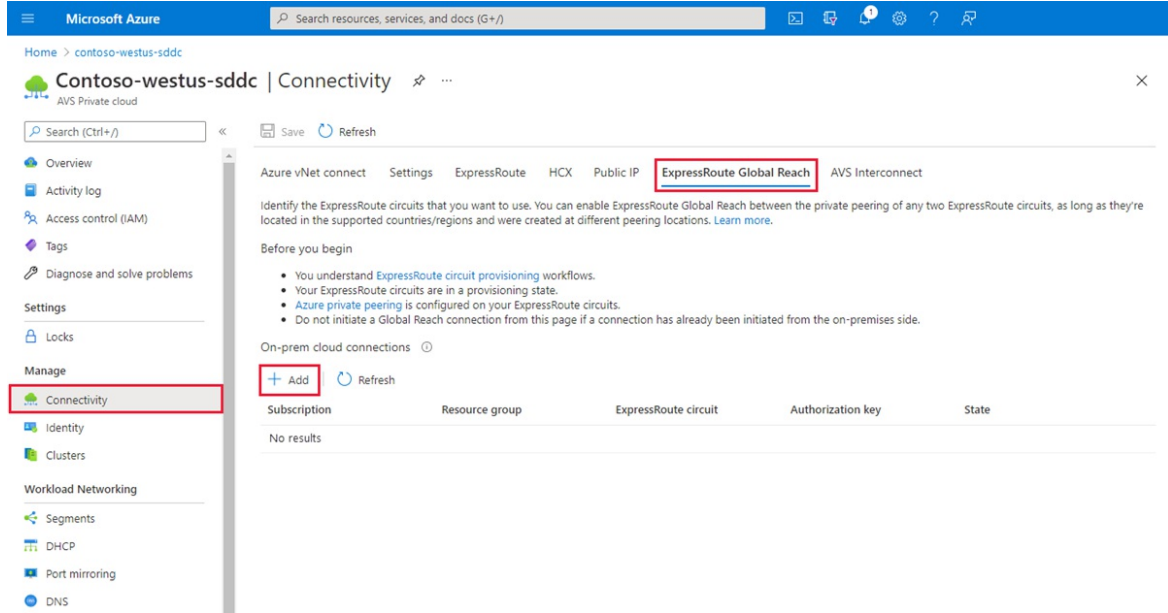
Once created, the new key appears in the list of authorization keys for the circuit.

3. Copy the authorization key and the ExpressRoute ID. You'll use them in the next step to complete the peering.

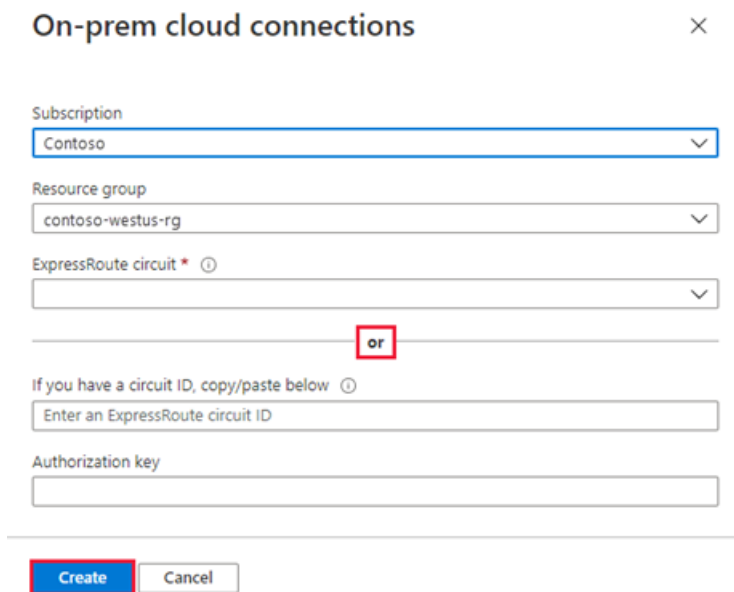
Peer private cloud to on-premises

Now that you've created an authorization key for the private cloud ExpressRoute circuit, you can peer it with your on-premises ExpressRoute circuit. The peering is done from the on-premises ExpressRoute circuit in the **Azure portal**. You'll use the resource ID (ExpressRoute circuit ID) and authorization key of your private cloud ExpressRoute circuit to finish the peering.

1. From the private cloud, under Manage, select **Connectivity > ExpressRoute Global Reach > Add**.



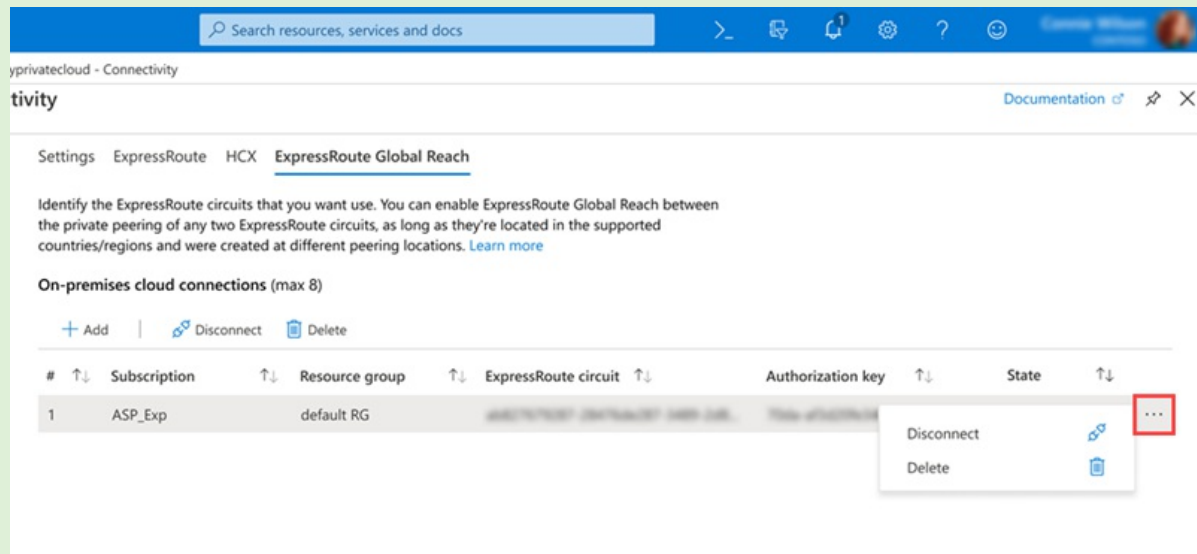
2. Enter the ExpressRoute ID and the authorization key created in the previous section.



3. Select **Create**. The new connection shows in the on-premises cloud connections list.

TIP

You can delete or disconnect a connection from the list by selecting **More**.



Search resources, services and docs

ypivatecloud - Connectivity

tivity Documentation

Settings ExpressRoute HCX **ExpressRoute Global Reach**

Identify the ExpressRoute circuits that you want use. You can enable ExpressRoute Global Reach between the private peering of any two ExpressRoute circuits, as long as they're located in the supported countries/regions and were created at different peering locations. [Learn more](#)

On-premises cloud connections (max 8)

+ Add | Disconnect Delete

#	Subscription	Resource group	ExpressRoute circuit	Authorization key	State
1	ASP_Exp	default RG			

Disconnect
Delete

Verify on-premises network connectivity

In your **on-premises edge router**, you should now see where the ExpressRoute connects the NSX-T network segments and the Azure VMware Solution management segments.

IMPORTANT

Everyone has a different environment, and some will need to allow these routes to propagate back into the on-premises network.

Next steps

Continue to the next tutorial to install VMware HCX add-on in your Azure VMware Solution private cloud.

[Install VMware HCX](#)

Tutorial: Scale clusters in a private cloud

12/16/2022 • 2 minutes to read • [Edit Online](#)

To get the most out of your Azure VMware Solution private cloud experience, scale the clusters and hosts to reflect what you need for planned workloads. You can scale the clusters and hosts in a private cloud as required for your application workload. You should address performance and availability limitations for specific services on a case-by-case basis.

The following table describes the maximum limits for Azure VMware Solution.

RESOURCE	LIMIT
vSphere clusters per private cloud	12
Minimum number of ESXi hosts per cluster	3 (hard-limit)
Maximum number of ESXi hosts per cluster	16 (hard-limit)
Maximum number of ESXi hosts per private cloud	96
Maximum number of vCenter Servers per private cloud	1 (hard-limit)
Maximum number of HCX site pairings	25 (any edition)
Maximum number of HCX service meshes	10 (any edition)
Maximum number of Azure VMware Solution ExpressRoute max linked private clouds	4 The virtual network gateway used determines the actual max linked private clouds. For more details, see About ExpressRoute virtual network gateways
Maximum Azure VMware Solution ExpressRoute port speed	10 Gbps The virtual network gateway used determines the actual bandwidth. For more details, see About ExpressRoute virtual network gateways
Maximum number of Azure Public IPv4 addresses assigned to NSX-T Data Center	2,000
Maximum number of Azure VMware Solution Interconnects per private cloud	10
vSAN capacity limits	75% of total usable (keep 25% available for SLA)
VMware Site Recovery Manager - Maximum number of protected Virtual Machines	3,000
VMware Site Recovery Manager - Maximum number of Virtual Machines per recovery plan	2,000

RESOURCE	LIMIT
VMware Site Recovery Manager - Maximum number of protection groups per recovery plan	250
VMware Site Recovery Manager - RPO Values	5 min or higher * (hard-limit)
VMware Site Recovery Manager - Maximum number of virtual machines per protection group	500
VMware Site Recovery Manager - Maximum number of recovery plans	250

* For information about Recovery Point Objective (RPO) lower than 15 minutes, see [How the 5 Minute Recovery Point Objective Works](#) in the *vSphere Replication Administration guide*.

For other VMware-specific limits, use the [VMware configuration maximum tool](#).

In this tutorial, you'll use the Azure portal to:

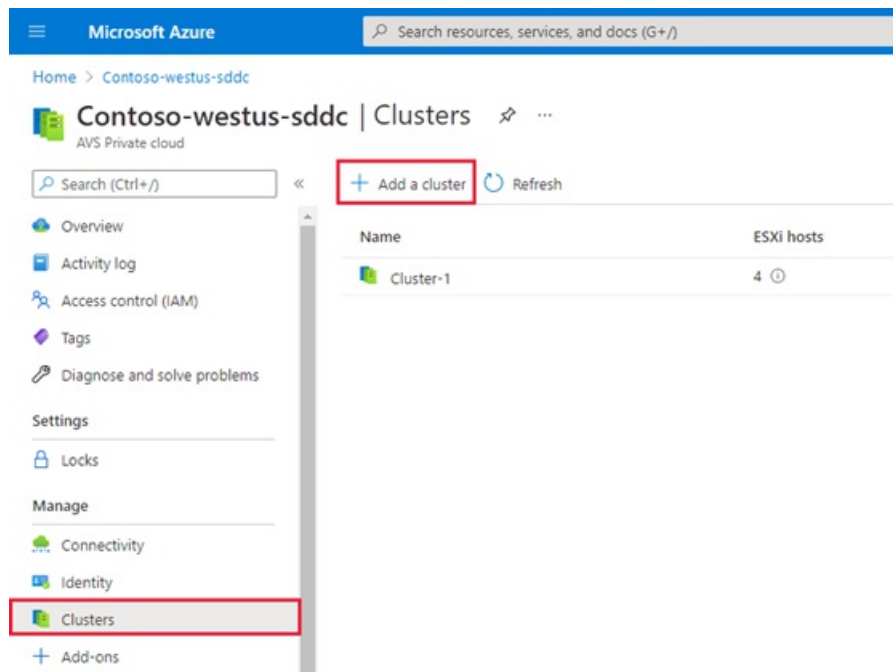
- Add a cluster to an existing private cloud
- Add hosts to an existing cluster

Prerequisites

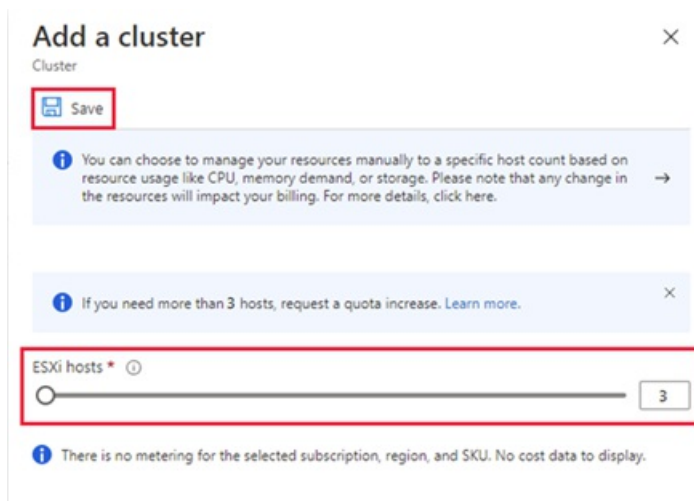
You'll need an existing private cloud to complete this tutorial. If you haven't created a private cloud, follow the [create a private cloud tutorial](#) to create one.

Add a new cluster

1. In your Azure VMware Solution private cloud, under **Manage**, select **Clusters** > **Add a cluster**.



2. Use the slider to select the number of hosts and then select **Save**.



The deployment of the new cluster begins.

Scale a cluster

1. In your Azure VMware Solution private cloud, under **Manage**, select **Clusters**.
2. Select the cluster you want to scale, select **More (...)**, then select **Edit**.



3. Use the slider to select the number of hosts and then select **Save**.

The addition of hosts to the cluster begins.

Next steps

If you require another Azure VMware Solution private cloud, [create another private cloud](#) following the same networking prerequisites, cluster, and host limits.

Tutorial: Delete an Azure VMware Solution private cloud

12/16/2022 • 2 minutes to read • [Edit Online](#)

If you have an Azure VMware Solution private cloud that you no longer need, you can delete it. The private cloud includes:

- An isolated network domain
- One or more provisioned vSphere clusters on dedicated server hosts
- Several virtual machines (VMs)

When you delete a private cloud, all VMs, their data, clusters, and network address space provisioned get deleted. The dedicated Azure VMware Solution hosts are securely wiped and returned to the free pool.

Caution

Deleting the private cloud terminates all running workloads and components and is an irreversible operation. Once you delete the private cloud, you cannot recover the data.

Prerequisites

If you require the VMs and their data later, make sure to backup the data before you delete the private cloud. Unfortunately, there's no way to recover the VMs and their data.

Delete the private cloud

1. Access the Azure VMware Solutions console in the [Azure portal](#).
2. Select the private cloud you want to delete.
3. Enter the name of the private cloud and select **Yes**.

NOTE

The deletion process takes a few hours to complete.

Azure VMware Solution identity concepts

12/16/2022 • 12 minutes to read • [Edit Online](#)

Azure VMware Solution private clouds are provisioned with a vCenter Server and NSX-T Manager. You'll use vCenter to manage virtual machine (VM) workloads and NSX-T Manager to manage and extend the private cloud. The CloudAdmin role is used for vCenter Server and the CloudAdmin role (with restricted permissions) is used for NSX-T Manager.

vCenter Server access and identity

In Azure VMware Solution, vCenter Server has a built-in local user called *cloudadmin* assigned to the CloudAdmin role. You can configure users and groups in Active Directory (AD) with the CloudAdmin role for your private cloud. In general, the CloudAdmin role creates and manages workloads in your private cloud. But in Azure VMware Solution, the CloudAdmin role has vCenter Server privileges that differ from other VMware cloud solutions and on-premises deployments.

IMPORTANT

The local *cloudadmin* user should be treated as an emergency access account for "break glass" scenarios in your private cloud. It's not for daily administrative activities or integration with other services.

- In a vCenter Server and ESXi on-premises deployment, the administrator has access to the vCenter Server administrator@vsphere.local account and the ESXi root account. They can also have more AD users and groups assigned.
- In an Azure VMware Solution deployment, the administrator doesn't have access to the administrator user account or the ESXi root account. They can, however, assign AD users and groups to the CloudAdmin role in vCenter Server. The CloudAdmin role doesn't have permissions to add an identity source like on-premises LDAP or LDAPS server to vCenter Server. However, you can use Run commands to add an identity source and assign cloudadmin role to users and groups.

The private cloud user doesn't have access to and can't configure specific management components Microsoft supports and manages. For example, clusters, hosts, datastores, and distributed virtual switches.

NOTE

In Azure VMware Solution, the *vsphere.local/SSO* domain is provided as a managed resource to support platform operations. It doesn't support the creation and management of local groups and users except for those provided by default with your private cloud.

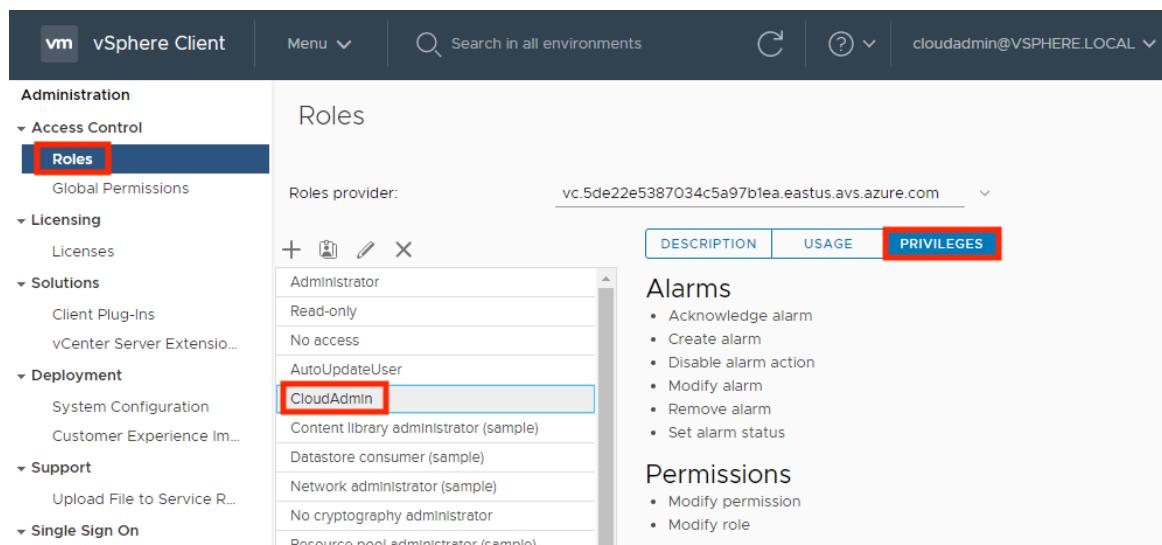
IMPORTANT

Azure VMware Solution offers custom roles on vCenter Server but currently doesn't offer them on the Azure VMware Solution portal. For more information, see the [Create custom roles on vCenter Server](#) section later in this article.

View the vCenter Server privileges

You can view the privileges granted to the Azure VMware Solution CloudAdmin role on your Azure VMware Solution private cloud vCenter Server.

1. Sign into the vSphere Client and go to **Menu > Administration**.
2. Under **Access Control**, select **Roles**.
3. From the list of roles, select **CloudAdmin** and then select **Privileges**.



The CloudAdmin role in Azure VMware Solution has the following privileges on vCenter Server. For more information, see the [VMware product documentation](#).

PRIVILEGE	DESCRIPTION
Alarms	<ul style="list-style-type: none"> Acknowledge alarm Create alarm Disable alarm action Modify alarm Remove alarm Set alarm status

PRIVILEGE	DESCRIPTION
Content Library	Add library item Add root certificate to trust store Check in a template Check out a template Create a subscription for a published library Create local library Create or delete a Harbor registry Create subscribed library Create, delete or purge a Harbor registry project Delete library item Delete local library Delete root certificate from trust store Delete subscribed library Delete subscription of a published library Download files Evict library items Evict subscribed library Import storage Manage Harbor registry resources on specified compute resource Probe subscription information Publish a library item to its subscribers Publish a library to its subscribers Read storage Sync library item Sync subscribed library Type introspection Update configuration settings Update files Update library Update library item Update local library Update subscribed library Update subscription of a published library View configuration settings
Cryptographic operations	Direct access
Datastore	Allocate space Browse datastore Configure datastore Low-level file operations Remove files Update virtual machine metadata
Folder	Create folder Delete folder Move folder Rename folder
Global	Cancel task Global tag Health Log event Manage custom attributes Service managers Set custom attribute System tag

PRIVILEGE	DESCRIPTION
Host	vSphere Replication Manage replication
Network	Assign network
Permissions	Modify permissions Modify role
Profile Driven Storage	Profile driven storage view
Resource	Apply recommendation Assign vApp to resource pool Assign virtual machine to resource pool Create resource pool Migrate powered off virtual machine Migrate powered on virtual machine Modify resource pool Move resource pool Query vMotion Remove resource pool Rename resource pool
Scheduled task	Create task Modify task Remove task Run task
Sessions	Message Validate session
Storage view	View
vApp	Add virtual machine Assign resource pool Assign vApp Clone Create Delete Export Import Move Power off Power on Rename Suspend Unregister View OVF environment vApp application configuration vApp instance configuration vApp managedBy configuration vApp resource configuration
Virtual machine	Change Configuration Acquire disk lease Add existing disk Add new disk Add or remove device Advanced configuration

PRIVILEGE	DESCRIPTION
	<ul style="list-style-type: none"> Change CPU count Change memory Change settings Change swapfile placement Change resource Configure host USB device Configure raw device Configure managedBy Display connection settings Extend virtual disk Modify device settings Query fault tolerance compatibility Query unowned files Reload from paths Remove disk Rename Reset guest information Set annotation Toggle disk change tracking Toggle fork parent Upgrade virtual machine compatibility Edit inventory <ul style="list-style-type: none"> Create from existing Create new Move Register Remove Unregister Guest operations <ul style="list-style-type: none"> Guest operation alias modification Guest operation alias query Guest operation modifications Guest operation program execution Guest operation queries Interaction <ul style="list-style-type: none"> Answer question Back up operation on virtual machine Configure CD media Configure floppy media Connect devices Console interaction Create screenshot Defragment all disks Drag and drop Guest operating system management by VIX API Inject USB HID scan codes Install VMware tools Pause or Unpause Wipe or shrink operations Power off Power on Record session on virtual machine Replay session on virtual machine Reset Resume Fault Tolerance Suspend Suspend fault tolerance Test failover Test restart secondary VM Turn off fault tolerance Turn on fault tolerance Provisioning <ul style="list-style-type: none"> Allow disk access Allow file access Allow read-only disk access

PRIVILEGE	DESCRIPTION
	Allow virtual machine download Clone template Clone virtual machine Create template from virtual machine Customize guest Deploy template Mark as template Modify customization specification Promote disks Read customization specifications Service configuration Allow notifications Allow polling of global event notifications Manage service configuration Modify service configuration Query service configurations Read service configuration Snapshot management Create snapshot Remove snapshot Rename snapshot Revert snapshot vSphere Replication Configure replication Manage replication Monitor replication
vService	Create dependency Destroy dependency Reconfigure dependency configuration Update dependency
vSphere tagging	Assign and unassign vSphere tag Create vSphere tag Create vSphere tag category Delete vSphere tag Delete vSphere tag category Edit vSphere tag Edit vSphere tag category Modify UsedBy field for category Modify UsedBy field for tag

Create custom roles on vCenter Server

Azure VMware Solution supports the use of custom roles with equal or lesser privileges than the CloudAdmin role. You'll use the CloudAdmin role to create, modify, or delete custom roles with privileges less than or equal to their current role.

NOTE

You can create roles with privileges greater than CloudAdmin. However, you can't assign the role to any users or groups or delete the role. Roles that have privileges greater than that of CloudAdmin is unsupported.

To prevent creating roles that can't be assigned or deleted, clone the CloudAdmin role as the basis for creating new custom roles.

Create a custom role

1. Sign in to vCenter Server with `cloudadmin@vsphere.local` or a user with the CloudAdmin role.
2. Navigate to the **Roles** configuration section and select **Menu > Administration > Access Control > Roles**.

3. Select the **CloudAdmin** role and select the **Clone role action** icon.

NOTE

Don't clone the **Administrator** role because you can't use it. Also, the custom role created can't be deleted by cloudadmin@vsphere.local.

4. Provide the name you want for the cloned role.
5. Remove privileges for the role and select **OK**. The cloned role is visible in the **Roles** list.

Apply a custom role

1. Navigate to the object that requires the added permission. For example, to apply permission to a folder, navigate to **Menu > VMs and Templates > Folder Name**.
2. Right-click the object and select **Add Permission**.
3. Select the Identity Source in the **User** drop-down where the group or user can be found.
4. Search for the user or group after selecting the Identity Source under the **User** section.
5. Select the role that you want to apply to the user or group.

NOTE

Attempting to apply a user or group to a role that has privileges greater than that of CloudAdmin will result in errors.

6. Check the **Propagate to children** if needed, and select **OK**. The added permission displays in the **Permissions** section.

VMware NSX-T Data Center NSX-T Manager access and identity

When a private cloud is provisioned using Azure portal, software-defined data center (SDDC) management components like vCenter Server and VMware NSX-T Data Center NSX-T Manager are provisioned for customers.

Microsoft is responsible for the lifecycle management of NSX-T appliances like, VMware NSX-T Data Center NSX-T Manager and VMware NSX-T Data Center Microsoft Edge appliances. They're responsible for bootstrapping network configuration, like creating the Tier-0 gateway.

You're responsible for VMware NSX-T Data Center software-defined networking (SDN) configuration, for example:

- Network segments
- Other Tier-1 gateways
- Distributed firewall rules
- Stateful services like gateway firewall
- Load balancer on Tier-1 gateways

You can access VMware NSX-T Data Center NSX-T Manager using the built-in local user "cloudadmin" assigned to a custom role that gives limited privileges to a user to manage VMware NSX-T Data Center. While Microsoft manages the lifecycle of VMware NSX-T Data Center, certain operations aren't allowed by a user. Operations not allowed include editing the configuration of host and edge transport nodes or starting an upgrade. For new users, Azure VMware Solution deploys them with a specific set of permissions needed by that user. The purpose is to provide a clear separation of control between the Azure VMware Solution control plane configuration and

Azure VMware Solution private cloud user.

For new private cloud deployments, VMware NSX-T Data Center access will be provided with a built-in local user cloudadmin assigned to the cloudadmin role with a specific set of permissions to use VMware NSX-T Data Center functionality for workloads.

VMware NSX-T Data Center cloudadmin user permissions

The following permissions are assigned to the cloudadmin user in Azure VMware Solution NSX-T Data Center.

NOTE

VMware NSX-T Data Center cloudadmin user on Azure VMware Solution is not the same as the cloudadmin user mentioned in the VMware product documentation.

CATEGORY	TYPE	OPERATION	PERMISSION
Networking	Connectivity	Tier-0 Gateways Tier-1 Gateways Segments	Read-only Full Access Full Access
Networking	Network Services	VPN NAT Load Balancing Forwarding Policy Statistics	Full Access Full Access Full Access Read-only Full Access
Networking	IP Management	DNS DHCP IP Address Pools	Full Access Full Access Full Access
Networking	Profiles		Full Access
Security	East West Security	Distributed Firewall Distributed IDS and IPS Identity Firewall	Full Access Full Access Full Access
Security	North South Security	Gateway Firewall URL Analysis	Full Access Full Access
Security	Network Introspection		Read-only
Security	Endpoint Protection		Read-only
Security	Settings		Full Access
Inventory			Full Access
Troubleshooting	IPFIX		Full Access
Troubleshooting	Port Mirroring		Full Access
Troubleshooting	Traceflow		Full Access

CATEGORY	TYPE	OPERATION	PERMISSION
System	Configuration Settings Settings Settings	Identity firewall Users and Roles Certificate Management (Service Certificate only) User Interface Settings	Full Access Full Access Full Access Full Access
System	All other		Read-only

You can view the permissions granted to the Azure VMware Solution cloudadmin role on your Azure VMware Solution private cloud VMware NSX-T Data Center.

1. Log in to the NSX-T Manager.
2. Navigate to **Systems** and locate **Users and Roles**.
3. Select and expand the **cloudadmin** role, found under **Roles**.
4. Select a category like, Networking or Security, to view the specific permissions.

NOTE

Private clouds created before June 2022 will switch from **admin** role to **cloudadmin** role. You'll receive a notification through Azure Service Health that includes the timeline of this change so you can change the NSX-T Data Center credentials you've used for other integration.

NSX-T Data Center LDAP integration for role-based access control (RBAC)

In an Azure VMware Solution deployment, the VMware NSX-T Data Center can be integrated with external LDAP directory service to add remote directory users or group, and assign them a VMware NSX-T Data Center RBAC role, like on-premises deployment. For more information on how to enable VMware NSX-T Data Center LDAP integration, see the [VMware product documentation](#).

Unlike on-premises deployment, not all pre-defined NSX-T Data Center RBAC roles are supported with Azure VMware solution to keep Azure VMware Solution IaaS control plane config management separate from tenant network and security configuration. See the next section, Supported NSX-T Data Center RBAC roles, for more details.

NOTE

VMware NSX-T Data Center LDAP Integration is supported only with SDDC's with VMware NSX-T Data Center "cloudadmin" user.

Supported and unsupported NSX-T Data Center RBAC roles

In an Azure VMware Solution deployment, the following VMware NSX-T Data Center predefined RBAC roles are supported with LDAP integration:

- Auditor
- Cloudadmin
- LB Admin
- LB Operator
- VPN Admin
- Network Operator

In an Azure VMware Solution deployment, the following VMware NSX-T Data Center predefined RBAC roles aren't supported with LDAP integration:

- Enterprise Admin
- Network Admin
- Security Admin
- Netx Partner Admin
- GI Partner Admin

You can create custom roles in NSX-T Data Center with permissions lesser than or equal to CloudAdmin role created by Microsoft. Following are examples on how to create a supported "Network Admin" and "Security Admin" role.

NOTE

Custom role creation will fail if you assign a permission not allowed by CloudAdmin role.

Create "AVS network admin" role

Use the following steps to create this custom role.

1. Navigate to **System > Users and Roles > Roles**.
2. Clone **Network Admin** and provide the name, **AVS Network Admin**.
3. **Modify** the following permissions to "Read Only" or "None" as seen in the **Permission** column in the following table.

CATEGORY	SUBCATEGORY	FEATURE	PERMISSION
Networking	Connectivity	Tier-0 Gateways	Read-only
	Network Services	Tier-0 Gateways > OSPF	None
		Forwarding Policy	None

4. **Apply** the changes and **Save** the Role.

Create "AVS security admin" role

Use the following steps to create this custom role.

1. Navigate to **System > Users and Roles > Roles**.
2. Clone **Security Admin** and provide the name, "AVS Security Admin".
3. **Modify** the following permissions to "Read Only" or "None" as seen in the **Permission** column in the following table.

CATEGORY	SUBCATEGORY	FEATURE	PERMISSION
Networking	Network Services	Forwarding Policy	None
Security	Network Introspection	Service profiles	None
	Endpoint Protection		None
	Settings		None

4. **Apply** the changes and **Save** the Role.

NOTE

The VMware NSX-T Data Center **System > Identity Firewall AD** configuration option isn't supported by the NSX custom role. The recommendation is to assign the **Security Operator** role to the user with the custom role to allow managing the Identity Firewall (IDFW) feature for that user.

NOTE

The VMware NSX-T Data Center Traceflow feature isn't supported by the VMware NSX-T Data Center custom role. The recommendation is to assign the **Auditor** role to the user along with above custom role to enable Traceflow feature for that user.

NOTE

VMware vRealize Automation(vRA) integration with the NSX-T Data Center component of the Azure VMware Solution requires the "auditor" role to be added to the user with the NSX-T Manager cloudadmin role.

Next steps

Now that you've covered Azure VMware Solution access and identity concepts, you may want to learn about:

- [How to configure external identity source for vCenter](#)
- [How to enable Azure VMware Solution resource](#)
- [Details of each privilege](#)
- [How Azure VMware Solution monitors and repairs private clouds](#)

Publish and protect APIs running on Azure VMware Solution VMs

12/16/2022 • 2 minutes to read • [Edit Online](#)

Microsoft Azure [API Management](#) lets you securely publish to external or internal consumers. Only the Developer (development) and Premium (production) SKUs allow Azure Virtual Network integration to publish APIs that run on Azure VMware Solution workloads. In addition, both SKUs enable the connectivity between the API Management service and the backend.

The API Management configuration is the same for backend services that run on Azure VMware Solution virtual machines (VMs) and on-premises. API Management also configures the virtual IP on the load balancer as the backend endpoint for both deployments when the backend server is placed behind an NSX Load Balancer on Azure VMware Solution.

External deployment

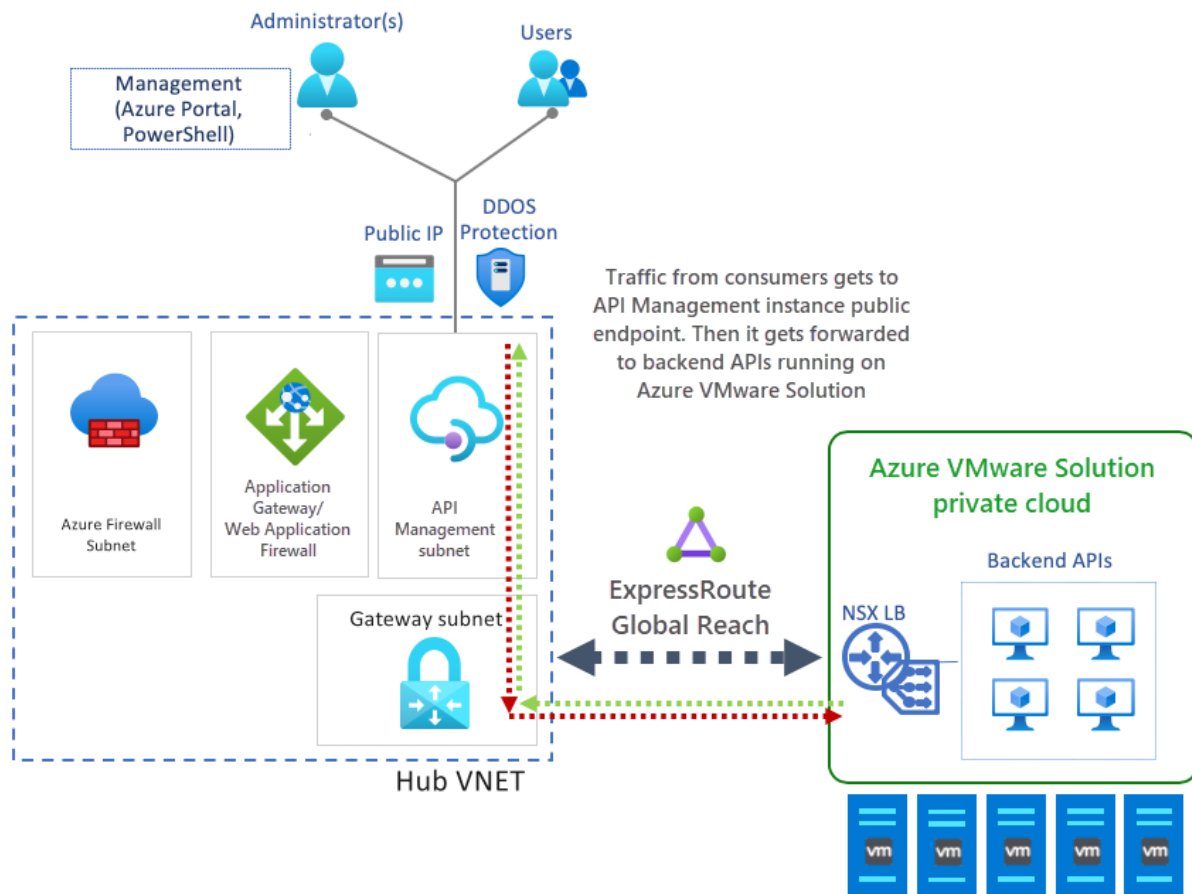
An external deployment publishes APIs consumed by external users that use a public endpoint. Developers and DevOps engineers can manage APIs through the Azure portal or PowerShell and the API Management developer portal.

The external deployment diagram shows the entire process and the actors involved (shown at the top). The actors are:

- **Administrator(s):** Represents the admin or DevOps team, which manages Azure VMware Solution through the Azure portal and automation mechanisms like PowerShell or Azure DevOps.
- **Users:** Represents the exposed APIs' consumers and represents both users and services consuming the APIs.

The traffic flow goes through the API Management instance, which abstracts the backend services, plugged into the Hub virtual network. The ExpressRoute Gateway routes the traffic to the ExpressRoute Global Reach channel and reaches an NSX Load Balancer distributing the incoming traffic to the different backend service instances.

API Management has an Azure Public API, and activating Azure DDoS Protection Service is recommended.



Internal deployment

An internal deployment publishes APIs consumed by internal users or systems. DevOps teams and API developers use the same management tools and developer portal as in the external deployment.

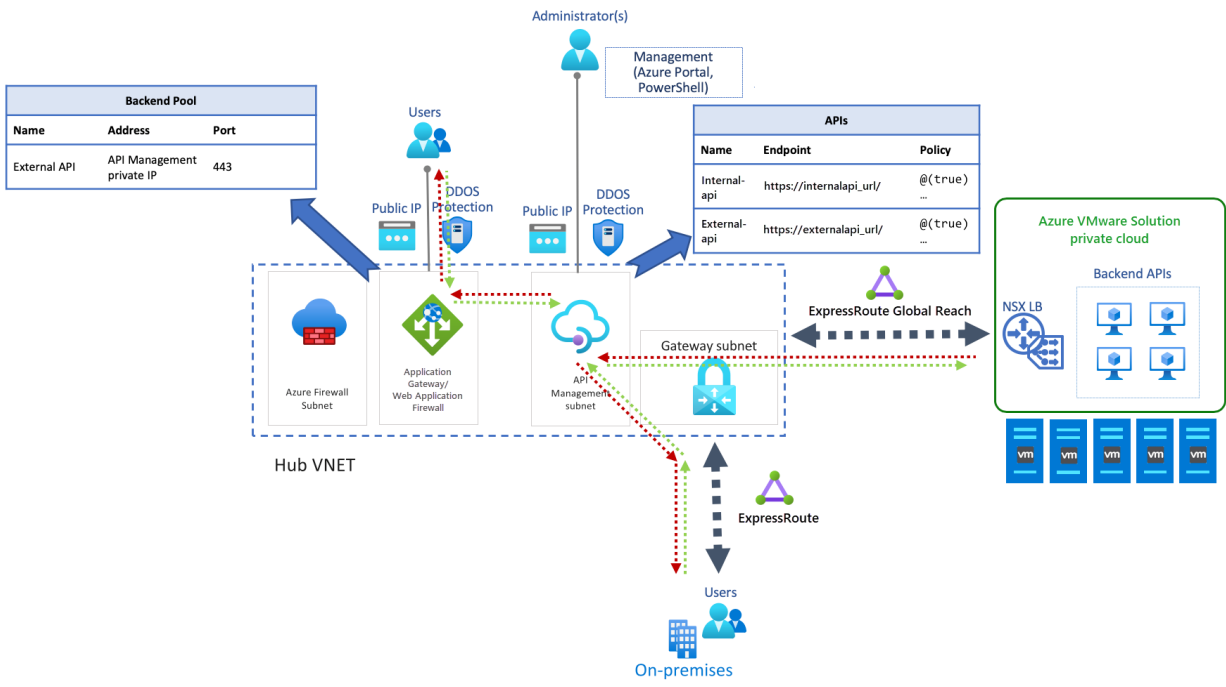
Use [Azure Application Gateway](#) for internal deployments to create a public and secure endpoint for the API. The gateway's capabilities are used to create a hybrid deployment that enables different scenarios.

- Use the same API Management resource for consumption by both internal and external consumers.
- Have a single API Management resource with a subset of APIs defined and available for external consumers.
- Provide an easy way to switch access to API Management from the public internet on and off.


The deployment diagram below shows consumers that can be internal or external, with each type accessing the same or different APIs.

In an internal deployment, APIs get exposed to the same API Management instance. In front of API Management, Application Gateway gets deployed with Azure Web Application Firewall (WAF) capability activated. Also deployed, a set of HTTP listeners and rules to filter the traffic, exposing only a subset of the backend services running on Azure VMware Solution.

- Internal traffic routes through ExpressRoute Gateway to Azure Firewall and then to API Management, directly or through traffic rules.
- External traffic enters Azure through Application Gateway, which uses the external protection layer for API Management.



Hub VNET



On-premises

Integrate Azure VMware Solution in a hub and spoke architecture

12/16/2022 • 6 minutes to read • [Edit Online](#)

This article provides recommendations for integrating an Azure VMware Solution deployment in an existing or a new [Hub and Spoke architecture](#) on Azure.

The Hub and Spoke scenario assume a hybrid cloud environment with workloads on:

- Native Azure using IaaS or PaaS services
- Azure VMware Solution
- vSphere on-premises

Architecture

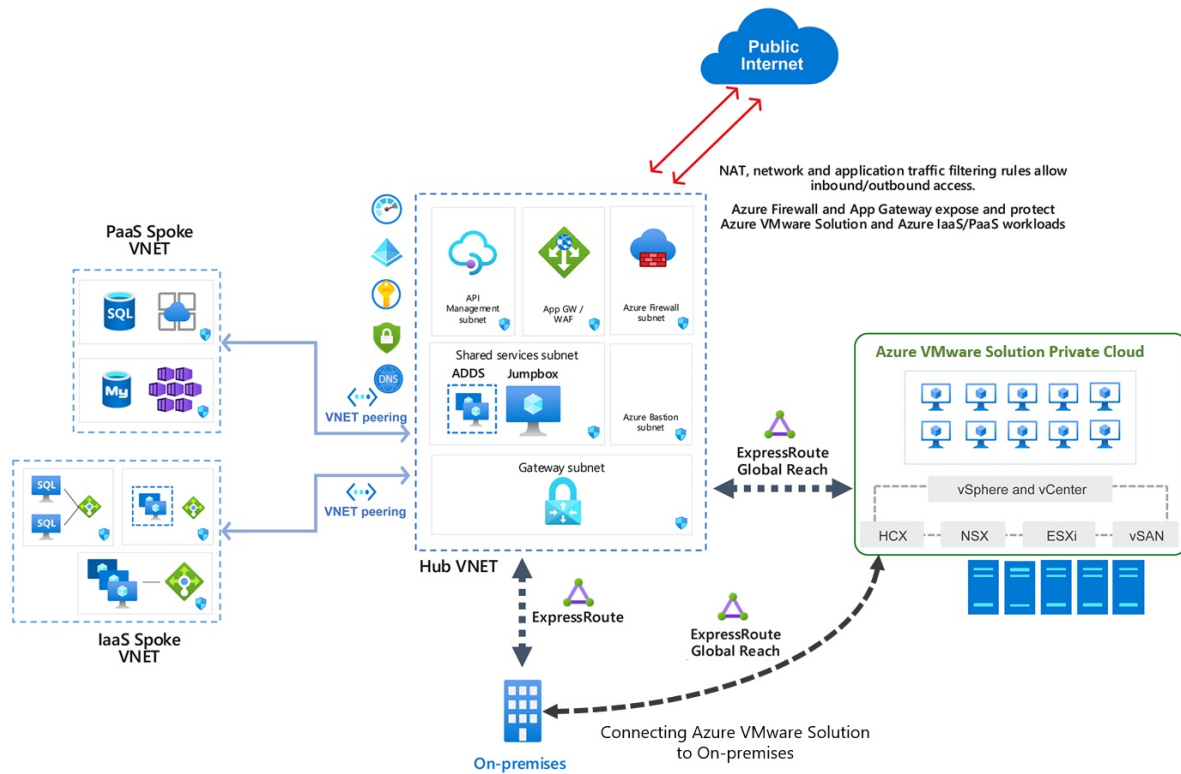
The *Hub* is an Azure Virtual Network that acts as a central point of connectivity to your on-premises and Azure VMware Solution private cloud. The *Spokes* are virtual networks peered with the Hub to enable cross-virtual network communication.

Traffic between the on-premises datacenter, Azure VMware Solution private cloud, and the Hub goes through Azure ExpressRoute connections. Spoke virtual networks usually contain IaaS based workloads but can have PaaS services like [App Service Environment](#), which has direct integration with Virtual Network, or other PaaS services with [Azure Private Link](#) enabled.

IMPORTANT

You can use an existing ExpressRoute Gateway to connect to Azure VMware Solution as long as it does not exceed the limit of four ExpressRoute circuits per virtual network. However, to access Azure VMware Solution from on-premises through ExpressRoute, you must have ExpressRoute Global Reach since the ExpressRoute gateway does not provide transitive routing between its connected circuits.

The diagram shows an example of a Hub and Spoke deployment in Azure connected to on-premises and Azure VMware Solution through ExpressRoute Global Reach.



The architecture has the following main components:

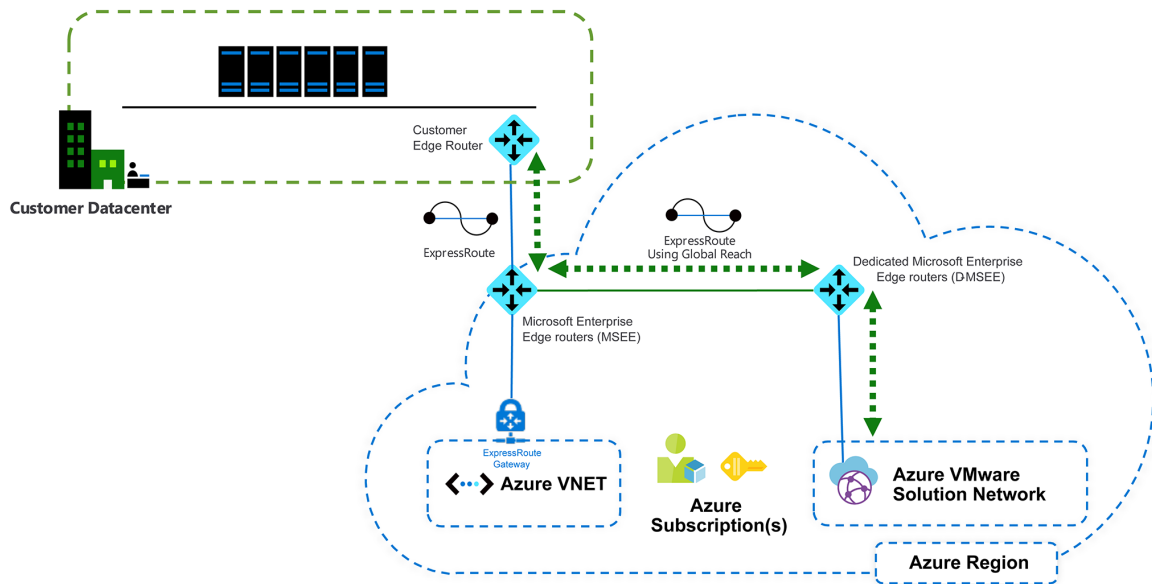
- **On-premises site:** Customer on-premises datacenter(s) connected to Azure through an ExpressRoute connection.
- **Azure VMware Solution private cloud:** Azure VMware Solution SDDC formed by one or more vSphere clusters, each one with a maximum of 16 hosts.
- **ExpressRoute gateway:** Enables the communication between Azure VMware Solution private cloud, shared services on Hub virtual network, and workloads running on Spoke virtual networks.
- **ExpressRoute Global Reach:** Enables the connectivity between on-premises and Azure VMware Solution private cloud. The connectivity between Azure VMware Solution and the Azure fabric is through ExpressRoute Global Reach only.
- **S2S VPN considerations:** Connectivity to Azure VMware Solution private cloud using Azure S2S VPN is supported as long as it meets the [minimum network requirements](#) for VMware HCX.
- **Hub virtual network:** Acts as the central point of connectivity to your on-premises network and Azure VMware Solution private cloud.
- **Spoke virtual network**
 - **IaaS Spoke:** Hosts Azure IaaS based workloads, including VM availability sets and Virtual Machine Scale Sets, and the corresponding network components.
 - **PaaS Spoke:** Hosts Azure PaaS services using private addressing thanks to [Private Endpoint](#) and [Private Link](#).
- **Azure Firewall:** Acts as the central piece to segment traffic between the Spokes and Azure VMware Solution.
- **Application Gateway:** Exposes and protects web apps that run either on Azure IaaS/PaaS or Azure VMware Solution virtual machines (VMs). It integrates with other services like API Management.

Network and security considerations

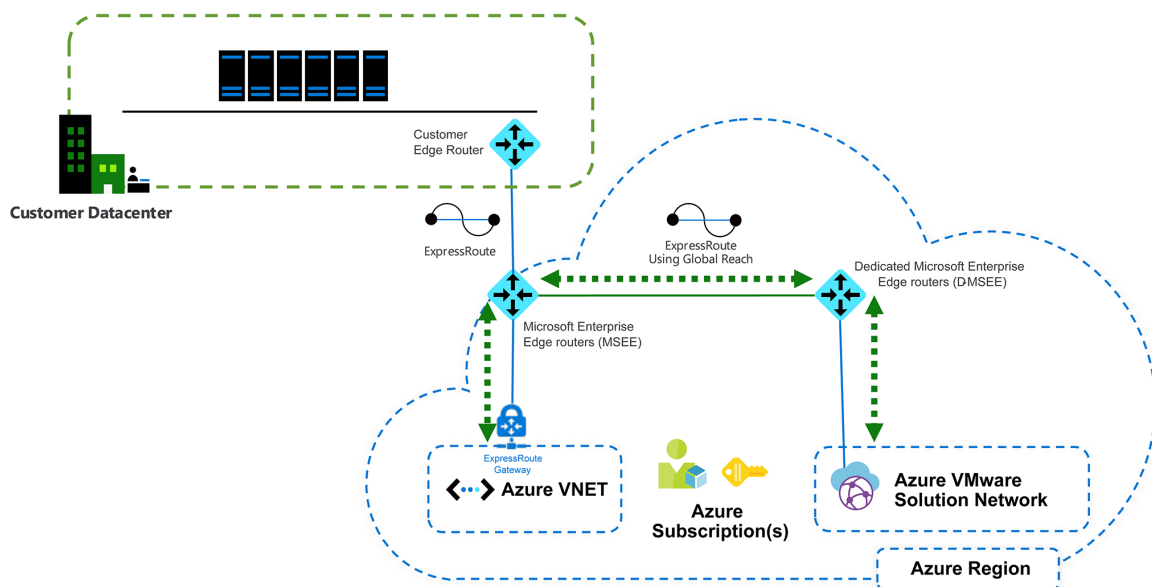
ExpressRoute connections enable traffic to flow between on-premises, Azure VMware Solution, and the Azure network fabric. Azure VMware Solution uses [ExpressRoute Global Reach](#) to implement this connectivity.

Because an ExpressRoute gateway doesn't provide transitive routing between its connected circuits, on-premises connectivity also must use ExpressRoute Global Reach to communicate between the on-premises vSphere environment and Azure VMware Solution.

- On-premises to Azure VMware Solution traffic flow



- Azure VMware Solution to Hub VNET traffic flow



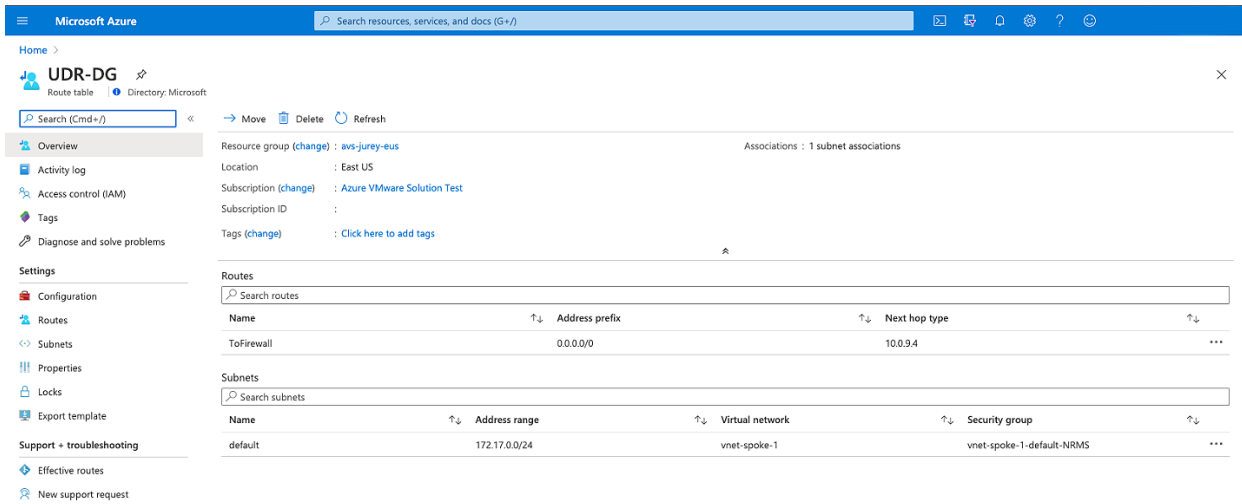
For more information on Azure VMware Solution networking and connectivity concepts, see the [Azure VMware Solution product documentation](#).

Traffic segmentation

[Azure Firewall](#) is the Hub and Spoke topology's central piece, deployed on the Hub virtual network. Use Azure Firewall, or another Azure supported network virtual appliance (NVA) to establish traffic rules and segment the communication between the different spokes and Azure VMware Solution workloads.

Create route tables to direct the traffic to Azure Firewall. For the Spoke virtual networks, create a route that sets the default route to the internal interface of the Azure Firewall. This way, when a workload in the Virtual

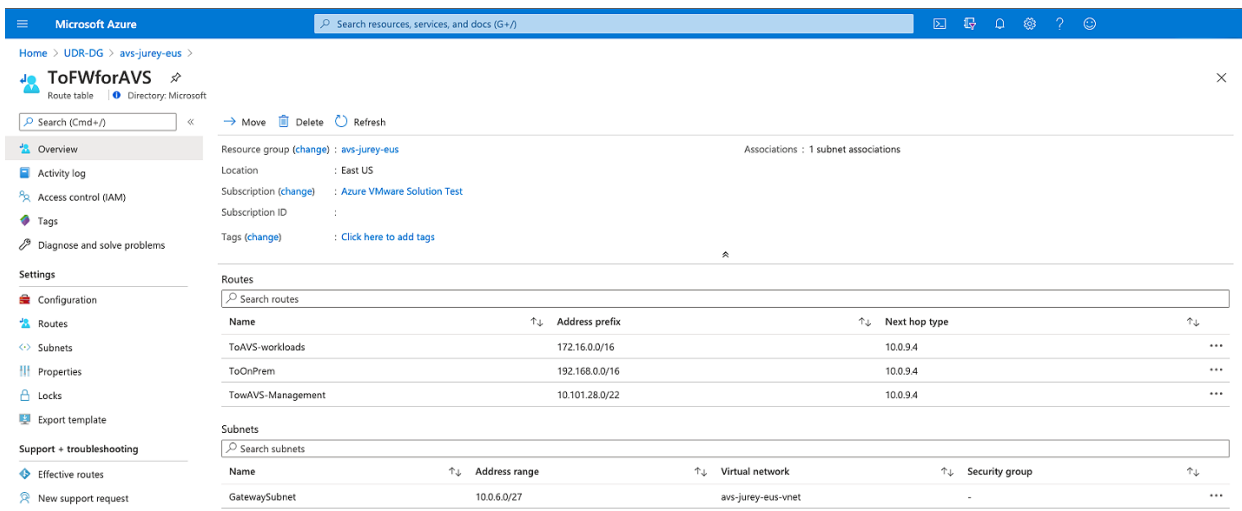
Network needs to reach the Azure VMware Solution address space, the firewall can evaluate it and apply the corresponding traffic rule to either allow or deny it.



IMPORTANT

A route with address prefix 0.0.0.0/0 on the **GatewaySubnet** setting is not supported.

Set routes for specific networks on the corresponding route table. For example, routes to reach Azure VMware Solution management and workloads IP prefixes from the spoke workloads and the other way around.



A second level of traffic segmentation using the network security groups within the Spokes and the Hub to create a more granular traffic policy.

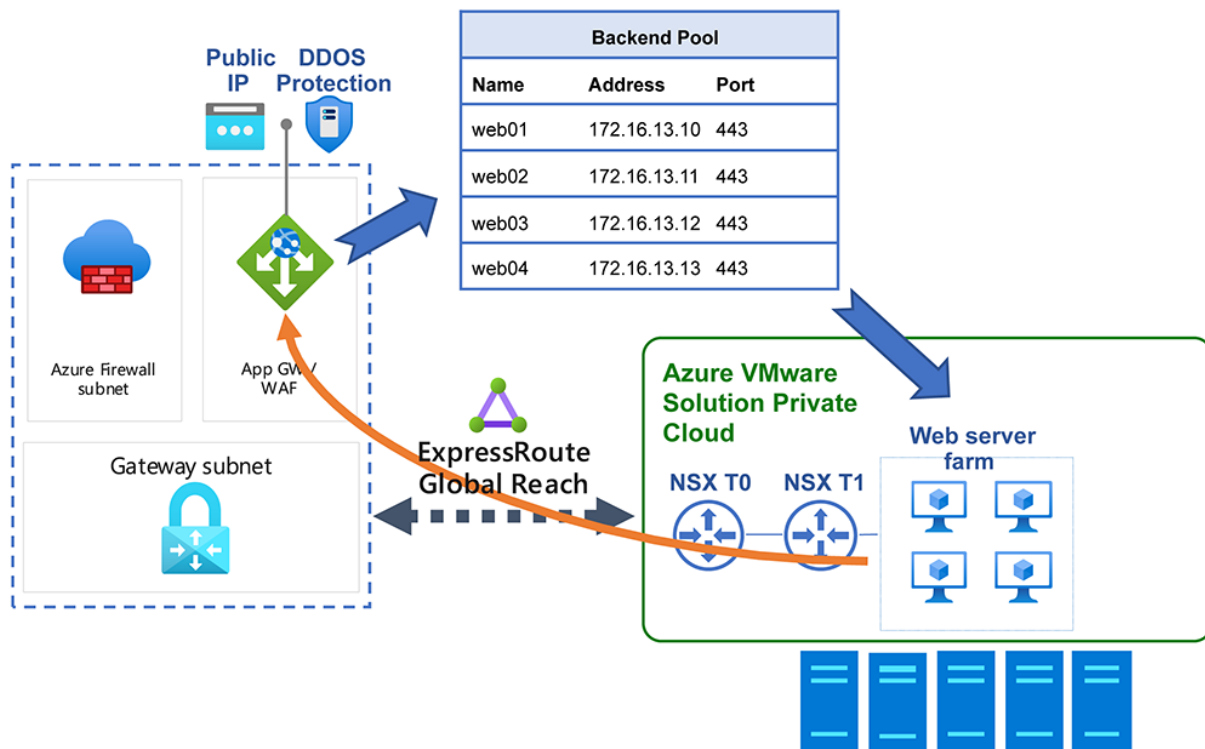
NOTE

Traffic from on-premises to Azure VMware Solution: Traffic between on-premises workloads, either vSphere-based or others, are enabled by Global Reach, but the traffic doesn't go through Azure Firewall on the hub. In this scenario, you must implement traffic segmentation mechanisms, either on-premises or in Azure VMware Solution.

Application Gateway

Azure Application Gateway V1 and V2 have been tested with web apps that run on Azure VMware Solution VMs as a backend pool. Application Gateway is currently the only supported method to expose web apps running on Azure VMware Solution VMs to the internet. It can also expose the apps to internal users securely.

For more information, see the Azure VMware Solution-specific article on [Application Gateway](#).



Jump box and Azure Bastion

Access Azure VMware Solution environment with a jump box, which is a Windows 10 or Windows Server VM deployed in the shared service subnet within the Hub virtual network.

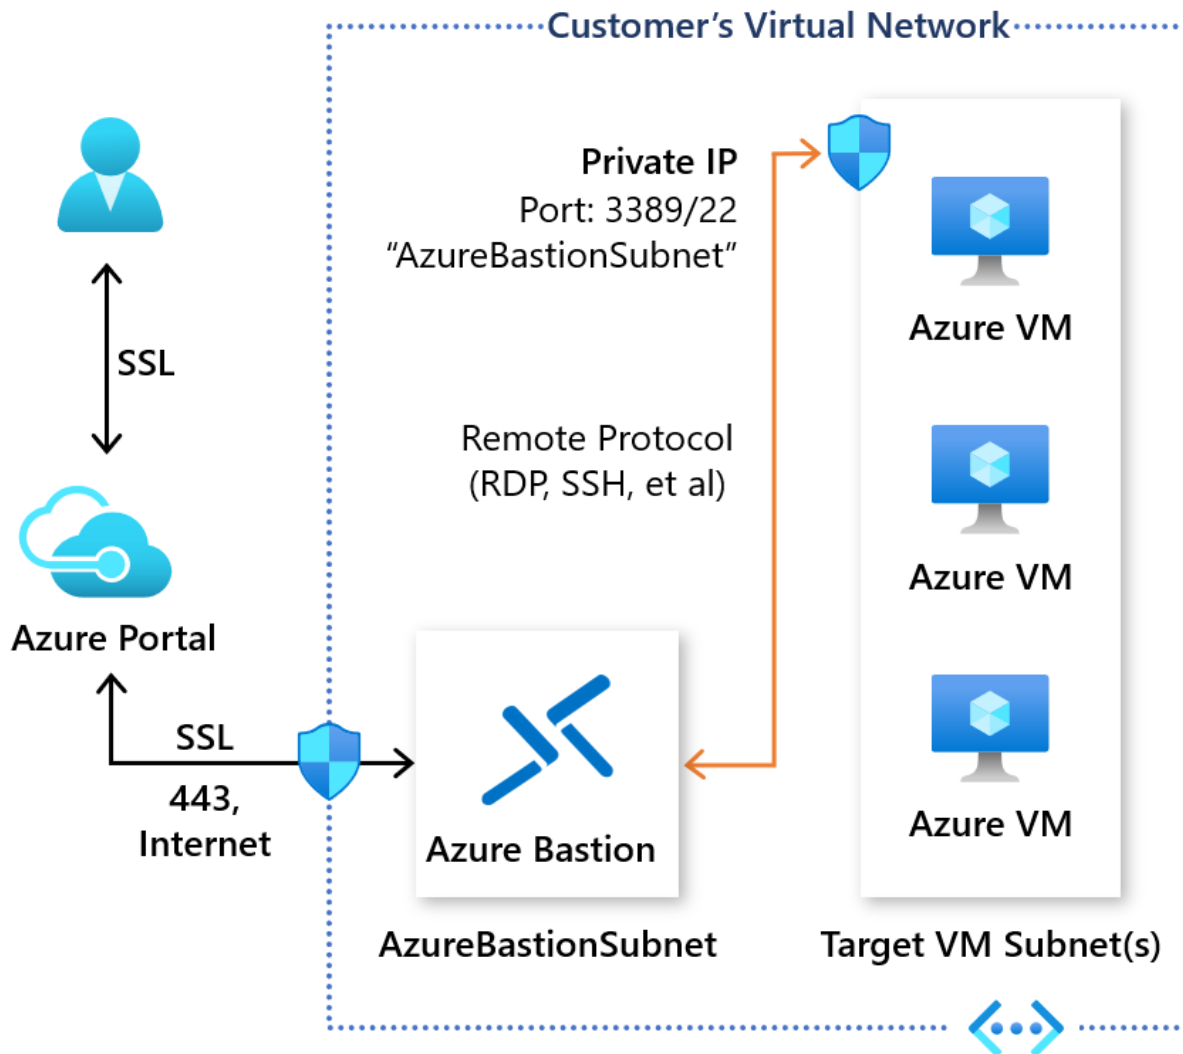
IMPORTANT

Azure Bastion is the service recommended to connect to the jump box to prevent exposing Azure VMware Solution to the internet. You cannot use Azure Bastion to connect to Azure VMware Solution VMs since they are not Azure IaaS objects.

As a security best practice, deploy [Microsoft Azure Bastion](#) service within the Hub virtual network. Azure Bastion provides seamless RDP and SSH access to VMs deployed on Azure without providing public IP addresses to those resources. Once you provision the Azure Bastion service, you can access the selected VM from the Azure portal. After establishing the connection, a new tab opens, showing the jump box desktop, and from that desktop, you can access the Azure VMware Solution private cloud management plane.

IMPORTANT

Do not give a public IP address to the jump box VM or expose 3389/TCP port to the public internet.



Azure DNS resolution considerations

For Azure DNS resolution, there are two options available:

- Use the domain controllers deployed on the Hub (described in [Identity considerations](#)) as name servers.
- Deploy and configure an Azure DNS private zone.

The best approach is to combine both to provide reliable name resolution for Azure VMware Solution, on-premises, and Azure.

As a general design recommendation, use the existing Active Directory-integrated DNS deployed onto at least two Azure VMs in the Hub virtual network and configured in the Spoke virtual networks to use those Azure DNS servers in the DNS settings.

You can use Azure Private DNS, where the Azure Private DNS zone links to the virtual network. The DNS servers are used as hybrid resolvers with conditional forwarding to on-premises or Azure VMware Solution running DNS using customer Azure Private DNS infrastructure.

To automatically manage the DNS records' lifecycle for the VMs deployed within the Spoke virtual networks, enable autoregistration. When enabled, the maximum number of private DNS zones is only one. If disabled, then the maximum number is 1000.

On-premises and Azure VMware Solution servers can be configured with conditional forwarders to resolver VMs in Azure for the Azure Private DNS zone.

Identity considerations

For identity purposes, the best approach is to deploy at least one domain controller on the Hub. Use two shared service subnets in zone-distributed fashion or a VM availability set. For more information on extending your on-premises Active Directory (AD) domain to Azure, see [Azure Architecture Center](#).

Additionally, deploy another domain controller on the Azure VMware Solution side to act as identity and DNS source within the vSphere environment.

As a recommended best practice, integrate [AD domain with Azure Active Directory](#).

Internet connectivity design considerations

12/16/2022 • 4 minutes to read • [Edit Online](#)

There are three primary patterns for creating outbound access to the Internet from Azure VMware Solution and to enable inbound Internet access to resources on your Azure VMware Solution private cloud.

- [Internet Service hosted in Azure](#)
- [Azure VMware Solution Managed SNAT](#)
- [Azure Public IPv4 address to NSX-T Data Center Edge](#)

Your requirements for security controls, visibility, capacity, and operations drive the selection of the appropriate method for delivery of Internet access to the Azure VMware Solution private cloud.

Internet Service hosted in Azure

There are multiple ways to generate a default route in Azure and send it towards your Azure VMware Solution private cloud or on-premises. The options are as follows:

- An Azure firewall in a Virtual WAN Hub.
- A third-party Network Virtual Appliance in a Virtual WAN Hub Spoke Virtual Network.
- A third-party Network Virtual Appliance in a Native Azure Virtual Network using Azure Route Server.
- A default route from on-premises transferred to Azure VMware Solution over Global Reach.

Use any of these patterns to provide an outbound SNAT service with the ability to control what sources are allowed out, to view the connection logs, and for some services, do further traffic inspection.

The same service can also consume an Azure Public IP and create an inbound DNAT from the Internet towards targets in Azure VMware Solution.

An environment can also be built that utilizes multiple paths for Internet traffic. One for outbound SNAT (for example, a third-party security NVA), and another for inbound DNAT (like a third party Load balancer NVA using SNAT pools for return traffic).

Azure VMware Solution Managed SNAT

A Managed SNAT service provides a simple method for outbound internet access from an Azure VMware Solution private cloud. Features of this service include the following.

- Easily enabled – select the radio button on the Internet Connectivity tab and all workload networks will have immediate outbound access to the Internet through a SNAT gateway.
- No control over SNAT rules, all sources that reach the SNAT service are allowed.
- No visibility into connection logs.
- Two Public IPs are used and rotated to support up to 128k simultaneous outbound connections.
- No inbound DNAT capability is available with the Azure VMware Solution Managed SNAT.

Azure Public IPv4 address to NSX-T Data Center Edge

This option brings an allocated Azure Public IPv4 address directly to the NSX-T Data Center Edge for consumption. It allows the Azure VMware Solution private cloud to directly consume and apply public network addresses in NSX-T Data Center as required. These addresses are used for the following types of connections:

- Outbound SNAT
- Inbound DNAT
- Load balancing using VMware NSX Advanced Load Balancer and other third-party Network Virtual Appliances
- Applications directly connected to a workload VM interface.

This option also lets you configure the public address on a third-party Network Virtual Appliance to create a DMZ within the Azure VMware Solution private cloud.

Features include:

- Scale – the soft limit of 64 Azure Public IPv4 addresses can be increased by request to 1,000s of Azure Public IPs allocated if required by an application.
- Flexibility – An Azure Public IPv4 address can be applied anywhere in the NSX-T Data Center ecosystem. It can be used to provide SNAT or DNAT, on load balancers like VMware’s NSX Advanced Load Balancer, or third-party Network Virtual Appliances. It can also be used on third-party Network Virtual Security Appliances on VMware segments or directly on VMs.
- Regionality – the Azure Public IPv4 address to the NSX-T Data Center Edge is unique to the local SDDC. For “multi private cloud in distributed regions,” with local exit to Internet intentions, it’s much easier to direct traffic locally versus trying to control default route propagation for a security or SNAT service hosted in Azure. If you’ve two or more Azure VMware Solution private clouds connected with a Public IP configured, they can both have a local exit.

Considerations for selecting an option

The option that you select depends on the following factors:

- To add an Azure VMware private cloud to a security inspection point provisioned in Azure native that inspects all Internet traffic from Azure native endpoints, use an Azure native construct and leak a default route from Azure to your Azure VMware Solution private cloud.
- If you need to run a third-party Network Virtual Appliance to conform to existing standards for security inspection or streamlined opex, you have two options. You can run your Azure Public IPv4 address in Azure native with the default route method or run it in Azure VMware Solution using Azure Public IPv4 address to NSX-T Data Center Edge.
- There are scale limits on how many Azure Public IPv4 addresses can be allocated to a Network Virtual Appliance running in native Azure or provisioned on Azure Firewall. The Azure Public IPv4 address to NSX-T Data Center Edge option allows for much higher allocations (1,000s versus 100s).
- Use an Azure Public IPv4 address to the NSX-T Data Center Edge for a localized exit to the internet from each private cloud in its local region. Using multiple Azure VMware Solution private clouds in several Azure regions that need to communicate with each other and the internet, it can be challenging to match an Azure VMware Solution private cloud with a security service in Azure. The difficulty is due to the way a default route from Azure works.

Next Steps

[Enable Managed SNAT for Azure VMware Solution Workloads](#)

[Enable Public IP to the NSX Edge for Azure VMware Solution](#)

[Disable Internet access or enable a default route](#)

Azure VMware Solution network design considerations

12/16/2022 • 6 minutes to read • [Edit Online](#)

Azure VMware Solution offers a VMware private cloud environment accessible for users and applications from on-premises and Azure-based environments or resources. The connectivity is delivered through networking services such as Azure ExpressRoute and VPN connections. There are several networking considerations to review before setting up your Azure VMware Solution environment. This article provides solutions for use cases you may encounter when configuring your networking with Azure VMware Solution.

Azure VMware Solution compatibility with AS-Path Prepend

Azure VMware Solution is incompatible with AS-Path Prepend for redundant ExpressRoute configurations and doesn't honor the outbound path selection from Azure towards on-premises. If you're running 2 or more ExpressRoute paths between on-premises and Azure, and the listed [Prerequisites](#) are true; you may experience impaired connectivity or no connectivity between your on-premises networks and Azure VMware Solution. The connectivity issue is caused when Azure VMware Solution doesn't see the AS-Path Prepend and uses ECMP to send traffic towards your environment over both ExR circuits. That action causes issues with stateful firewall inspection.

Prerequisites

For AS-Path Prepend, you'll need to verify that all of the following listed connections are true:

- Both or all circuits are connected to Azure VMware Solution with global reach.
- The same netblocks are being advertised from two or more circuits.
- Stateful firewalls are in the network path.
- You're using AS-Path Prepend to force Azure to prefer one path over others.

Either 2 or 4 byte Public ASN numbers should be used and be compatible with Azure VMware Solution. If you don't own a Public ASN to use for prepending, open a [Microsoft Customer Support Ticket](#) to view options.

Management VMs and default routes from on-premises

IMPORTANT

Azure VMware Solution Management VMs don't honor a default route from on-premises.

If you're routing back to your on-premises networks using only a default route advertised towards Azure, the vCenter Server and NSX-T Manager VMs won't be compatible with that route.

Solution

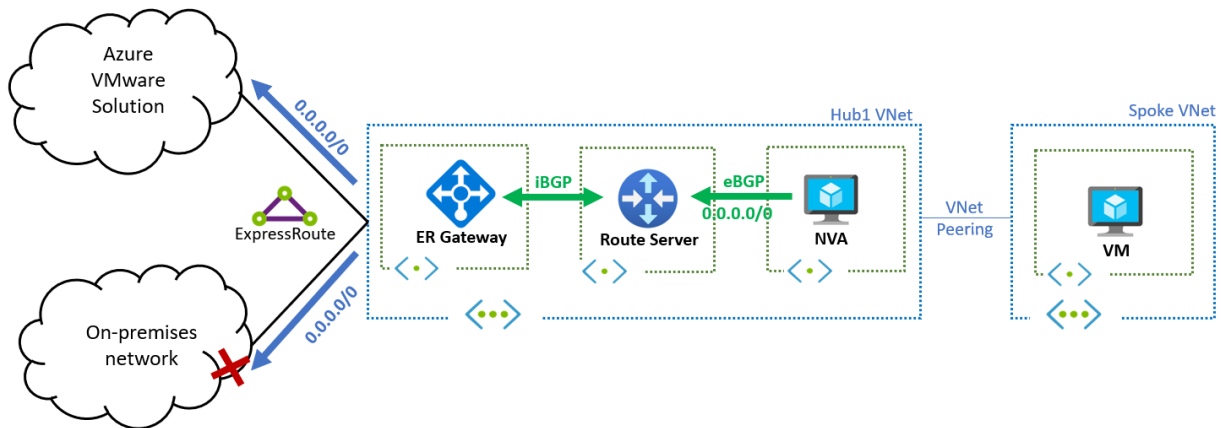
To reach vCenter Server and NSX-T Manager, more specific routes from on-premises need to be provided to allow traffic to have a return path route to those networks.

Use a default route to Azure VMware Solution for internet traffic inspection

Certain deployments require inspecting all egress traffic from Azure VMware Solution towards the Internet.

While it's possible to create Network Virtual Appliances (NVAs) in Azure VMware Solution, there are use cases when these appliances already exist in Azure that can be applied to inspect Internet traffic from Azure VMware Solution. In this case, a default route can be injected from the NVA in Azure to attract traffic from Azure VMware Solution and inspect it before sending it out to the public Internet.

The following diagram describes a basic hub and spoke topology connected to an Azure VMware Solution cloud and to an on-premises network through ExpressRoute. The diagram shows how the default route ($0.0.0.0/0$) is originated by the NVA in Azure, and propagated by Azure Route Server to Azure VMware Solution through ExpressRoute.



IMPORTANT

The default route advertised by the NVA will be propagated to the on-premises network. Because of that, UDRs will need to be added to ensure traffic from Azure VMware Solution is transiting through the NVA.

Communication between Azure VMware Solution and the on-premises network usually occurs over ExpressRoute Global Reach, as described in [Peer on-premises environments to Azure VMware Solution](#).

Connectivity between Azure VMware Solution and on-premises network via a third party network virtual appliance

There are two main scenarios for this connectivity pattern:

- Organizations may have the requirement to send traffic between Azure VMware Solution and the on-premises network through an NVA (typically a firewall).
- ExpressRoute Global Reach might not be available in a particular region to interconnect the ExpressRoute circuits of Azure VMware Solution and the on-premises network.

There are two topologies you can apply to meet all requirements for these two scenarios. The first is a [Supernet topology](#) and the second is a [Transit spoke virtual network topology](#).

IMPORTANT

The preferred option to connect Azure VMware Solution and on-premises environments is a direct ExpressRoute Global Reach connection. The patterns described in this document add considerable complexity to the environment.

Supernet design topology

If both ExpressRoute circuits (to Azure VMware Solution and to on-premises) are terminated in the same ExpressRoute gateway, you can assume that the gateway is going to route packets across them. However, an ExpressRoute gateway isn't designed to do that. You need to hairpin the traffic to an NVA that can route the

traffic. There are two requirements to hairpin network traffic to an NVA:

- The NVA should advertise a supernet for the Azure VMware Solution and on-premises prefixes.

You could use a supernet that includes both Azure VMware Solution and on-premises prefixes, or individual prefixes for Azure VMware Solution and on-premises (always less specific than the actual prefixes advertised over ExpressRoute). Keep in mind that all supernet prefixes advertised to Route Server are going to be propagated both to Azure VMware Solution and on-premises.

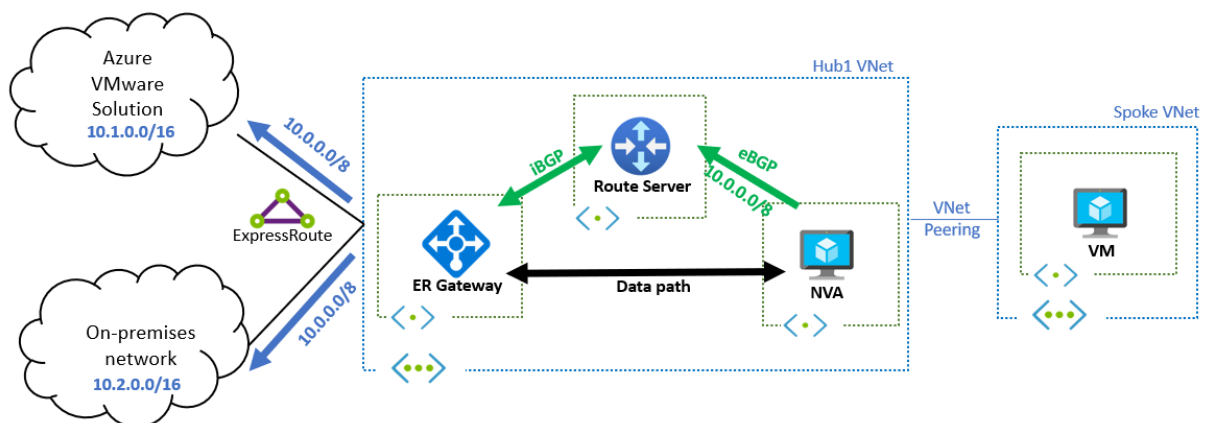
- UDRs in the GatewaySubnet that exactly match the prefixes advertised from Azure VMware Solution and on-premises will cause hairpin traffic from the GatewaySubnet to the NVA.

This topology results in high management overhead for large networks that change over time. Note that there are specific limitations to be considered.

Limitations

- Anytime a workload segment is created in Azure VMware Solution, UDRs may need to be added to ensure traffic from Azure VMware Solution is transiting through the NVA.
- If your on-premises environment has a large number of routes that change, BGP and UDR configuration in the supernet may need to be updated.
- Since there's a single ExpressRoute Gateway that processes network traffic in both directions, performance may be limited.
- There's an Azure Virtual Network limit of 400 UDRs.

The following diagram demonstrates how the NVA needs to advertise more generic (less specific) prefixes that include the networks from on-premises and Azure VMware Solution. Be careful with this approach as the NVA could potentially attract traffic that it shouldn't (since it's advertising wider ranges, for example: the whole `10.0.0.0/8` network).



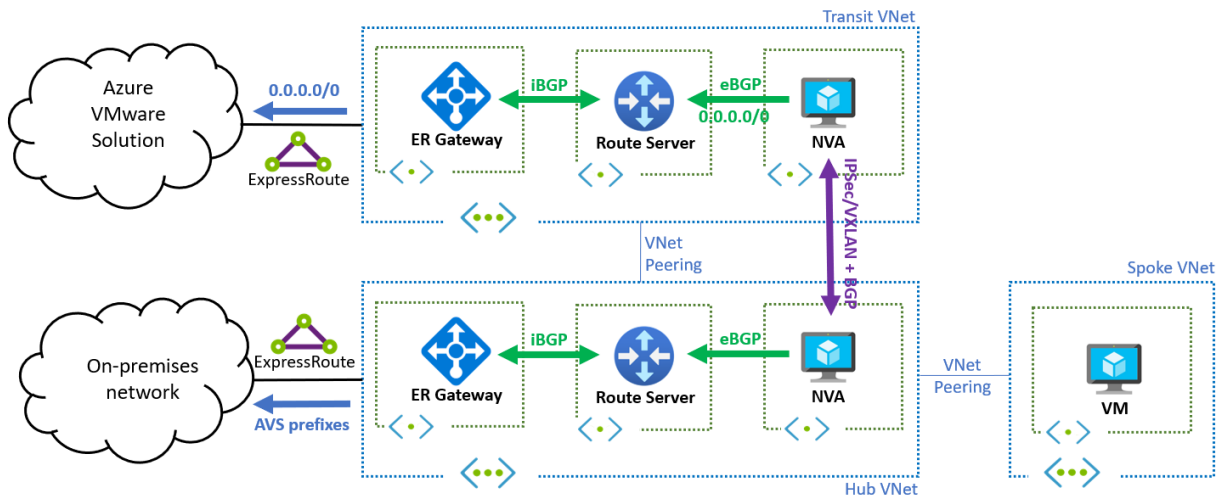
Transit spoke virtual network topology

NOTE

If advertising less specific prefixes is not possible due to the limits previously described, you can implement an alternative design using two separate Virtual Networks.

In this topology, instead of propagating less specific routes to attract traffic to the ExpressRoute gateway, two different NVAs in separate Virtual Networks can exchange routes between each other. The Virtual Networks can propagate these routes to their respective ExpressRoute circuits via BGP and Azure Route Server, as the following diagram shows. Each NVA has full control on which prefixes are propagated to each ExpressRoute circuit.

The following diagram demonstrates how a single 0.0.0.0/0 is advertised to Azure VMware Solution. It also shows how the individual Azure VMware Solution prefixes are propagated to the on-premises network.



IMPORTANT

An encapsulation protocol such as VXLAN or IPsec is required between the NVAs. Encapsulation is needed because the NVA NICs would learn the routes from Azure Route Server with the NVA as next hop and create a routing loop.

There's an alternative to using an overlay. Apply secondary NICs in the NVA that won't learn the routes from Azure Route Server and configure UDRs so that Azure can route traffic to the remote environment over those NICs. You can find more details in [Enterprise-scale network topology and connectivity for Azure VMware Solution](#).

This topology requires a complex initial set-up. Once the set-up is complete, the topology works as expected with minimal management overhead. See the following list of specific set-up complexities.

- There's an extra cost for an additional transit Virtual Network that includes an Azure Route Server, ExpressRoute Gateway, and another NVA. The NVAs may also need to use large VM sizes to meet throughput requirements.
- There's IPsec or VxLAN tunneling between the two NVAs required which means that the NVAs are also in the datapath. Depending on the type of NVA you're using, it can result in custom and complex configuration on those NVAs.

Next steps

Now that you've covered Azure VMware Solution network design considerations, you may want to learn more about:

- [Network interconnectivity concepts - Azure VMware Solution](#)
- [Plan the Azure VMware Solution deployment](#)
- [Networking planning checklist for Azure VMware Solution](#)

Recommended content

- [Tutorial - Configure networking for your VMware private cloud in Azure - Azure VMware Solution](#)

Azure VMware Solution networking and interconnectivity concepts

12/16/2022 • 5 minutes to read • [Edit Online](#)

Azure VMware Solution offers a private cloud environment accessible from on-premises sites and Azure-based resources. Services such as Azure ExpressRoute, VPN connections, or Azure Virtual WAN deliver the connectivity. However, these services require specific network address ranges and firewall ports for enabling the services.

When you deploy a private cloud; private networks for management, provisioning, and vMotion get created. You'll use these private networks to access VMware vCenter Server and VMware NSX-T Data Center NSX-T Manager and virtual machine vMotion or deployment.

[ExpressRoute Global Reach](#) is used to connect private clouds to on-premises environments. It connects circuits directly at the Microsoft Enterprise Edge (MSEE) level. The connection requires a virtual network (vNet) with an ExpressRoute circuit to on-premises in your subscription. The reason is that vNet gateways (ExpressRoute Gateways) can't transit traffic, which means you can attach two circuits to the same gateway, but it won't send the traffic from one circuit to the other.

Each Azure VMware Solution environment is its own ExpressRoute region (its own virtual MSEE device), which lets you connect Global Reach to the 'local' peering location. It allows you to connect multiple Azure VMware Solution instances in one region to the same peering location.

NOTE

For locations where ExpressRoute Global Reach isn't enabled, for example, because of local regulations, you have to build a routing solution using Azure IaaS VMs. For some examples, see [Azure Cloud Adoption Framework - Network topology and connectivity for Azure VMware Solution](#).

Virtual machines deployed on the private cloud are accessible to the internet through the [Azure Virtual WAN public IP](#) functionality. For new private clouds, internet access is disabled by default.

There are two ways to interconnectivity in the Azure VMware Solution private cloud:

- **Basic Azure-only interconnectivity** lets you manage and use your private cloud with only a single virtual network in Azure. This implementation is best suited for Azure VMware Solution evaluations or implementations that don't require access from on-premises environments.
- **Full on-premises to private cloud interconnectivity** extends the basic Azure-only implementation to include interconnectivity between on-premises and Azure VMware Solution private clouds.

This article covers the key concepts that establish networking and interconnectivity, including requirements and limitations. In addition, this article provides you with the information you need to know to work with Azure VMware Solution to configure your networking.

Azure VMware Solution private cloud use cases

The use cases for Azure VMware Solution private clouds include:

- New VMware vSphere VM workloads in the cloud
- VM workload bursting to the cloud (on-premises to Azure VMware Solution only)
- VM workload migration to the cloud (on-premises to Azure VMware Solution only)

- Disaster recovery (Azure VMware Solution to Azure VMware Solution or on-premises to Azure VMware Solution)
- Consumption of Azure services

TIP

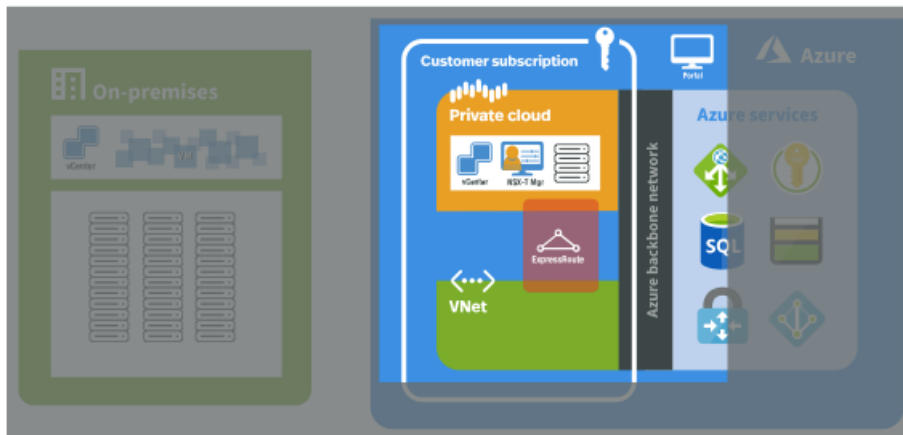
All use cases for the Azure VMware Solution service are enabled with on-premises to private cloud connectivity.

Azure virtual network interconnectivity

You can interconnect your Azure virtual network with the Azure VMware Solution private cloud implementation. You can manage your Azure VMware Solution private cloud, consume workloads in your private cloud, and access other Azure services.

The diagram below shows the basic network interconnectivity established at the time of a private cloud deployment. It shows the logical networking between a virtual network in Azure and a private cloud. This connectivity is established via a backend ExpressRoute that is part of the Azure VMware Solution service. The interconnectivity fulfills the following primary use cases:

- Inbound access to vCenter Server and NSX-T Manager that is accessible from VMs in your Azure subscription.
- Outbound access from VMs on the private cloud to Azure services.
- Inbound access of workloads running in the private cloud.

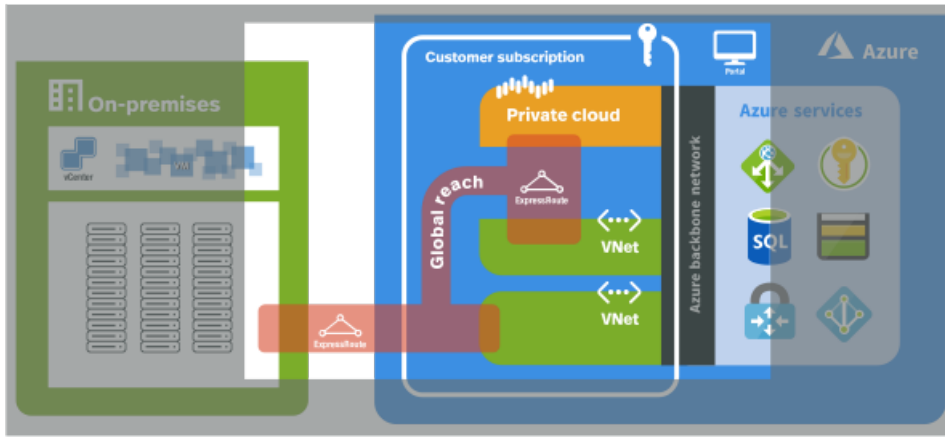


On-premises interconnectivity

In the fully interconnected scenario, you can access the Azure VMware Solution from your Azure virtual network(s) and on-premises. This implementation is an extension of the basic implementation described in the previous section. An ExpressRoute circuit is required to connect from on-premises to your Azure VMware Solution private cloud in Azure.

The diagram below shows the on-premises to private cloud interconnectivity, which enables the following use cases:

- Hot/Cold vSphere vMotion between on-premises and Azure VMware Solution.
- On-premises to Azure VMware Solution private cloud management access.



For full interconnectivity to your private cloud, you need to enable ExpressRoute Global Reach and then request an authorization key and private peering ID for Global Reach in the Azure portal. The authorization key and peering ID are used to establish Global Reach between an ExpressRoute circuit in your subscription and the ExpressRoute circuit for your private cloud. Once linked, the two ExpressRoute circuits route network traffic between your on-premises environments to your private cloud. For more information on the procedures, see the [tutorial for creating an ExpressRoute Global Reach peering to a private cloud](#).

Limitations

The following table describes the maximum limits for Azure VMware Solution.

RESOURCE	LIMIT
vSphere clusters per private cloud	12
Minimum number of ESXi hosts per cluster	3 (hard-limit)
Maximum number of ESXi hosts per cluster	16 (hard-limit)
Maximum number of ESXi hosts per private cloud	96
Maximum number of vCenter Servers per private cloud	1 (hard-limit)
Maximum number of HCX site pairings	25 (any edition)
Maximum number of HCX service meshes	10 (any edition)
Maximum number of Azure VMware Solution ExpressRoute max linked private clouds	4 The virtual network gateway used determines the actual max linked private clouds. For more details, see About ExpressRoute virtual network gateways
Maximum Azure VMware Solution ExpressRoute port speed	10 Gbps The virtual network gateway used determines the actual bandwidth. For more details, see About ExpressRoute virtual network gateways
Maximum number of Azure Public IPv4 addresses assigned to NSX-T Data Center	2,000

RESOURCE	LIMIT
Maximum number of Azure VMware Solution Interconnects per private cloud	10
vSAN capacity limits	75% of total usable (keep 25% available for SLA)
VMware Site Recovery Manager - Maximum number of protected Virtual Machines	3,000
VMware Site Recovery Manager - Maximum number of Virtual Machines per recovery plan	2,000
VMware Site Recovery Manager - Maximum number of protection groups per recovery plan	250
VMware Site Recovery Manager - RPO Values	5 min or higher * (hard-limit)
VMware Site Recovery Manager - Maximum number of virtual machines per protection group	500
VMware Site Recovery Manager - Maximum number of recovery plans	250

* For information about Recovery Point Objective (RPO) lower than 15 minutes, see [How the 5 Minute Recovery Point Objective Works](#) in the *vSphere Replication Administration guide*.

For other VMware-specific limits, use the [VMware configuration maximum tool](#).

Next steps

Now that you've covered Azure VMware Solution network and interconnectivity concepts, you may want to learn about:

- [Azure VMware Solution storage concepts](#)
- [Azure VMware Solution identity concepts](#)
- [Enabling the Azure VMware Solution resource provider](#)

Azure VMware Solution private cloud and cluster concepts

12/16/2022 • 11 minutes to read • [Edit Online](#)

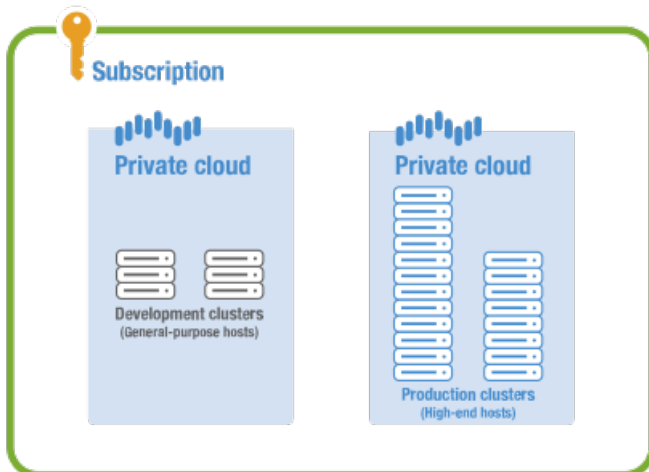
Azure VMware Solution delivers VMware-based private clouds in Azure. The private cloud hardware and software deployments are fully integrated and automated in Azure. You deploy and manage the private cloud through the Azure portal, CLI, or PowerShell.

A private cloud includes clusters with:

- Dedicated bare-metal server hosts provisioned with VMware ESXi hypervisor
- VMware vCenter Server for managing ESXi and vSAN
- VMware NSX-T Data Center software-defined networking for vSphere workload VMs
- VMware vSAN datastore for vSphere workload VMs
- VMware HCX for workload mobility
- Resources in the Azure underlay (required for connectivity and to operate the private cloud)

As with other resources, private clouds are installed and managed from within an Azure subscription. The number of private clouds within a subscription is scalable. Initially, there's a limit of one private cloud per subscription. There's a logical relationship between Azure subscriptions, Azure VMware Solution private clouds, vSAN clusters, and hosts.

The diagram shows a single Azure subscription with two private clouds that represent a development and production environment. In each of those private clouds are two clusters.



Hosts

Azure VMware Solution clusters are based upon hyper-converged infrastructure. The following table shows the CPU, memory, disk and network specifications of the host.

HOST TYPE	CPU (GHZ)	RAM (GB)	VSAN CACHE TIER (TB, RAW)	VSAN CAPACITY TIER (TB, RAW)	NETWORK INTERFACE CARDS	REGIONAL AVAILABILITY
AV36	Dual Intel Xeon Gold 6140 CPUs with 18 cores/CPU @ 2.3 GHz, Total 36 physical cores (72 logical cores with hyperthreading)	576	3.2 (NVMe)	15.20 (SSD)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	All product regions
AV36P	Dual Intel Xeon Gold 6240 CPUs with 18 cores/CPU @ 2.6 GHz / 3.9 GHz Turbo, Total 36 physical cores (72 logical cores with hyperthreading)	768	1.5 (Intel Cache)	19.20 (NVMe)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	Selected regions (*)
AV52	Dual Intel Xeon Platinum 8270 CPUs with 26 cores/CPU @ 2.7 GHz / 4.0 GHz Turbo, Total 52 physical cores (104 logical cores with hyperthreading)	1,536	1.5 (Intel Cache)	38.40 (NVMe)	4x 25 Gb/s NICs (2 for management & control plane, 2 for customer traffic)	Selected regions (*)

An Azure VMware Solution cluster requires a minimum number of three hosts. You can only use hosts of the same type in a single Azure VMware Solution private cloud. Hosts used to build or scale clusters come from an isolated pool of hosts. Those hosts have passed hardware tests and have had all data securely deleted before being added to a cluster.

(*) details available via the Azure pricing calculator.

Clusters

For each private cloud created, there's one vSAN cluster by default. You can add, delete, and scale clusters. The minimum number of hosts per cluster and the initial deployment is three.

You use vCenter Server and NSX-T Manager to manage most aspects of cluster configuration and operation. All local storage of each host in a cluster is under the control of vSAN.

The Azure VMware Solution management and control plane has the following resource requirements that need to be accounted for during solution sizing.

AREA	DESCRIPTION	PROVISIONED VCPUS	PROVISIONED VRAM (GB)	PROVISIONED VDISK (GB)	TYPICAL CPU USAGE (GHZ)	TYPICAL VRAM USAGE (GB)	TYPICAL RAW VSAN DATASTORE USAGE (GB)
VMware vSphere	vCenter Server	8	28	915	1.1	3.6	1,925
VMware vSphere	vSphere Cluster Service VM 1	1	0.1	5	0.1	0.1	2
VMware vSphere	vSphere Cluster Service VM 2	1	0.1	5	0.1	0.1	2
VMware vSphere	vSphere Cluster Service VM 3	1	0.1	5	0.1	0.1	2
VMware vSphere	ESXi node 1	N/A	N/A	N/A	9.4	0.4	N/A
VMware vSphere	ESXi node 2	N/A	N/A	N/A	9.4	0.4	N/A
VMware vSphere	ESXi node 3	N/A	N/A	N/A	9.4	0.4	N/A
VMware vSAN	vSAN System Usage	N/A	N/A	N/A	N/A	N/A	6,574
VMware NSX-T Data Center	NSX-T Unified Appliance Node 1	6	24	300	5.5	8.5	613
VMware NSX-T Data Center	NSX-T Unified Appliance Node 2	6	24	300	5.5	8.5	613
VMware NSX-T Data Center	NSX-T Unified Appliance Node 3	6	24	300	5.5	8.5	613
VMware NSX-T Data Center	NSX-T Edge VM 1	8	32	200	1.3	0.6	409

AREA	DESCRIPTION	PROVISIONED VCPUS	PROVISIONED VRAM (GB)	PROVISIONED VDISK (GB)	TYPICAL CPU USAGE (GHZ)	TYPICAL VRAM USAGE (GB)	TYPICAL RAW VSAN DATASTORE USAGE (GB)
VMware NSX-T Data Center	NSX-T Edge VM 2	8	32	200	1.3	0.6	409
VMware HCX (Optional Add-On)	HCX Manager	4	12	65	1	3.2	152
VMware Site Recovery Manager (Optional Add-On)	SRM Appliance	4	12	33	1	1	93
VMware vSphere (Optional Add-On)	vSphere Replication Manager Appliance	4	8	33	4.3	2.2	84
VMware vSphere (Optional Add-On)	vSphere Replication Server Appliance	2	1	33	1	0.1	84
	Total	59 vCPUs	197.3 GB	2,394 GB	56 GHz	38.3 GB	11,575 GB (9,646 GB with expected 1.2x Data Reduction ratio)

These resource requirements only apply to the first cluster deployed in an Azure VMware Solution private cloud. Subsequent clusters only need to account for the vSphere Cluster Service, ESXi resource requirements and vSAN System Usage in solution sizing.

The virtual appliance **Typical Raw vSAN Datastore Usage** values account for the space occupied by virtual machine files, including configuration and log files, snapshots, virtual disks and swap files.

The VMware ESXi nodes have compute usage values that account for the vSphere VMkernel hypervisor overhead, vSAN overhead and NSX-T distributed router, firewall and bridging overhead. These are estimates for a standard three cluster configuration. The storage requirements are listed as not applicable (N/A) since a boot volume separate from the vSAN Datastore is used.

The VMware vSAN System Usage storage overhead accounts for vSAN performance management objects, vSAN file system overhead, vSAN checksum overhead and vSAN deduplication and compression overhead. To view this consumption, select the Monitor, vSAN Capacity object for the vSphere Cluster in the vSphere Client.

The VMware HCX and VMware Site Recovery Manager resource requirements are optional Add-Ons to the Azure VMware Solution service. Discount these requirements in the solution sizing if they are not being used.

The VMware Site Recovery Manager Add-On has the option of configuring multiple VMware vSphere Replication Server Appliances. The table above assumes one vSphere Replication Server appliance is used.

Sizing an Azure VMware Solution is an estimate; the sizing calculations from the design phase should be validated during the testing phase of a project to ensure the Azure VMware Solution has been sized correctly for the application workload.

TIP

You can always extend the cluster and add additional clusters later if you need to go beyond the initial deployment number.

The following table describes the maximum limits for Azure VMware Solution.

RESOURCE	LIMIT
vSphere clusters per private cloud	12
Minimum number of ESXi hosts per cluster	3 (hard-limit)
Maximum number of ESXi hosts per cluster	16 (hard-limit)
Maximum number of ESXi hosts per private cloud	96
Maximum number of vCenter Servers per private cloud	1 (hard-limit)
Maximum number of HCX site pairings	25 (any edition)
Maximum number of HCX service meshes	10 (any edition)
Maximum number of Azure VMware Solution ExpressRoute max linked private clouds	4 The virtual network gateway used determines the actual max linked private clouds. For more details, see About ExpressRoute virtual network gateways
Maximum Azure VMware Solution ExpressRoute port speed	10 Gbps The virtual network gateway used determines the actual bandwidth. For more details, see About ExpressRoute virtual network gateways
Maximum number of Azure Public IPv4 addresses assigned to NSX-T Data Center	2,000
Maximum number of Azure VMware Solution Interconnects per private cloud	10
vSAN capacity limits	75% of total usable (keep 25% available for SLA)
VMware Site Recovery Manager - Maximum number of protected Virtual Machines	3,000
VMware Site Recovery Manager - Maximum number of Virtual Machines per recovery plan	2,000
VMware Site Recovery Manager - Maximum number of protection groups per recovery plan	250

RESOURCE	LIMIT
VMware Site Recovery Manager - RPO Values	5 min or higher * (hard-limit)
VMware Site Recovery Manager - Maximum number of virtual machines per protection group	500
VMware Site Recovery Manager - Maximum number of recovery plans	250

* For information about Recovery Point Objective (RPO) lower than 15 minutes, see [How the 5 Minute Recovery Point Objective Works](#) in the *vSphere Replication Administration guide*.

For other VMware-specific limits, use the [VMware configuration maximum tool](#).

VMware software versions

The VMware solution software versions used in new deployments of Azure VMware Solution private cloud clusters are:

SOFTWARE	VERSION
VMware vCenter Server	7.0 U3c
ESXi	7.0 U3c
vSAN	7.0 U3c
vSAN on-disk format	10
HCX	4.4.2
VMware NSX-T Data Center NOTE: VMware NSX-T Data Center is the only supported version of NSX Data Center.	3.1.2

The current running software version is applied to new clusters added to an existing private cloud. For more information, see the [VMware software version requirements for HCX](#) and [Understanding vSAN on-disk format versions and compatibility](#).

Host maintenance and lifecycle management

One benefit of Azure VMware Solution private clouds is that the platform is maintained for you. Microsoft is responsible for the lifecycle management of VMware software (ESXi, vCenter Server, and vSAN). Microsoft is also responsible for the lifecycle management of the NSX-T Data Center appliances and bootstrapping the network configuration, like creating the Tier-0 gateway and enabling North-South routing. You're responsible for the NSX-T Data Center SDN configuration: network segments, distributed firewall rules, Tier 1 gateways, and load balancers.

NOTE

A T0 gateway is created and configured as part of a private cloud deployment. Any modification to that logical router or the NSX-T Data Center edge node VMs could affect connectivity to your private cloud and should be avoided.

Microsoft is responsible for applying any patches, updates, or upgrades to ESXi, vCenter Server, vSAN, and NSX-T Data Center in your private cloud. The impact of patches, updates, and upgrades on ESXi, vCenter Server, and NSX-T Data Center is different.

- **ESXi** - There's no impact to workloads running in your private cloud. Access to vCenter Server and NSX-T Data Center isn't blocked during this time. It's recommended that, during this time, you don't plan any other activities like: scaling up private cloud, scheduling or initiating active HCX migrations, making HCX configuration changes, and so on, in your private cloud.
- **vCenter Server** - There's no impact to workloads running in your private cloud. During this time, vCenter Server will be unavailable and you won't be able to manage VMs (stop, start, create, or delete). It's recommended that, during this time, you don't plan any other activities like scaling up private cloud, creating new networks, and so on, in your private cloud. If you're using VMware Site Recovery Manager or vSphere Replication user interfaces, it's recommended to not configure vSphere Replication, and configure or execute site recovery plans during the vCenter Server upgrade.
- **NSX-T Data Center** - There's workload impact and when a particular host is being upgraded, the VMs on that host might lose connectivity from 2 seconds to maximum 1 minute with any of the following symptoms:
 - Ping errors
 - Packet loss
 - Error messages (for example, *Destination Host Unreachable* and *Net unreachable*)

During this upgrade window, all access to the NSX-T Data Center management plane will be blocked. You can't make configuration changes to the NSX-T Data Center environment for the duration. However, your workloads will continue to run as normal, subject to the upgrade impact detailed above.

It's recommended that, during the upgrade time, you don't plan any other activities like scaling up private cloud, and so on, in your private cloud. Other activities can prevent the upgrade from starting or could have adverse impacts on the upgrade and the environment.

You'll be notified before patches/updates or upgrades are applied to your private clouds. We'll also work with you to schedule a maintenance window before applying updates or upgrades to your private cloud.

Software updates include:

- **Patches** - Security patches or bug fixes released by VMware
- **Updates** - Minor version change of a VMware stack component
- **Upgrades** - Major version change of a VMware stack component

NOTE

Microsoft tests a critical security patch as soon as it becomes available from VMware.

Documented VMware workarounds are implemented in lieu of installing a corresponding patch until the next scheduled updates are deployed.

Host monitoring and remediation

Azure VMware Solution continuously monitors the health of both the underlay and the VMware components. When Azure VMware Solution detects a failure, it takes action to repair the failed components. When Azure VMware Solution detects a degradation or failure on an Azure VMware Solution node, it triggers the host remediation process.

Host remediation involves replacing the faulty node with a new healthy node in the cluster. Then, when possible, the faulty host is placed in VMware vSphere maintenance mode. VMware vMotion moves the VMs off the faulty host to other available servers in the cluster, potentially allowing zero downtime for live migration of workloads. If the faulty host can't be placed in maintenance mode, the host is removed from the cluster.

Azure VMware Solution monitors the following conditions on the host:

- Processor status
- Memory status
- Connection and power state
- Hardware fan status
- Network connectivity loss
- Hardware system board status
- Errors occurred on the disk(s) of a vSAN host
- Hardware voltage
- Hardware temperature status
- Hardware power status
- Storage status
- Connection failure

NOTE

Azure VMware Solution tenant admins must not edit or delete the above defined VMware vCenter Server alarms, as these are managed by the Azure VMware Solution control plane on vCenter Server. These alarms are used by Azure VMware Solution monitoring to trigger the Azure VMware Solution host remediation process.

Backup and restoration

Private cloud vCenter Server and NSX-T Data Center configurations are on an hourly backup schedule. Backups are kept for three days. If you need to restore from a backup, open a [support request](#) in the Azure portal to request restoration.

Azure VMware Solution continuously monitors the health of both the physical underlay and the VMware Solution components. When Azure VMware Solution detects a failure, it takes action to repair the failed components.

Next steps

Now that you've covered Azure VMware Solution private cloud concepts, you may want to learn about:

- [Azure VMware Solution networking and interconnectivity concepts](#)
- [Azure VMware Solution storage concepts](#)
- [How to enable Azure VMware Solution resource](#)

Run command in Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

In Azure VMware Solution, vCenter Server has a built-in local user called *cloudadmin* assigned to the CloudAdmin role. The CloudAdmin role has vCenter Server [privileges](#) that differ from other VMware cloud solutions and on-premises deployments. The Run command feature lets you perform operations that would normally require elevated privileges through a collection of PowerShell cmdlets.

Azure VMware Solution supports the following operations:

- [Configure an external identity source](#)
- [View and set storage policies](#)
- [Deploy disaster recovery using JetStream](#)

NOTE

Run commands are executed one at a time in the order submitted.

View the status of an execution

You can view the status of any executed run command, including the output, errors, warnings, and information logs of the cmdlets.

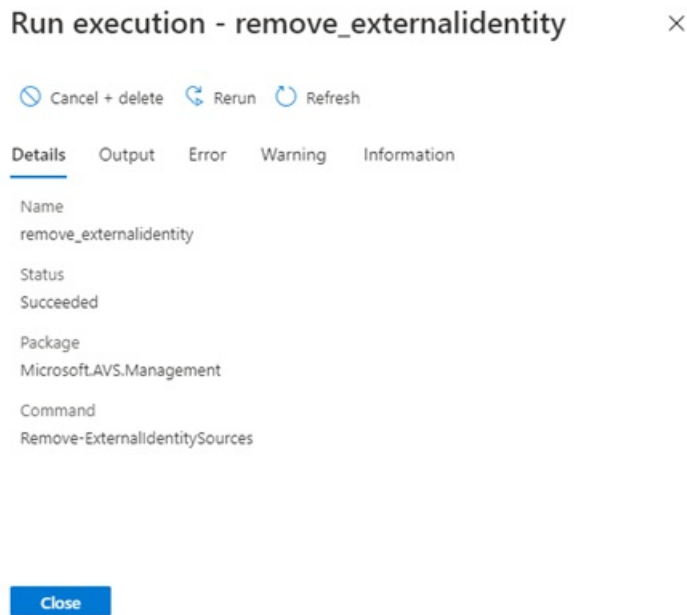
1. Sign in to the [Azure portal](#).
2. Select **Run command** > **Run execution status**.

You can sort by the various columns by selecting the column.

The screenshot shows the Azure portal interface for a resource named 'Contoso-westus-sddc'. The 'Run command' section is active, and the 'Run execution status' link is highlighted. The table below lists the execution details for various commands.

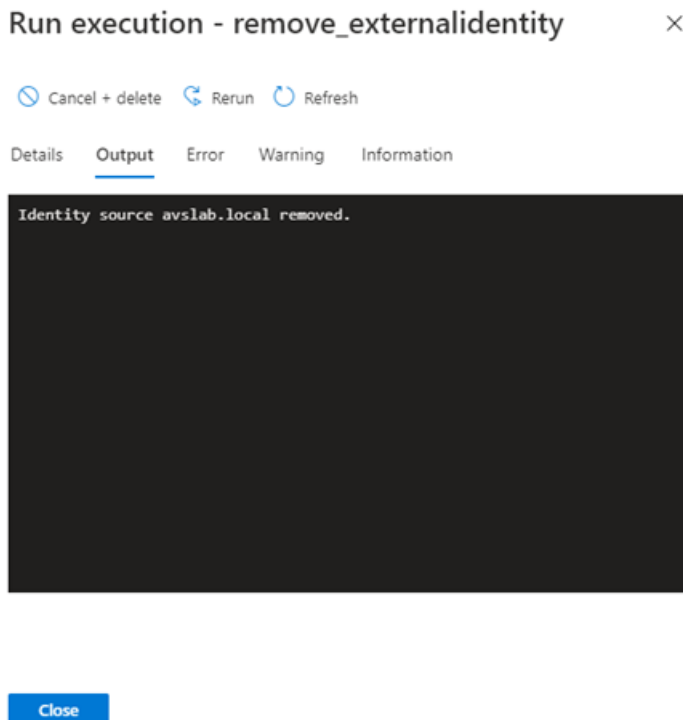
Execution name	Package name	Packag...	Command name	Started time...	End time stamp	Status
remove_externalidentity	Microsoft.AVS.Management	1.0.30	Remove-ExternalId	7/20/2021, 12:58:4	7/20/2021, 12:59:3	✔ Succeeded
removeGroup	Microsoft.AVS.Management	1.0.30	Remove-GroupFroi	7/20/2021, 12:52:0	7/20/2021, 12:53:4	✔ Succeeded
addADgroup	Microsoft.AVS.Management	1.0.30	Add-GroupToClou	7/20/2021, 12:08:4	7/20/2021, 12:09:2	✔ Succeeded
addexternalidentity	Microsoft.AVS.Management	1.0.30	New-AvsLDAPiden	7/20/2021, 11:13:2	7/20/2021, 11:14:0	✔ Succeeded
getidentitysource	Microsoft.AVS.Management	1.0.30	Get-ExternalIdentit	7/20/2021, 10:40:3	7/20/2021, 10:45:3	✔ Succeeded
check_jetserverdetails	JSDR.Configuration	1.0.20	Invoke-PreflightJet	7/20/2021, 7:52:56	7/20/2021, 8:04:57	✘ Failed
checkDRsystem	JSDR.Configuration	1.0.20	Invoke-PreflightJet	7/20/2021, 5:56:33	7/20/2021, 5:57:54	✔ Succeeded
amaneja-jsdrcheck	JSDR.Configuration	1.0.20	Invoke-PreflightJet	7/19/2021, 2:11:41	7/19/2021, 2:13:09	✔ Succeeded
installJSDR_withDNSRegistered	JSDR.Configuration	1.0.20	Install-JetDR	7/16/2021, 8:26:40	7/16/2021, 8:33:09	✘ Failed
del_jetdr_4	JSDR.Configuration	1.0.20	Uninstall-JetDR	7/16/2021, 8:17:22	7/16/2021, 8:18:06	✘ Failed
del_jetdr_3	JSDR.Configuration	1.0.20	Uninstall-JetDR	7/16/2021, 8:07:16	7/16/2021, 8:08:26	✔ Succeeded

3. Select the execution you want to view. A pane opens with details about the execution, and other tabs for the various types of output generated by the cmdlet.



You can view more details about the execution including the output, errors, warnings, and information.

- **Details** - Summary of the execution details, such as the name, status, package, cmdlet name, and error if the command failed.
- **Output** - Messages output by the cmdlet. May include progress or the result of the operation. Not all cmdlets have output.



- **Error** - Error messages generated in the execution of the cmdlet. This is in addition to the terminating error message on the details pane.

Run execution - remove_externalidentity ×

 Cancel + delete  Rerun  Refresh

Details Output **Error** Warning Information

```
VCenter already has JetDR plugin. Please unregister vCenter or cleanup any
```

Close

- **Warning** - Warning messages generated during the execution.

Run execution - remove_externalidentity ×

 Cancel + delete  Rerun  Refresh

Details Output Error **Warning** Information

```
Task 1 finished w/ error: 1046 @id=1; target=ClusterComputeResource:domain  
Task 2 finished w/ error: 1046 @id=2; target=ClusterComputeResource:domain  
Task 3 finished w/ error: 1046 @id=3; target=ClusterComputeResource:domain
```

Close

- **Information** - Progress and diagnostic generated messages during the execution of a cmdlet.

Run execution - remove_externalidentity ×

 Cancel + delete  Rerun  Refresh

Details Output Error Warning Information

```
SUCCESS: Check for powershell version passed
SUCCESS: VCenter Address is set.
SUCCESS: Check for CloudAdmin role in vCenter passed
SUCCESS: Check for module VMware.vSphere.SsoAdmin passed
SUCCESS: Check for module VMware.VimAutomation.Core passed
Invoke-PreflightJetDRSystemCheck option only checks required configuration
### get_system_state_install ###
SUCCESS: Cluster 'Cluster-1' provided for protection exists in the Datacenter
SUCCESS: Protected cluster satisfies the number of host requirement
```

Close

Cancel or delete a job

Method 1

This method attempts to cancel the execution, and then deletes it upon completion.

IMPORTANT

Method 1 is irreversible.

1. Select **Run command** > **Run execution status** and then select the job you want to cancel.

Run execution - remove_externalidentity ×

 Cancel + delete  Rerun  Refresh

Details Output Error Warning Information

Name
remove_externalidentity

Status
Succeeded

Package
Microsoft.AVS.Management

Command
Remove-ExternalIdentitySources

Close

2. Select **Yes** to cancel and remove the job for all users.

Method 2

1. Select **Run command** > **Packages** > **Run execution status**.

2. Select **More (...)** for the job you want to cancel and delete.

The screenshot shows the Microsoft Azure portal interface. The top navigation bar includes the Microsoft Azure logo and a search bar. The main content area is titled 'Contoso-westus-sddc | Run command'. The left-hand navigation pane is expanded to show 'Run command' under the 'Operations' section. The main area displays a table of 'Run execution status' for various packages. The table has columns for Execution name, Package name, Packa..., Command name, Started time..., End time stamp, and Status. A dropdown menu is open for the 'del_jetdr_3' job, showing options 'Cancel + delete' and 'Rerun'. The 'Run command' option in the left navigation pane is highlighted with a red box.

Execution name	Package name	Packa...	Command name	Started time...	End time stamp	Status
remove_externalidentity	Microsoft.AVS.Management	1.0.30	Remove-ExternalId	7/20/2021, 12:58:4	7/20/2021, 12:59:3	✔ Succeeded
removeGroup	Microsoft.AVS.Management	1.0.30	Remove-GroupFroi	7/20/2021, 12:52:0	7/20/2021, 12:53:4	✔ Succeeded
addADgroup	Microsoft.AVS.Management	1.0.30	Add-GroupToClou	7/20/2021, 12:08:4	7/20/2021, 12:09:2	✔ Succeeded
addexternalidentity	Microsoft.AVS.Management	1.0.30	New-AvsLDAPIden	7/20/2021, 11:13:2	7/20/2021, 11:14:0	✔ Succeeded
getidentitysource	Microsoft.AVS.Management	1.0.30	Get-ExternalIdent	7/20/2021, 10:40:3	7/20/2021, 10:45:3	✔ Succeeded
check_jetserverdetails	JSDR.Configuration	1.0.20	Invoke-PreflightJet	7/20/2021, 7:52:56	7/20/2021, 8:04:57	✘ Failed
checkDRsystem	JSDR.Configuration	1.0.20	Invoke-PreflightJet	7/20/2021, 5:56:33	7/20/2021, 5:57:54	✔ Succeeded
amaneja-jsdrcheck	JSDR.Configuration	1.0.20	Invoke-PreflightJet	7/19/2021, 2:11:41	7/19/2021, 2:13:09	✔ Succeeded
installJSDR_withDNSRegistered	JSDR.Configuration	1.0.20	Install-JetDR	7/16/2021, 8:26:40	7/16/2021, 8:33:09	✘ Failed
del_jetdr_4	JSDR.Configuration	1.0.20	Uninstall-JetDR	7/16/2021, 8:17:22	7/16/2021, 8:18:06	✘ Failed
del_jetdr_3	JSDR.Configuration	1.0.20	Uninstall-JetDR	7/16/2021, 8:07:16	7/16/2021, 8:08:26	✔ Cancel + delete
checksystem	JSDR.Configuration	1.0.20	Invoke-PreflightJet	7/16/2021, 8:04:24	7/16/2021, 8:05:58	✔ Rerun
del_jetdr_2	JSDR.Configuration	1.0.20	Uninstall-JetDR	7/16/2021, 7:53:16	7/16/2021, 7:59:06	✘ Failed

3. Select **Yes** to cancel and remove the job for all users.

Next steps

Now that you've learned about the Run command concepts, you can use the Run command feature to:

- [Configure storage policy](#) - Each VM deployed to a vSAN datastore is assigned a vSAN storage policy. You can assign a vSAN storage policy in an initial deployment of a VM or when you do other VM operations, such as cloning or migrating.
- [Configure external identity source for vCenter \(Run command\)](#) - Configure Active Directory over LDAP or LDAPS for vCenter Server, which enables the use of an external identity source as an Active Directory. Then, you can add groups from the external identity source to the CloudAdmin role.
- [Deploy disaster recovery using JetStream](#) - Store data directly to a recovery cluster in vSAN. The data gets captured through I/O filters that run within vSphere. The underlying data store can be VMFS, vSAN, vVol, or any HCI platform.

Security recommendations for Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

It's important that proper measures are taken to secure your Azure VMware Solution deployments. Use this information as a high-level guide to achieve your security goals.

General

Use the following guidelines and links for general security recommendations for both Azure VMware Solution and VMware best practices.

RECOMMENDATION	COMMENTS
Review and follow VMware Security Best Practices	It's important to stay updated on Azure security practices and VMware Security Best Practices .
Keep up to date on VMware Security Advisories	Subscribe to VMware notifications in my.vmware.com and regularly review and remediate any VMware Security Advisories .
Enable Microsoft Defender for Cloud	Microsoft Defender for Cloud provides unified security management and advanced threat protection across hybrid cloud workloads.
Follow the Microsoft Security Response Center blog	Microsoft Security Response Center
Review and implement recommendations within the Azure Security Baseline for Azure VMware Solution	Azure security baseline for VMware Solution

Network

The following are network-related security recommendations for Azure VMware Solution.

RECOMMENDATION	COMMENTS
Only allow trusted networks	Only allow access to your environments over ExpressRoute or other secured networks. Avoid exposing your management services like vCenter Server, for example, on the internet.
Use Azure Firewall Premium	If you must expose management services on the internet, use Azure Firewall Premium with both IDPS Alert and Deny mode along with TLS inspection for proactive threat detection.
Deploy and configure Network Security Groups on VNET	Ensure any VNET deployed has Network Security Groups configured to control ingress and egress to your environment.

RECOMMENDATION	COMMENTS
Review and implement recommendations within the Azure security baseline for Azure VMware Solution	Azure security baseline for Azure VMware Solution

HCX

See the following information for recommendations to secure your HCX deployment.

RECOMMENDATION	COMMENTS
Stay current with HCX service updates	HCX service updates can include new features, software fixes, and security patches. Apply service updates during a maintenance window where no new HCX operations are queued up by following these steps .

Azure VMware Solution storage concepts

12/16/2022 • 4 minutes to read • [Edit Online](#)

Azure VMware Solution private clouds provide native, cluster-wide storage with VMware vSAN. Local storage from each host in a cluster is used in a vSAN datastore, and data-at-rest encryption is available and enabled by default. You can use Azure Storage resources to extend storage capabilities of your private clouds.

vSAN clusters

Local storage in each cluster host is claimed as part of a vSAN datastore. For the AV36 SKU, all diskgroups use an NVMe cache tier of 1.6 TB with the raw, per host, SSD-based capacity of 15.4 TB. The size of the raw capacity tier of a cluster is the per host capacity times the number of hosts. For example, a four host cluster provides 61.6-TB raw capacity in the vSAN capacity tier. Check the hardware specification for the [AV36P and AV52 SKU](#) storage device details.

Local storage in cluster hosts is used in the cluster-wide vSAN datastore. All datastores are created as part of private cloud deployment and are available for use immediately. The **cloudadmin** user and all users assigned to the CloudAdmin role can manage datastores with these vSAN privileges:

- Datastore.AllocateSpace
- Datastore.Browse
- Datastore.Config
- Datastore.DeleteFile
- Datastore.FileManagement
- Datastore.UpdateVirtualMachineMetadata

IMPORTANT

You can't change the name of datastores or clusters. Azure CLI and PowerShell support changing the name of the resource clusters (Cluster-2 to Cluster-12), however this should not be used, because it creates a meta-data mismatch between the Azure portal resource cluster name and the vSphere cluster name.

Storage policies and fault tolerance

The default storage policy is set to **RAID-1 FTT-1**, with Object Space Reservation set to Thin provisioning. Unless you adjust the storage policy or apply a new policy, the cluster grows with this configuration. This is the policy that will be applied to the workload VMs. To set a different storage policy, see [Configure storage policy](#).

In a three-host cluster, FTT-1 accommodates a single host's failure. Microsoft governs failures regularly and replaces the hardware when events are detected from an operations perspective.

NOTE

When you log on to the vSphere Client, you may notice a VM Storage Policy called **vSAN Default Storage Policy** with **Object Space Reservation** set to **Thick** provisioning. Please note that this is not the default storage policy applied to the cluster. This policy exists for historical purposes and will eventually be modified to **Thin** provisioning.

NOTE

All of the software-defined data center (SDDC) management VMs (vCenter Server, NSX-T Manager, NSX-T Data Center Edges, and others) use the **Microsoft vSAN Management Storage Policy**, with **Object Space Reservation** set to **Thick** provisioning.

TIP

If you're unsure if the cluster will grow to four or more, then deploy using the default policy. If you're sure your cluster will grow, then instead of expanding the cluster after your initial deployment, we recommend deploying the extra hosts during deployment. As the VMs are deployed to the cluster, change the disk's storage policy in the VM settings to either RAID-5 FTT-1 or RAID-6 FTT-2. In reference to [SLA for Azure VMware Solution](#), note that more than 6 hosts should be configured in the cluster to use an FTT-2 policy (RAID-1, or RAID-6). Also note that the storage policy is not automatically updated based on cluster size. Similarly, changing the default does not automatically update the running VM policies.

Data-at-rest encryption

vSAN datastores use data-at-rest encryption by default using keys stored in Azure Key Vault. The encryption solution is KMS-based and supports vCenter Server operations for key management. When a host is removed from a cluster, all data on SSDs is invalidated immediately.

Datastore capacity expansion options

The existing cluster vSAN storage capacity can be expanded by connecting Azure storage resources such as [Azure NetApp Files volumes as additional datastores](#). Virtual machines can be migrated between vSAN and Azure NetApp Files datastores using storage vMotion. Azure NetApp Files is available in [Ultra, Premium and Standard performance tiers](#) to allow for adjusting performance and cost to the requirements of the workloads.

Azure storage integration

You can use Azure storage services in workloads running in your private cloud. The Azure storage services include Storage Accounts, Table Storage, and Blob Storage. The connection of workloads to Azure storage services doesn't traverse the internet. This connectivity provides more security and enables you to use SLA-based Azure storage services in your private cloud workloads.

Alerts and monitoring

Microsoft provides alerts when capacity consumption exceeds 75%. In addition, you can monitor capacity consumption metrics that are integrated into Azure Monitor. For more information, see [Configure Azure Alerts in Azure VMware Solution](#).

Next steps

Now that you've covered Azure VMware Solution storage concepts, you may want to learn about:

- [Attach disk pools to Azure VMware Solution hosts \(Preview\)](#) - You can use disks as the persistent storage for Azure VMware Solution for optimal cost and performance.
- [Configure storage policy](#) - Each VM deployed to a vSAN datastore is assigned at least one VM storage policy. You can assign a VM storage policy in an initial deployment of a VM or when you perform other VM operations, such as cloning or migrating.
- [Scale clusters in the private cloud](#) - You can scale the clusters and hosts in a private cloud as required for

your application workload. Performance and availability limitations for specific services should be addressed on a case by case basis.

- [Azure NetApp Files with Azure VMware Solution](#) - You can use Azure NetApp Files to migrate and run the most demanding enterprise file-workloads in the cloud: databases, and general purpose computing applications, with no code changes. Azure NetApp Files volumes can be attached to virtual machines, and as [datastores](#) to extend the vSAN datastore capacity without adding more nodes.
- [vSphere role-based access control for Azure VMware Solution](#) - You use vCenter Server to manage VM workloads and NSX-T Manager to manage and extend the private cloud. Access and identity management use the CloudAdmin role for vCenter Server and restricted administrator rights for NSX-T Manager.

Monitor and protect VMs with Azure native services

12/16/2022 • 3 minutes to read • [Edit Online](#)

Microsoft Azure native services let you monitor, manage, and protect your virtual machines (VMs) in a hybrid environment (Azure, Azure VMware Solution, and on-premises). In this article, you'll integrate Azure native services in your Azure VMware Solution private cloud. You'll also learn how to use the tools to manage your VMs throughout their lifecycle.

The Azure native services that you can integrate with Azure VMware Solution include:

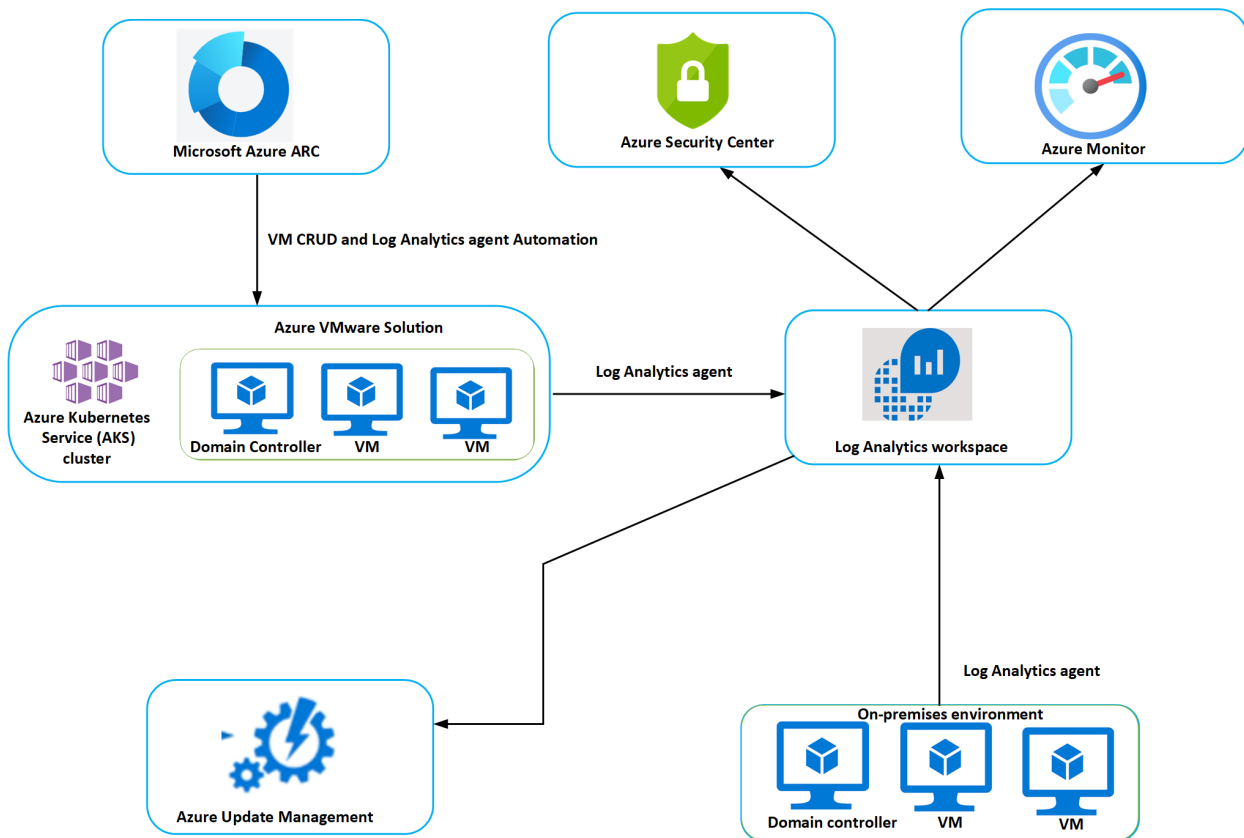
- Azure Arc extends Azure management Azure VMware Solution. After your Azure VMware Solution private cloud is deployed to Arc, you'll be ready to execute operations in Azure VMware Solution vCenter Server from the Azure portal. Operations are related to Create, Read, Update, and Delete (CRUD) virtual machines (VMs) in an Arc-enabled Azure VMware Solution private cloud. Users can also enable guest management and install Azure extensions after the private cloud is Arc-enabled.
- Azure Monitor collects, analyzes, and acts on data from your cloud and on-premises environments. Your Log Analytics workspace in Azure Monitor enables log collection and performance counter collection using the Log Analytics agent or extensions. You can send logs from your Azure VMware Solution private cloud to your Log Analytics workspace, allowing you to take advantage of the Log Analytics feature set, including:
 - system patches, security misconfigurations, and endpoint protection. You can also define security policies in Microsoft Defender for Cloud.
- Log Analytics workspace stores log data. Each workspace has its own data repository and configuration to store data. You can monitor Azure VMware Solution VMs through the Log Analytics agent. Machines connected to the Log Analytics Workspace use the Log Analytics agent to collect data about changes to installed software, Microsoft services, Windows registry and files, and Linux daemons on monitored servers. When data is available, the agent sends it to Azure Monitor Logs for processing. Azure Monitor Logs applies logic to the received data, records it, and makes it available for analysis.

Benefits

- Azure native services can be used to manage your VMs in a hybrid environment (Azure, Azure VMware Solution, and on-premises).
- Integrated monitoring and visibility of your Azure, Azure VMware Solution, and on-premises VMs.
 - Fileless security alerts
 - Operating system patch assessment
 - Security misconfigurations assessment
 - Endpoint protection assessment
- Easily deploy the Log Analytics extension after enabling guest management on VMware vSphere virtual machine (VM).
- Your Log Analytics workspace in Azure Monitor enables log collection and performance counter collection using the Log Analytics extensions. Collect data and logs to a single point and present that data to different Azure native services.
- Added benefits of Azure Monitor include:
 - Seamless monitoring
 - Better infrastructure visibility
 - Instant notifications
 - Automatic resolution
 - Cost efficiency

Topology

The diagram shows the integrated monitoring architecture for Azure VMware Solution VMs.



NOTE

If you're new to Azure or not familiar with the services previously mentioned, see [Enable Azure Monitor for VMs overview](#) for guidance.

Enable guest management and install extension

The guest management must be enabled on the VMware vSphere virtual machine (VM) before you can install an extension. Use the following prerequisite steps to enable guest management.

Prerequisites

- Navigate to Azure portal.
- Locate the VMware vSphere VM you want to check for guest management and install extensions on, select the name of the VM.
- Select **Configuration** from the left navigation for a VMware VM.
- Verify **Enable guest management** has been checked.

The following conditions are necessary to enable guest management on a VM.

- The machine must be running a supported operating system.
- The machine needs to connect through the firewall to communicate over the internet. Make sure the URLs listed aren't blocked.
- The machine can't be behind a proxy, it's not supported yet.
- If you're using Linux VM, the account must not prompt to sign in on pseudo commands.
- To avoid pseudo commands, follow these steps:

1. Sign into Linux VM.
2. Open terminal and run the following command: `sudo visudo`.
3. Add the line `username ALL=(ALL) NOPASSWD: ALL` at the end of the file.
4. Replace username with the appropriate user-name. If your VM template already has these changes incorporated, you won't need to do the steps for the VM created from that template.

Install extensions

1. Go to **Azure** portal.
2. Find the Arc-enabled Azure VMware Solution VM that you want to install an extension on and select the VM name.
3. Navigate to **Extensions** in the left navigation, select **Add**.
4. Select the extension you want to install.
Based on the extension, you'll need to provide details.
For example, workspace ID and key for Log Analytics extension.
5. When you're done, select **Review + create**.

When the extension installation steps are completed, they trigger deployment and install the selected extension on the VM.

Next steps

Now that you've covered how to integrate services and monitor VMware Solution VMs, you may want to learn about:

- [Using the workload protection dashboard](#)
- [Advanced multistage attack detection in Microsoft Sentinel](#)

Integrate Microsoft Defender for Cloud with Azure VMware Solution

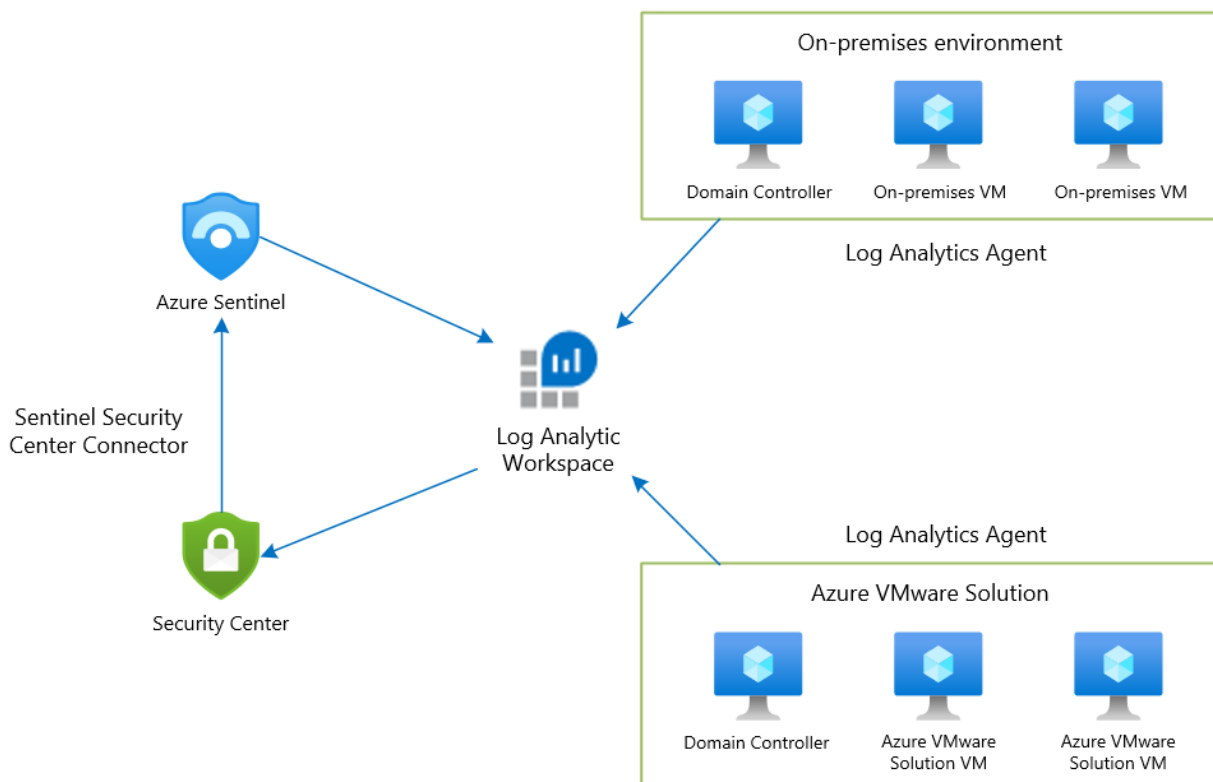
12/16/2022 • 4 minutes to read • [Edit Online](#)

Microsoft Defender for Cloud provides advanced threat protection across your Azure VMware Solution and on-premises virtual machines (VMs). It assesses the vulnerability of Azure VMware Solution VMs and raises alerts as needed. These security alerts can be forwarded to Azure Monitor for resolution. You can define security policies in Microsoft Defender for Cloud. For more information, see [Working with security policies](#).

Microsoft Defender for Cloud offers many features, including:

- File integrity monitoring
- Fileless attack detection
- Operating system patch assessment
- Security misconfigurations assessment
- Endpoint protection assessment

The diagram shows the integrated monitoring architecture of integrated security for Azure VMware Solution VMs.



Log Analytics agent collects log data from Azure, Azure VMware Solution, and on-premises VMs. The log data is sent to Azure Monitor Logs and stored in a **Log Analytics Workspace**. Each workspace has its own data repository and configuration to store data. Once the logs are collected, **Microsoft Defender for Cloud** assesses the vulnerability status of Azure VMware Solution VMs and raises an alert for any critical vulnerability. Once assessed, Microsoft Defender for Cloud forwards the vulnerability status to Microsoft Sentinel to create an incident and map with other threats. Microsoft Defender for Cloud is connected to Microsoft Sentinel using Microsoft Defender for Cloud Connector.

Prerequisites

- [Plan for optimized use of Defender for Cloud.](#)
- [Review the supported platforms in Defender for Cloud.](#)
- [Create a Log Analytics workspace](#) to collect data from various sources.
- [Enable Microsoft Defender for Cloud in your subscription.](#)


NOTE

Microsoft Defender for Cloud is a pre-configured tool that doesn't require deployment, but you'll need to enable it.

- [Enable Microsoft Defender for Cloud.](#)

Add Azure VMware Solution VMs to Defender for Cloud

1. In the Azure portal, search on **Azure Arc** and select it.
2. Under Resources, select **Servers** and then **+ Add**.



3. Select **Generate script**.

Home > Azure Arc >

Select a method

To connect servers (both from on-premises and other clouds) to Azure, deploy the Azure Connected Machine agent to your servers. Select from one of the options below to onboard your servers. [Learn more](#)

Add servers using interactive script

Generate a script to onboard the target server. Use this option to run the script which will prompt for your Azure login during deployment time.

[Generate script](#)

[Learn more](#)

Add servers at scale

If you are running the deployment at scale, you will need to provide a Service Principal Name with the minimum set of Azure permissions to onboard your servers.

[View instructions](#)

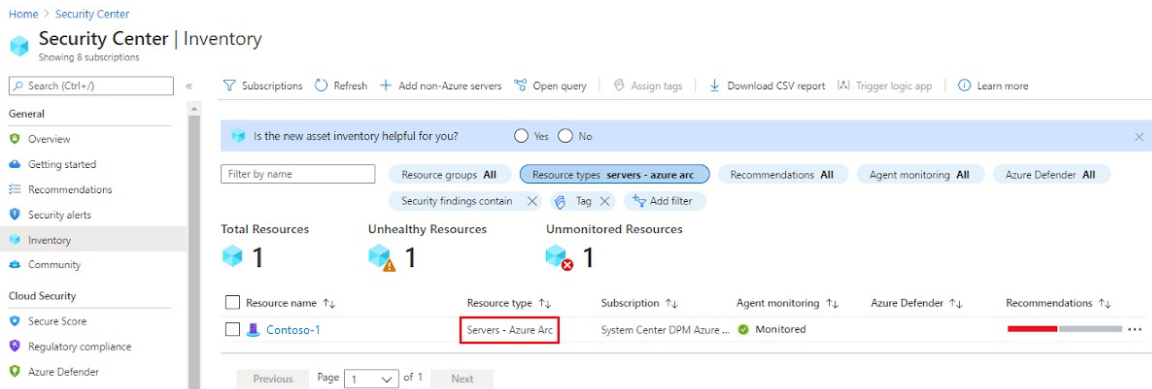
4. On the **Prerequisites** tab, select **Next**.
5. On the **Resource details** tab, fill in the following details and then select **Next**. **Tags**:
 - Subscription
 - Resource group
 - Region
 - Operating system
 - Proxy Server details

- On the **Tags** tab, select **Next**.
- On the **Download and run script** tab, select **Download**.
- Specify your operating system and run the script on your Azure VMware Solution VM.

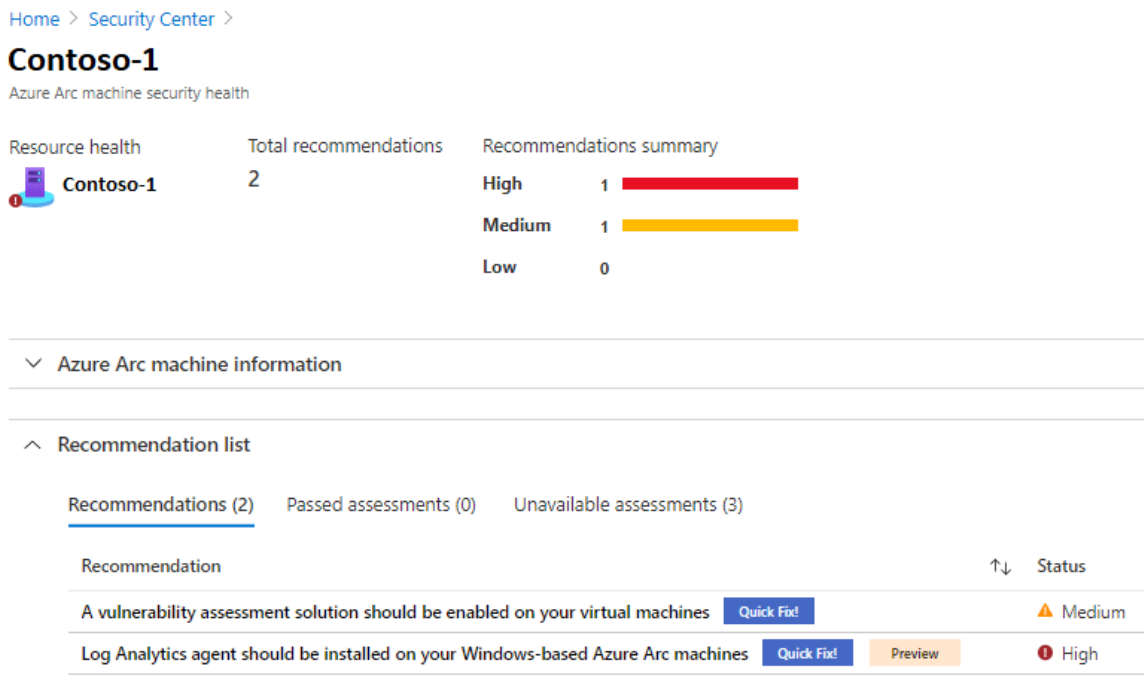
View recommendations and passed assessments

Recommendations and assessments provide you with the security health details of your resource.

- In Microsoft Defender for Cloud, select **Inventory** from the left pane.
- For Resource type, select **Servers - Azure Arc**.



- Select the name of your resource. A page opens showing the security health details of your resource.
- Under **Recommendation list**, select the **Recommendations**, **Passed assessments**, and **Unavailable assessments** tabs to view these details.



Deploy a Microsoft Sentinel workspace

Microsoft Sentinel provides security analytics, alert detection, and automated threat response across an environment. It's a cloud-native, security information event management (SIEM) solution that's built on top of a Log Analytics workspace.

Since Microsoft Sentinel is built on top of a Log Analytics workspace, you'll only need to select the workspace

you want to use.

1. In the Azure portal, search for **Microsoft Sentinel**, and select it.
2. On the Microsoft Sentinel workspaces page, select **+ Add**.
3. Select the Log Analytics workspace and select **Add**.

Enable data collector for security events

1. On the Microsoft Sentinel workspaces page, select the configured workspace.
2. Under Configuration, select **Data connectors**.
3. Under the Connector Name column, select **Security Events** from the list, then select **Open connector page**.
4. On the connector page, select the events you wish to stream, then select **Apply Changes**.

The screenshot shows the 'Security Events' configuration page in the Azure portal. On the left, there's a summary card showing 'Connected Status' as 'Microsoft Provider' with a last log received '5 minutes ago'. It also displays '5 Workbooks', '1 Queries', and '31 Analytic rules templates'. Below this is a line graph titled 'Data received' showing data volume over time, with a total of '105.58k' data types. The main content area is titled 'Instructions' and 'Next steps'. Step 1 is 'Download and install the agent', noting that logs are collected from Windows agents. It offers two installation options: 'Install agent on Azure Windows Virtual Machine' and 'Install agent on non-Azure Windows Machine'. Step 2 is 'Select which events to stream', with options: 'All events - All Windows security and AppLocker events.', 'Common - A standard set of events for auditing purposes.', 'Minimal - A small set of events that might indicate potential threats.', and 'None - No security or AppLocker events.'. The 'All Events' radio button is selected and highlighted with a red box. An 'Apply Changes' button is at the bottom.

Connect Microsoft Sentinel with Microsoft Defender for Cloud

1. On the Microsoft Sentinel workspace page, select the configured workspace.
2. Under Configuration, select **Data connectors**.
3. Select **Microsoft Defender for Cloud** from the list, then select **Open connector page**.

The screenshot shows the 'Data connectors' page in the Azure Sentinel portal. The page title is 'Azure Sentinel | Data connectors' for the workspace 'logupdatemanagement'. It shows a search bar and filters for 'Providers: All', 'Data Types: All', and 'Status: All'. A table lists connectors with columns for 'Status', 'Connector name', and 'Provider'. The 'Azure Security Center' connector is highlighted. To the right, a detailed view of the 'Azure Security Center' connector is shown, including its status ('Connected'), provider ('Microsoft'), and last log received ('23 hours ...'). A description explains that Azure Security Center is a security management tool that provides insight into security state across hybrid cloud workloads. An 'Open connector page' button is visible at the bottom of the details pane.

4. Select **Connect** to connect the Microsoft Defender for Cloud with Microsoft Sentinel.
5. Enable **Create incident** to generate an incident for Microsoft Defender for Cloud.

Create rules to identify security threats

After connecting data sources to Microsoft Sentinel, you can create rules to generate alerts for detected threats. In the following example, we'll create a rule for attempts to sign into Windows server with the wrong password.

1. On the Microsoft Sentinel overview page, under Configurations, select **Analytics**.
2. Under Configurations, select **Analytics**.
3. Select **+Create** and on the drop-down, select **Scheduled query rule**.
4. On the **General** tab, enter the required information and then select **Next: Set rule logic**.
 - Name
 - Description
 - Tactics
 - Severity
 - Status
5. On the **Set rule logic** tab, enter the required information, then select **Next**.
 - Rule query (here showing our example query)

```
SecurityEvent
|where Activity startswith '4625'
|summarize count ( ) by IPAddress,Computer
|where count_ > 3
```

- Map entities
 - Query scheduling
 - Alert threshold
 - Event grouping
 - Suppression
6. On the **Incident settings** tab, enable **Create incidents from alerts triggered by this analytics rule** and select **Next: Automated response**.

Analytics rule wizard - Create new rule



- General**
- Set rule logic
- Incident settings (Preview)
- Automated response
- Review and create

Create an analytics rule that will run on your data to detect threats.

Analytics rule details

Name *

Description

Tactics and techniques

Severity

Status

 Enabled Disabled

Next : Set rule logic >

7. Select **Next: Review**.
8. On the **Review and create** tab, review the information, and select **Create**.

TIP

After the third failed attempt to sign into Windows server, the created rule triggers an incident for every unsuccessful attempt.

View alerts

You can view generated incidents with Microsoft Sentinel. You can also assign incidents and close them once they're resolved, all from within Microsoft Sentinel.

1. Go to the Microsoft Sentinel overview page.
2. Under Threat Management, select **Incidents**.
3. Select an incident and then assign it to a team for resolution.

Azure Sentinel | Incidents

Selected workspace: 'logupdatemanagement'

Search (Ctrl+/) Refresh Last 30 days Actions Security efficiency workbook (Preview)

General

- Overview
- Logs
- News & guides

Threat management

- Incidents
- Workbooks
- Hunting
- Notebooks (Preview)
- Entity behavior
- Threat intelligence (Preview)

Configuration

- Data connectors

60 Open incidents 60 New incidents 0 Active incidents

Open incidents by severity: High (0) | Medium (60) | Low (0) | Informational (0)

Search by id or title

Severity: All Status: New, Active Product name: All Owner: All

Auto-refresh incidents

Incident id	Title	Alerts
119	Login Failure	1
118	Login Failure	1
117	Login Failure	1
116	Login Failure	1
115	Login Failure	1

Search users

- Assign to me amb@microsoft.com
- adam adam@ipvm.com

Apply Cancel

< Previous 1 - 50 Next >

TIP

After resolving the issue, you can close it.

Hunt security threats with queries

You can create queries or use the available pre-defined query in Microsoft Sentinel to identify threats in your environment. The following steps run a pre-defined query.

1. On the Microsoft Sentinel overview page, under Threat management, select **Hunting**. A list of pre-defined queries is displayed.

TIP

You can also create a new query by selecting **New Query**.

Home > Microsoft Sentinel

Microsoft Sentinel | Hunting

Selected workspace: 'CyberSecuritySOC'

Search (Ctrl+/) Refresh Last 24 hours + New Query Run all queries (Preview) Columns

General

- Overview
- Logs
- News & guides

Threat management

- Incidents
- Workbooks
- Hunting
- Notebooks
- Entity behavior
- Threat intelligence (Preview)

224 / 249 Active / total queries 0 / 0 Result count / queries run 0 Livestream Results 0 My bookmarks

Queries Livestream Bookmarks

Query	Provider	Data Source	Results	Results delta (Pre...)
Changes made to AWS ...	Microsoft	AWSCloudTrail	--	N/A
Consent to Application ...	Microsoft	AuditLogs +1	--	N/A

Search queries Favorites: All Provider: All Data sources: All Tactics: All

2. Select a query and then select **Run Query**.

3. Select **View Results** to check the results.

Next steps

Now that you've covered how to protect your Azure VMware Solution VMs, you may want to learn about:

- [Using the workload protection dashboard](#)
- [Advanced multistage attack detection in Microsoft Sentinel](#)
- [Integrating Azure native services in Azure VMware Solution](#)

Protect web apps on Azure VMware Solution with Azure Application Gateway

12/16/2022 • 5 minutes to read • [Edit Online](#)

[Azure Application Gateway](#) is a layer 7 web traffic load balancer that lets you manage traffic to your web applications. It's offered in both Azure VMware Solution v1.0 and v2.0. Both versions tested with web apps running on Azure VMware Solution.

The capabilities include:

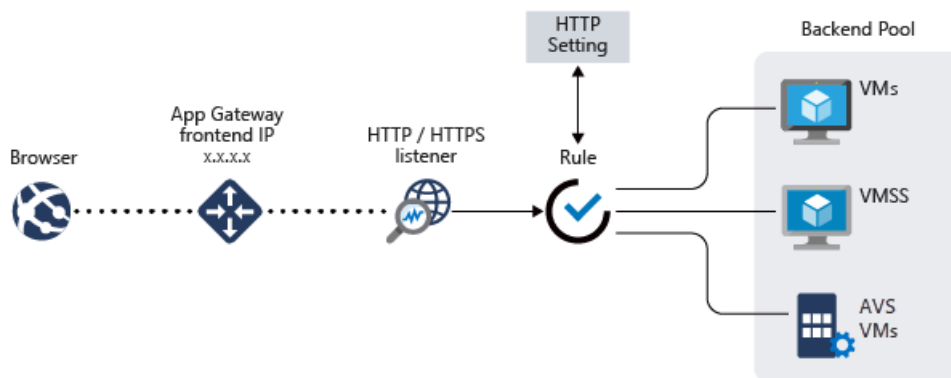
- Cookie-based session affinity
- URL-based routing
- Web Application Firewall (WAF)

For a complete list of features, see [Azure Application Gateway features](#).

This article shows you how to use Application Gateway in front of a web server farm to protect a web app running on Azure VMware Solution.

Topology


The diagram shows how Application Gateway is used to protect Azure IaaS virtual machines (VMs), Azure Virtual Machine Scale Sets, or on-premises servers. Application Gateway treats Azure VMware Solution VMs as on-premises servers.



IMPORTANT

Azure Application Gateway is currently the only supported method to expose web apps running on Azure VMware Solution VMs.

The diagram shows the testing scenario used to validate the Application Gateway with Azure VMware Solution web applications.



The Application Gateway instance gets deployed on the hub in a dedicated subnet with an Azure public IP address. Activating the [Azure DDoS Protection](#) for the virtual network is recommended. The web server is hosted on an Azure VMware Solution private cloud behind NSX T0 and T1 Gateways. Additionally, Azure VMware Solution uses [ExpressRoute Global Reach](#) to enable communication with the hub and on-premises systems.

Prerequisites

- An Azure account with an active subscription.
- An Azure VMware Solution private cloud deployed and running.

Deployment and configuration

1. In the Azure portal, search for **Application Gateway** and select **Create application gateway**.
2. Provide the basic details as in the following figure; then select **Next: Frontends >**.



Create application gateway

An application gateway is a web traffic load balancer that enables you to manage traffic to your web application. [Learn more about application gateway](#)

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ

Resource group * ⓘ

[Create new](#)

Instance details

Application gateway name *

Region *

Tier ⓘ

Enable autoscaling

 Yes No

Minimum scale units * ⓘ

Maximum scale units

Availability zone ⓘ

HTTP2 ⓘ

 Disabled Enabled

Configure virtual network

Virtual network * ⓘ

[Create new](#)

Subnet * ⓘ

[Manage subnet configuration](#)

Previous

Next : Frontends >

3. Choose the frontend IP address type. For public, choose an existing public IP address or create a new one. Select **Next: Backends >**.

NOTE

Only standard and Web Application Firewall (WAF) SKUs are supported for private frontends.

4. Add a backend pool of the VMs that run on Azure VMware Solution infrastructure. Provide the details of web servers that run on the Azure VMware Solution private cloud and select **Add**. Then select **Next: Configuration** > .
5. On the **Configuration** tab, select **Add a routing rule**.
6. On the **Listener** tab, provide the details for the listener. If HTTPS is selected, you must provide a certificate, either from a PFX file or an existing Azure Key Vault certificate.
7. Select the **Backend targets** tab and select the backend pool previously created. For the **HTTP settings** field, select **Add new**.
8. Configure the parameters for the HTTP settings. Select **Add**.
9. If you want to configure path-based rules, select **Add multiple targets to create a path-based rule**.
10. Add a path-based rule and select **Add**. Repeat to add more path-based rules.
11. When you have finished adding path-based rules, select **Add** again; then select **Next: Tags** > .
12. Add tags and then select **Next: Review + Create** > .
13. A validation runs on your Application Gateway. If it's successful, select **Create** to deploy.

Configuration examples

Now you'll configure Application Gateway with Azure VMware Solution VMs as backend pools for the following use cases:

- [Hosting multiple sites](#)
- [Routing by URL](#)

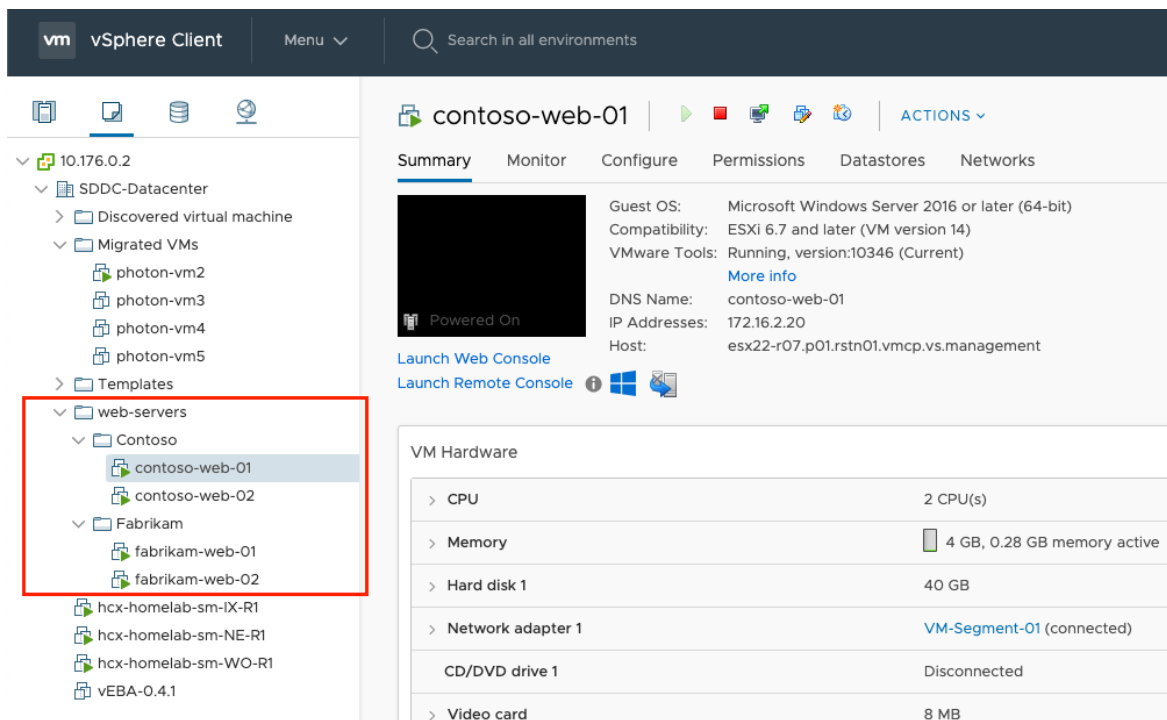
Hosting multiple sites

This procedure shows you how to define backend address pools using VMs running on an Azure VMware Solution private cloud on an existing application gateway.

NOTE

This procedure assumes you have multiple domains, so we'll use examples of www.contoso.com and www.fabrikam.com.

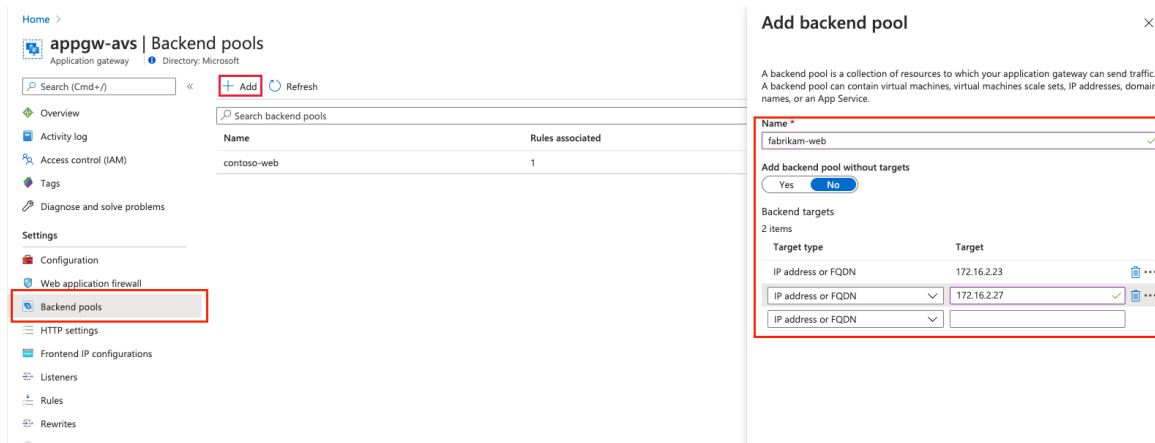
1. In your private cloud, create two different pools of VMs. One represents Contoso and the second Fabrikam.



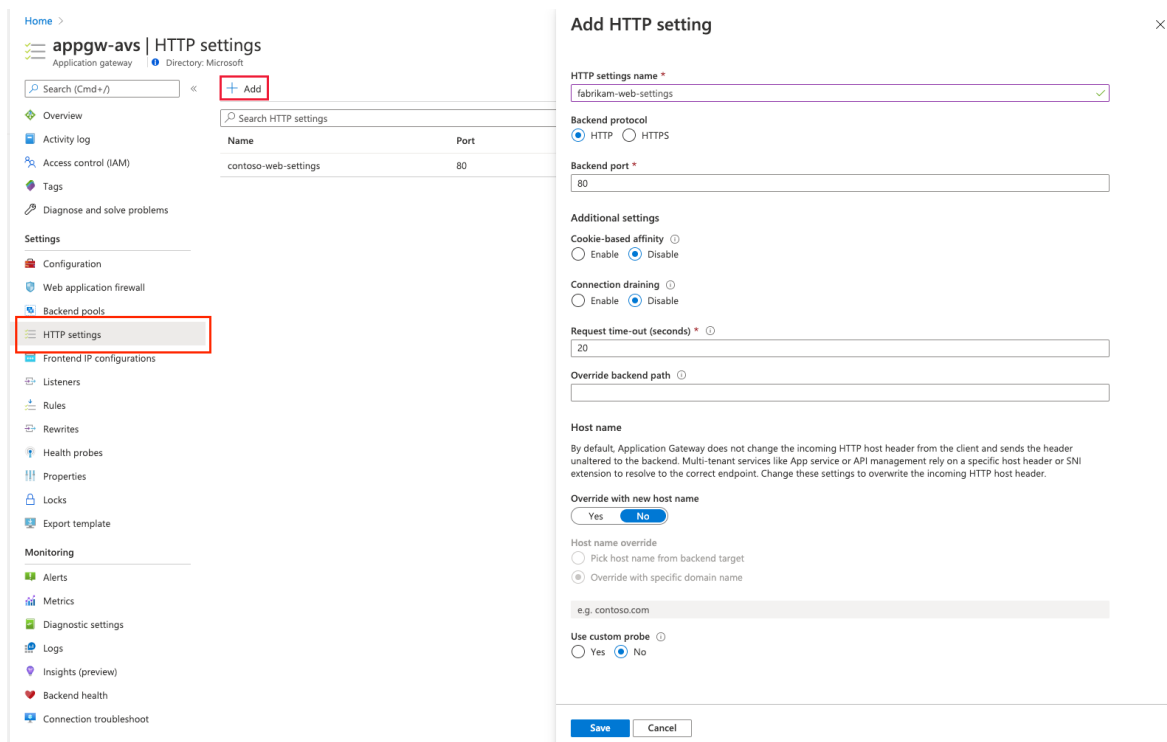
We've used Windows Server 2016 with the Internet Information Services (IIS) role installed. Once the VMs are installed, run the following PowerShell commands to configure IIS on each of the VMs.

```
Install-WindowsFeature -Name Web-Server
Add-Content -Path C:\inetpub\wwwroot\Default.htm -Value $($env:computername)
```


2. In an existing application gateway instance, select **Backend pools** from the left menu, select **Add**, and enter the new pools' details. Select **Add** in the right pane.



3. In the **Listeners** section, create a new listener for each website. Enter the details for each listener and select **Add**.
4. On the left, select **HTTP settings** and select **Add** in the left pane. Fill in the details to create a new HTTP setting and select **Save**.



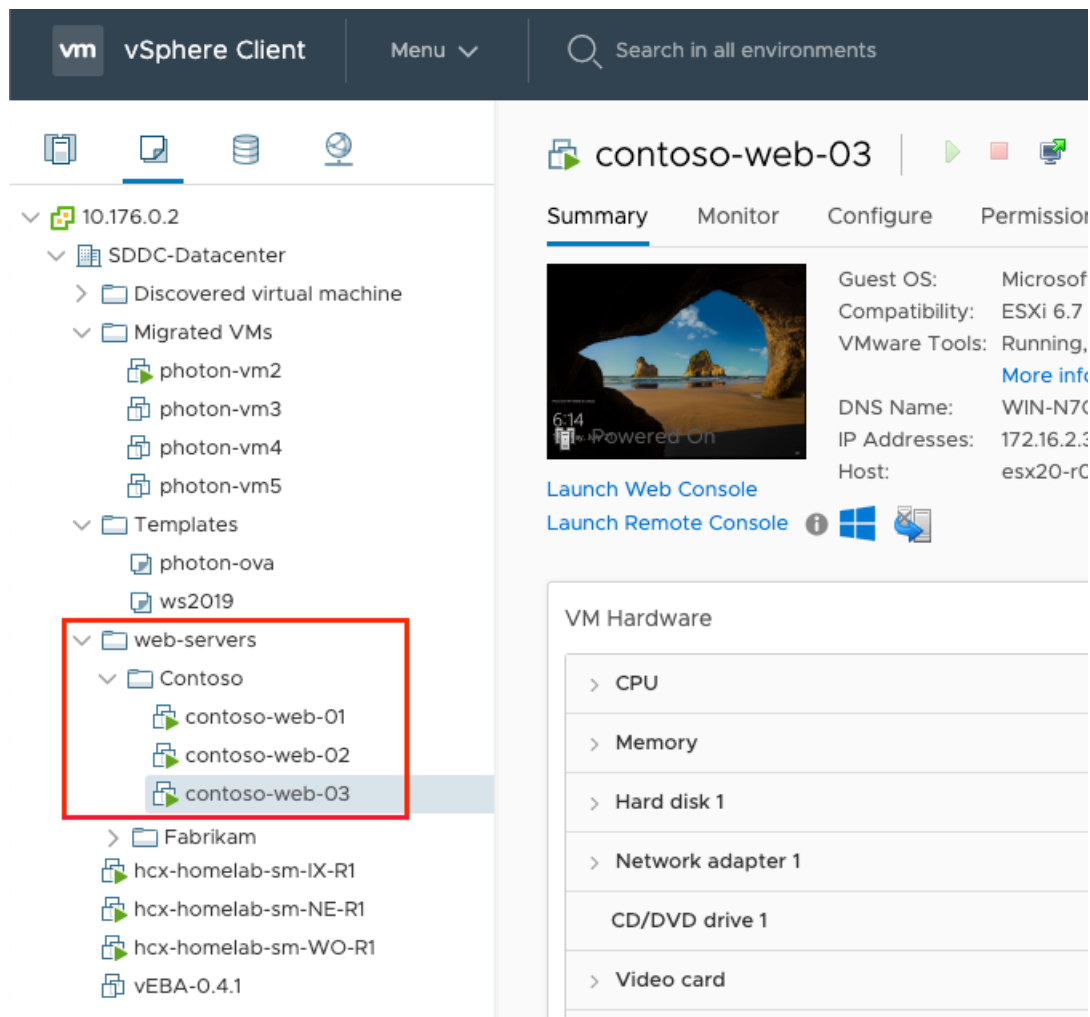
5. Create the rules in the **Rules** section of the left menu. Associate each rule with the corresponding listener. Select **Add**.
6. Configure the corresponding backend pool and HTTP settings. Select **Add**.
7. Test the connection. Open your preferred browser and navigate to the different websites hosted on your Azure VMware Solution environment, for example, <http://www.fabrikam.com>.



Routing by URL

The following steps define backend address pools using VMs running on an Azure VMware Solution private cloud. The private cloud is on an existing application gateway. You then create routing rules that make sure web traffic arrives at the appropriate servers in the pools.

1. In your private cloud, create a virtual machine pool to represent the web farm.



Windows Server 2016 with IIS role installed has been used to illustrate this tutorial. Once the VMs are installed, run the following PowerShell commands to configure IIS for each VM tutorial.

The first virtual machine, contoso-web-01, hosts the main website.

```
Install-WindowsFeature -Name Web-Server
Add-Content -Path C:\inetpub\wwwroot\Default.htm -Value $($env:computername)
```

The second virtual machine, contoso-web-02, hosts the images site.

```
Install-WindowsFeature -Name Web-Server
New-Item -Path "C:\inetpub\wwwroot\" -Name "images" -ItemType "directory"
Add-Content -Path C:\inetpub\wwwroot\images\test.htm -Value $($env:computername)
```

The third virtual machine, contoso-web-03, hosts the video site.

```
Install-WindowsFeature -Name Web-Server
New-Item -Path "C:\inetpub\wwwroot\" -Name "video" -ItemType "directory"
Add-Content -Path C:\inetpub\wwwroot\video\test.htm -Value $($env:computername)
```

2. Add three new backend pools in an existing application gateway instance.

- a. Select **Backend pools** from the left menu.
- b. Select **Add** and enter the details of the first pool, **contoso-web**.
- c. Add one VM as the target.
- d. Select **Add**.

e. Repeat this process for **contoso-images** and **contoso-video**, adding one unique VM as the target.



3. In the **Listeners** section, create a new listener of type **Basic** using port **8080**.

4. On the left navigation, select **HTTP settings** and select **Add** in the left pane. Fill in the details to create a new HTTP setting and select **Save**.

Add HTTP setting

HTTP settings name *

 ✓

Backend protocol

HTTP HTTPS

Backend port *

Additional settings

Cookie-based affinity ⓘ

Enable Disable

Connection draining ⓘ

Enable Disable

Request time-out (seconds) * ⓘ

Override backend path ⓘ

Host name

By default, Application Gateway does not change the incoming HTTP host header from the client and sends the header unaltered to the backend. Multi-tenant services like App service or API management rely on a specific host header or SNI extension to resolve to the correct endpoint. Change these settings to overwrite the incoming HTTP host header.

Override with new host name

Yes No

Host name override

Pick host name from backend target

Override with specific domain name

e.g. contoso.com

Use custom probe ⓘ

Yes No

5. Create the rules in the **Rules** section of the left menu and associate each rule with the previously created listener. Then configure the main backend pool and HTTP settings, and then select **Add**.

Add a routing rule



appgw-avs

Configure a routing rule to send traffic from a given frontend IP address to one or more backend targets. A routing rule must contain a listener and at least one backend target.

Rule name *

* Listener * Backend targets

Choose a backend pool to which this routing rule will send traffic. You will also need to specify a set of HTTP settings that define the behavior of the routing rule.

Target type Backend pool Redirection

Backend target *

HTTP settings *

Path-based routing

You can route traffic from this rule's listener to different backend targets based on the URL path of the request. You can also apply a different set of HTTP settings based on the URL path.


Path based rules

Path	Target name	HTTP setting name	Backend pool	
/images/*	contoso-images	contoso-web-setting	contoso-images	***
/video/*	contoso-video	contoso-web-setting	contoso-video	***


[Add multiple targets to create a path-based rule](#)

6. Test the configuration. Access the application gateway on the Azure portal and copy the public IP address in the **Overview** section.


a. Open a new browser window and enter the URL .



b. Change the URL to .



c. Change the URL again to .



CONTOSO-WEB-03

Next Steps

Now that you've covered using Application Gateway to protect a web app running on Azure VMware Solution, you may want to learn about:

- [Configuring Azure Application Gateway for different scenarios.](#)
- [Deploying Traffic Manager to balance Azure VMware Solution workloads.](#)
- [Integrating Azure NetApp Files with Azure VMware Solution-based workloads.](#)
- [Protecting Azure resources in virtual networks.](#)

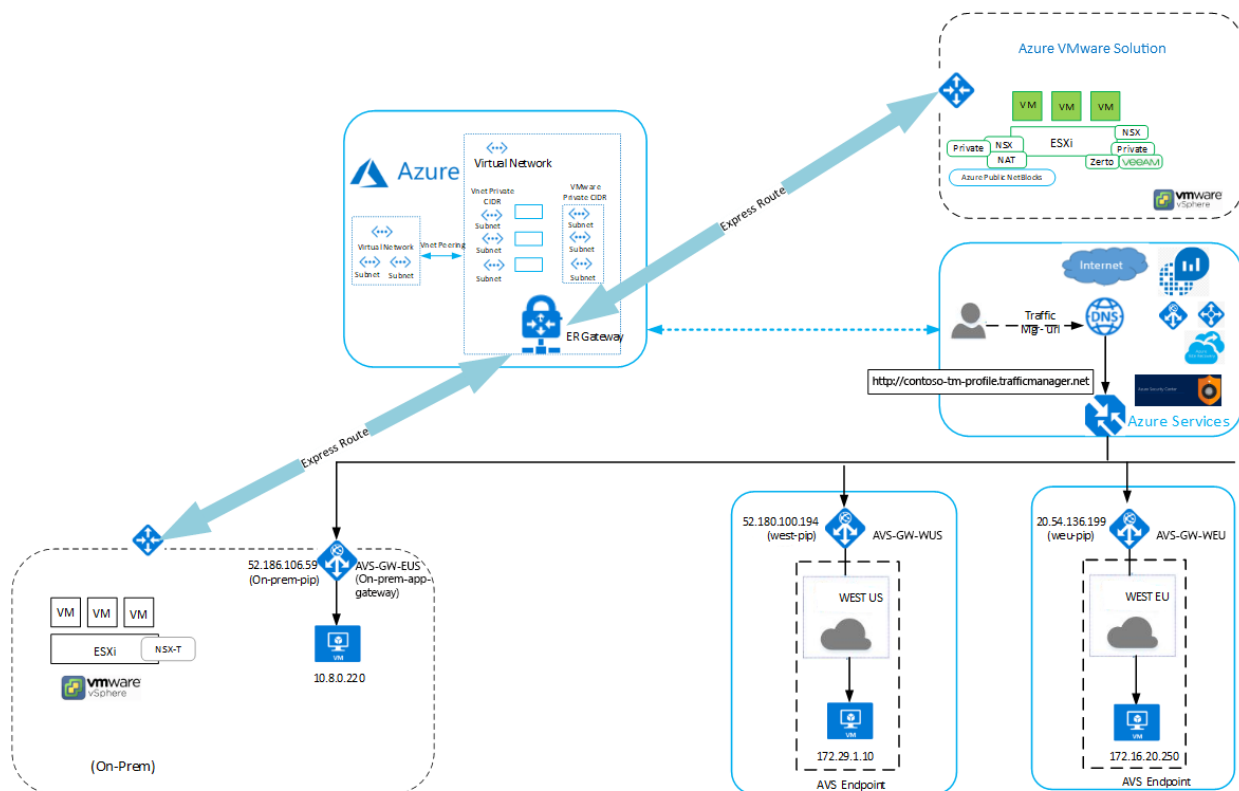
Deploy Azure Traffic Manager to balance Azure VMware Solution workloads

12/16/2022 • 3 minutes to read • [Edit Online](#)

This article walks through the steps of how to integrate [Azure Traffic Manager](#) with Azure VMware Solution. The integration balances application workloads across multiple endpoints. This article also walks through the steps of how to configure Traffic Manager to direct traffic between three [Azure Application Gateway](#) spanning several Azure VMware Solution regions.

The gateways have Azure VMware Solution virtual machines (VMs) configured as backend pool members to load balance the incoming layer 7 requests. For more information, see [Use Azure Application Gateway to protect your web apps on Azure VMware Solution](#)

The diagram shows how Traffic Manager provides load balancing for the applications at the DNS level between regional endpoints. The gateways have backend pool members configured as IIS Servers and referenced as Azure VMware Solution external endpoints. Connection over the virtual network between the two private cloud regions uses an ExpressRoute gateway.



Before you begin, first review the [Prerequisites](#) and then we'll walk through the procedures to:

- Verify configuration of your application gateways and the NSX-T segment
- Create your Traffic Manager profile
- Add external endpoints into your Traffic Manager profile

Prerequisites

- Three VMs configured as Microsoft IIS Servers running in different Azure VMware Solution regions:
 - West US


- West Europe
- East US (on-premises)
- An application gateway with external endpoints in the Azure VMware Solution regions mentioned above.
- Host with internet connectivity for verification.
- An [NSX-T network segment created in Azure VMware Solution](#).

Verify your application gateways configuration

The following steps verify the configuration of your application gateways.

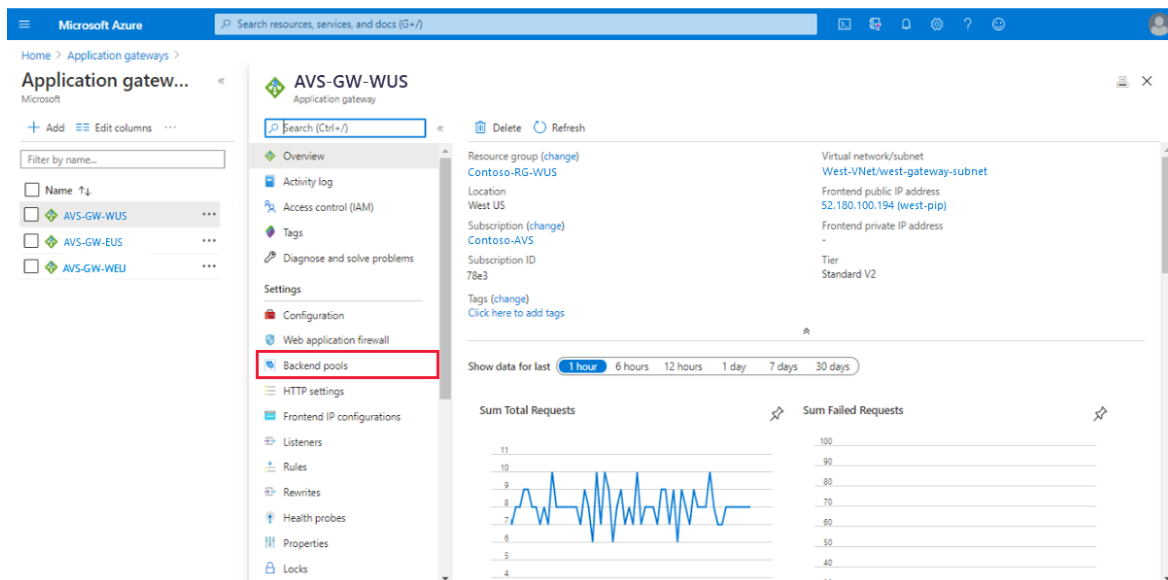
1. In the Azure portal, select **Application gateways** to view a list of your current application gateways:

- AVS-GW-WUS
- AVS-GW-EUS (on-premises)
- AVS-GW-WEU

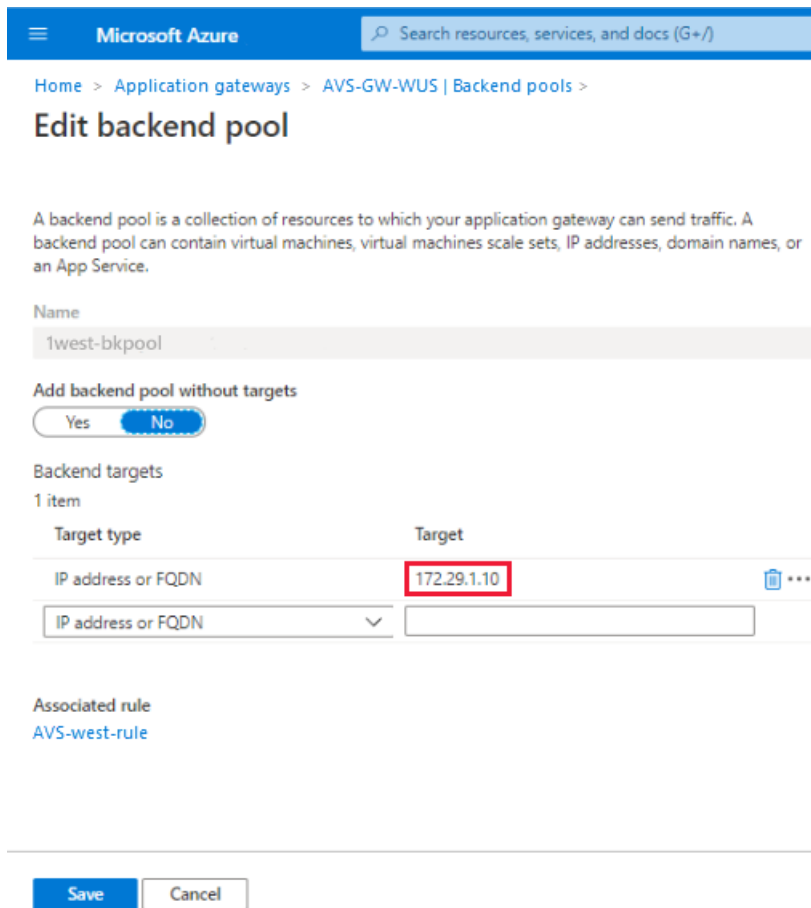


2. Select one of your previously deployed application gateways.

A window opens showing various information on the application gateway.



3. Select **Backend pools** to verify the configuration of one of the backend pools. You see one VM backend pool member configured as a web server with an IP address of 172.29.1.10.

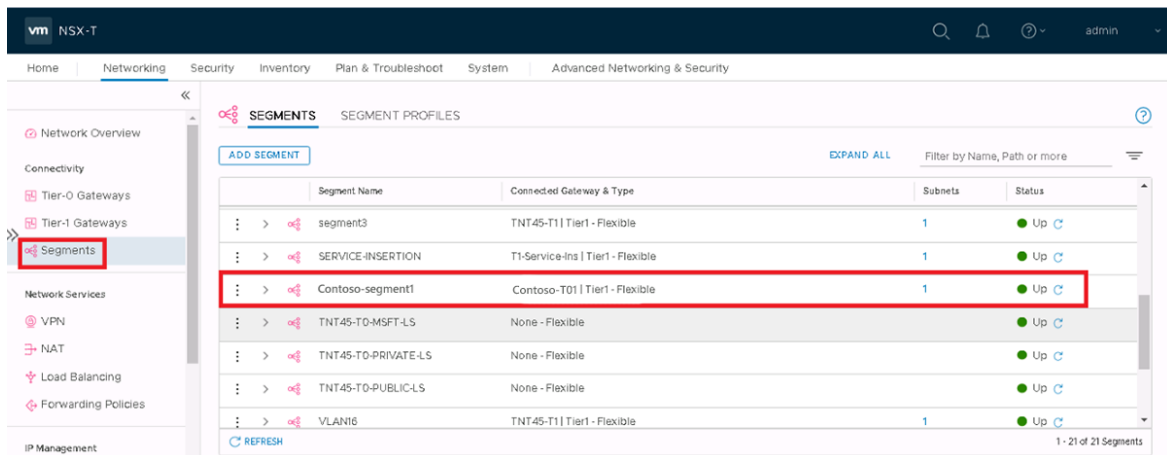


4. Verify the configuration of the other application gateways and backend pool members.

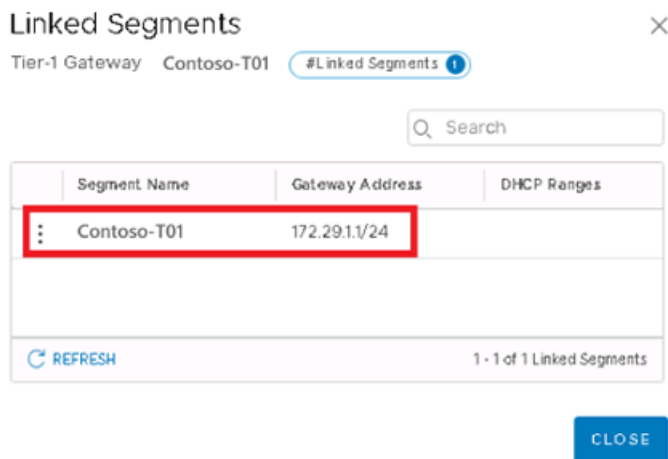
Verify the NSX-T segment configuration

The following steps verify the configuration of the NSX-T segment in the Azure VMware Solution environment.

1. Select **Segments** to view your configured segments. You see Contoso-segment1 connected to Contoso-T01 gateway, a Tier-1 flexible router.



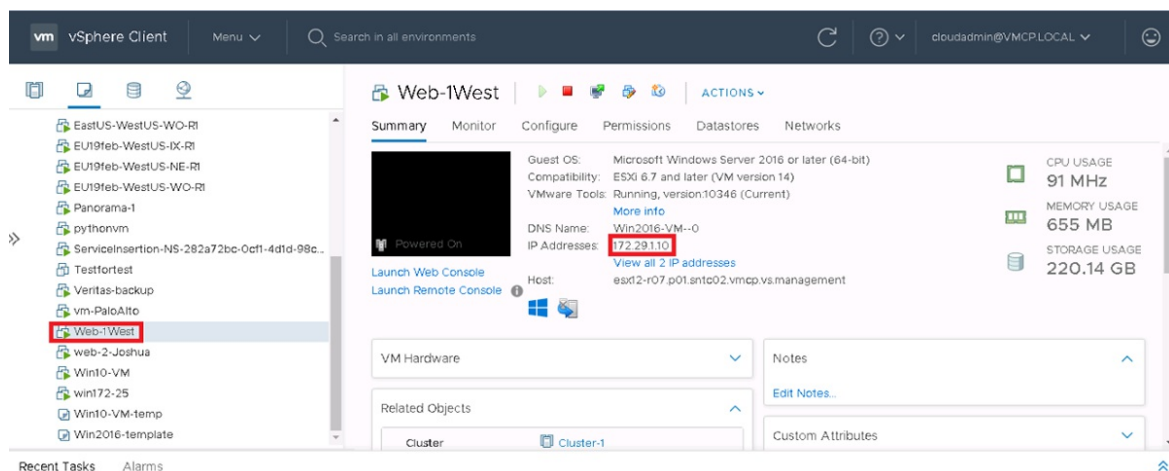
2. Select **Tier-1 Gateways** to see a list of Tier-1 gateways with the number of linked segments.



3. Select the segment linked to Contoso-T01. A window opens showing the logical interface configured on the Tier-01 router. It serves as a gateway to the backend pool member VM connected to the segment.
4. In the vSphere client, select the VM to view its details.

NOTE

Its IP address matches VM backend pool member configured as a web server from the preceding section: 172.29.1.10.



5. Select the VM, then select **ACTIONS > Edit Settings** to verify connection to the NSX-T segment.

Create your Traffic Manager profile

1. Sign in to the [Azure portal](#). Under **Azure Services > Networking**, select **Traffic Manager profiles**.
2. Select **+ Add** to create a new Traffic Manager profile.
3. Provide the following information and then select **Create**:
 - Profile name
 - Routing method (use [weighted](#))
 - Subscription
 - Resource group

Add external endpoints into the Traffic Manager profile

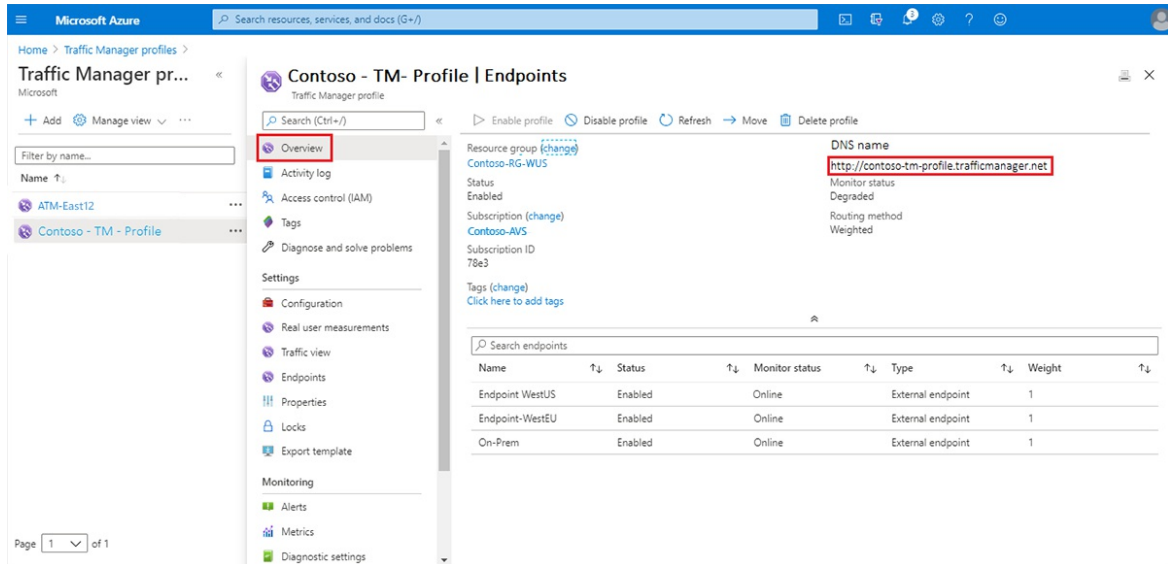
1. Select the Traffic Manager profile from the search results pane, select **Endpoints**, and then **+ Add**.

2. For each of the external endpoints in the different regions, enter the required details and then select **Add**:


- Type
- Name
- Fully Qualified domain name (FQDN) or IP
- Weight (assign a weight of 1 to each endpoint).

Once created, all three shows in the Traffic Manager profile. The monitor status of all three must be **Online**.


3. Select **Overview** and copy the URL under **DNS Name**.




4. Paste the DNS name URL in a browser. The screenshot shows traffic directing to the West Europe region.



5. Refresh your browser. The screenshot shows traffic directing to another set of backend pool members in the West US region.



6. Refresh your browser again. The screenshot shows traffic directing to the final set of backend pool members on-premises.



IIS-OnPrem

Next steps

Now that you've covered integrating Azure Traffic Manager with Azure VMware Solution, you may want to learn about:

- [Using Azure Application Gateway on Azure VMware Solution](#)
- [Traffic Manager routing methods](#)
- [Combining load-balancing services in Azure](#)
- [Measuring Traffic Manager performance](#)

Configure Azure Alerts in Azure VMware Solution

12/16/2022 • 3 minutes to read • [Edit Online](#)

In this article, you'll learn how to configure [Azure Action Groups](#) in [Microsoft Azure Alerts](#) to receive notifications of triggered events that you define. You'll also learn about using [Azure Monitor Metrics](#) to gain deeper insights into your Azure VMware Solution private cloud.

NOTE

Incidents affecting the availability of an Azure VMware Solution host and its corresponding restoration are sent automatically to the Account Administrator, Service Administrator (Classic Permission), Co-Admins (Classic Permission), and Owners (RBAC Role) of the subscription(s) containing Azure VMware Solution private clouds.

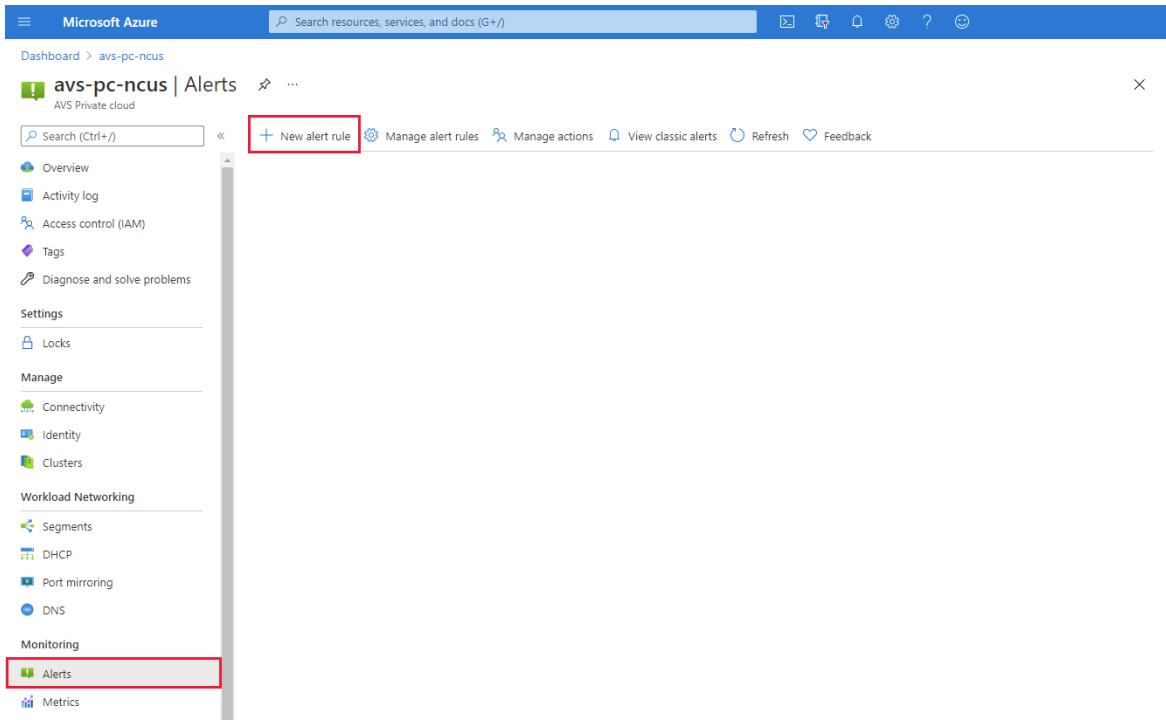
Supported metrics and activities

The following metrics are visible through Azure Monitor Metrics.

SIGNAL NAME	SIGNAL TYPE	MONITOR SERVICE
Datstore Disk Total Capacity	Metric	Platform
Percentage Datstore Disk Used	Metric	Platform
Percentage CPU	Metric	Platform
Average Effective Memory	Metric	Platform
Average Memory Overhead	Metric	Platform
Average Total Memory	Metric	Platform
Average Memory Usage	Metric	Platform
Datstore Disk Used	Metric	Platform
All Administrative operations	Activity Log	Administrative
Register Microsoft.AVS resource provider. (Microsoft.AVS/privateClouds)	Activity Log	Administrative
Create or update a PrivateCloud. (Microsoft.AVS/privateClouds)	Activity Log	Administrative
Delete a PrivateCloud. (Microsoft.AVS/privateClouds)	Activity Log	Administrative

Configure an alert rule

1. From your Azure VMware Solution private cloud, select **Monitoring > Alerts**, and then **New alert rule**.



A new configuration screen opens where you'll:

- Define the Scope
- Configure a Condition
- Set up the Action Group
- Define the Alert rule details

Create alert rule

Rules management

Condition name

Whenever the Activity Log has an event with Category='Administrative', Signal name='All Administrative operations'

Select condition

i You can only define one Activity Log signal per alert rule. To alert on more signals, please create additional alert rules.

Action group

Send notifications or invoke actions when the alert rule triggers, by selecting or creating a new action group. [Learn more](#)

Action group name

Contains actions

Send email action group

1 Email Azure Resource Manager Role ⓘ

[Select action group](#)

Alert rule details

Provide details on your alert rule so that you can identify and manage it later.

Alert rule name * ⓘ

Email alert

Description

Send email to owners on all administrative operations on the lab

Save alert rule to resource group * ⓘ

contosorg

Enable alert rule upon creation

[Create alert rule](#)

- Under **Scope**, select the target resource you want to monitor. By default, the Azure VMware Solution private cloud from where you opened the Alerts menu has been defined.
- Under **Condition**, select **Add condition**, and in the window that opens, selects the signal you want to create for the alert rule.

In our example, we've selected **Percentage Datastore Disk Used**, which is relevant from an [Azure VMware Solution SLA](#) perspective.

Configure signal logic ✕

Choose a signal below and configure the logic on the next screen to define the alert condition.

Signal type ⓘ

Monitor service ⓘ

Displaying 1 - 12 signals out of total 12 signals

Signal name	↑↓	Signal type	↑↓	Monitor service	↑↓
Datastore Disk Total Capacity	↕	Metric		Platform	
Percentage Datastore Disk Used	↕	Metric		Platform	
Percentage CPU	↕	Metric		Platform	
Average Effective Memory	↕	Metric		Platform	
Average Memory Overhead	↕	Metric		Platform	
Average Total Memory	↕	Metric		Platform	
Average Memory Usage	↕	Metric		Platform	
Datastore Disk Used	↕	Metric		Platform	
All Administrative operations	📄	Activity Log		Administrative	
Register Microsoft.AVS resource provider. (Microsoft.AVS/privateClouds)	📄	Activity Log		Administrative	
Create or update a PrivateCloud. (Microsoft.AVS/privateClouds)	📄	Activity Log		Administrative	
Delete a PrivateCloud. (Microsoft.AVS/privateClouds)	📄	Activity Log		Administrative	

- Define the logic that will trigger the alert and then select **Done**.

In our example, only the **Threshold** and **Frequency of evaluation** have been adjusted.


Configure signal logic



Define the logic for triggering an alert. Use the chart to view trends in the data.

[← Edit signal](#)

Selected signal: Percentage Datastore Disk Used (Platform)



Split by dimensions

Use dimensions to monitor specific time series. If you select more than one dimension value, each time series that results from the combination will trigger its own alert and will be charged separately. [About monitoring multiple time series](#) ⓘ

Dimension name	Operator	Dimension values
Select dimension	=	0 selected

Add custom value

Alert logic ⓘ Monitoring 1 time series (\$0.1/time series)

Threshold ⓘ

Static Dynamic

Operator ⓘ

Aggregation type * ⓘ

Threshold value * ⓘ %

Condition preview

Whenever the average percentage datastore disk used is greater than 65%

Evaluated based on

Aggregation granularity (Period) * ⓘ

Frequency of evaluation ⓘ

Done

5. Under **Actions**, select **Add action groups**. The action group defines *how* the notification is received and *who* receives it. You can receive notifications by email, SMS, [Azure Mobile App Push Notification](#) or voice message.
6. Select an existing action group or select **Create action group** to create a new one.
7. In the window that opens, on the **Basics** tab, give the action group a name and a display name.
8. Select the **Notifications** tab, select a **Notification Type** and **Name**. Then select **OK**.

Our example is based on email notification.

Home > Alerts > Manage actions >

Create action group

Basics **Notifications** Actions Tags Review + create

Notifications

Configure the method in which users will be notified when the action group triggers. Select notification types, provide receiver details and add a unique description. This step is optional.

Notification type	Name	Selected
Email/SMS message/Push/Voice	Notify on-call team	Email
Email Azure Resource Manager Role	Notify subscription owners	Owner

Email
 Email * on-call@contoso.com

SMS (Carrier charges may apply)
 Country code 1
 Phone number

Azure app Push Notifications
 Azure account email

Voice
 Country code 1
 Phone number

Enable the common alert schema. [Learn more](#)

Yes **No**

OK

Review + create Previous Next: Actions >

9. (Optional) Configure the **Actions** if you want to take proactive actions and receive notification on the event. Select an available **Action type** and then select **Review + create**.

- **Automation Runbooks** - to automate tasks based on alerts
- **Azure Functions** – for custom event-driven serverless code execution
- **ITSM** – to integrate with a service provider like ServiceNow to create a ticket
- **Logic App** - for more complex workflow orchestration
- **Webhooks** - to trigger a process in another service

10. Under the **Alert rule details**, provide a name, description, resource group to store the alert rule, the severity. Then select **Create alert rule**.

The alert rule is visible and can be managed from the Azure portal.

Microsoft Azure Search resources, services, and docs (G+)

Dashboard > avs-pc-ncus > avs-ncus > avs-pc-ncus >

Rules

Rules management

+ New alert rule Edit columns Manage actions View classic alerts Refresh | Enable Disable Delete

Subscription: Microsoft Azure Resource group: avs-ncus Resource type: All Resource: avs-pc-ncus Signal type: All signal types

Status: Enabled

Displaying 1 - 1 rules out of total 1 rules

Search alert rules based on rule name and condition...

Name	Condition	Status	Target resource	Target resource type	Signal type
<input type="checkbox"/> AVS - Datastore disk used greater...	Whenever the average diskusedp...	Enabled	avs-pc-ncus	AVS Private clouds	Metrics

As soon as a metric reaches the threshold as defined in an alert rule, the **Alerts** menu is updated and made visible.

Microsoft Azure | Search resources, services, and docs (G+)

Dashboard > avs-pc-ncus

avs-pc-ncus | Alerts

AVS Private cloud

Search (Ctrl+/) << + New alert rule Manage alert rules Manage actions View classic alerts Refresh Feedback

Subscription: Microsoft Azure Resource group: avs-ncus Time range: Past 24 hours Resource: avs-pc-ncus

Total alerts: 1 (Since 3/23/2021, 7:09 PM) Smart groups (preview): 0 (0% Reduction) Total alert rules: 1 (Enabled 1) Action rules (preview): 0 (Enabled 0) [Learn more About alerts](#)

Severity	Total alerts	New	Acknowledged	Closed
0 - Critical	0	0	0	0
1 - Error	0	0	0	0
2 - Warning	1	1	0	0
3 - Informational	0	0	0	0
4 - Verbose	0	0	0	0

Navigation: Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Settings (Locks), Manage (Connectivity, Identity, Clusters), Workload Networking (Segments, DHCP, Port mirroring, DNS), Monitoring (Alerts, Metrics)

Depending on the configured Action Group, you'll receive a notification through the configured medium. In our example, we've configured email.



Microsoft Azure
To: [redacted]

Reply Reply All Forward

If there are problems with how this message is displayed, click here to view it in a web browser.

Microsoft Azure

Fired:Sev2 Azure Monitor Alert AVS - Datastore disk used greater then on avs-pc-ncus (microsoft.avs/privateclouds) at 3/24/2021 3:47:12 PM

[View the alert in Azure Monitor >](#)

Summary

Alert name	AVS - Datastore disk used greater then
Severity	Sev2
Monitor condition	Fired
Affected resource	avs-pc-ncus
Resource type	microsoft.avs/privateclouds
Resource group	avs-ncus
Subscription	Microsoft Azure
Monitoring service	Platform
Signal type	Metric
Fired time	March 24, 2021 15:47 UTC
Alert ID	35a9ebaf-712e-427c-47d6-55933942e411
Alert rule ID	https://portal.azure.com/#blade/Microsoft_Azure_Monitoring/UpdateVNextAlertRuleBlade/ruleInputs
Metric alert condition type	SingleResourceMultipleMetricCriteria
Time aggregation	Average
Metric name	DiskUsedPercentage
Metric namespace	microsoft.avs/privateclouds
Metric value (when alert fired)	37
Operator	GreaterThan
Threshold	26

You're receiving this notification as a member of the AVS-ActionGr action group. To unsubscribe from emails directed to this action group, [click here](#).

[f](#) [t](#) [v](#) [in](#)

[Privacy Statement](#)

Microsoft Corporation, One Microsoft Way, Redmond, WA 98052

Microsoft

Work with metrics

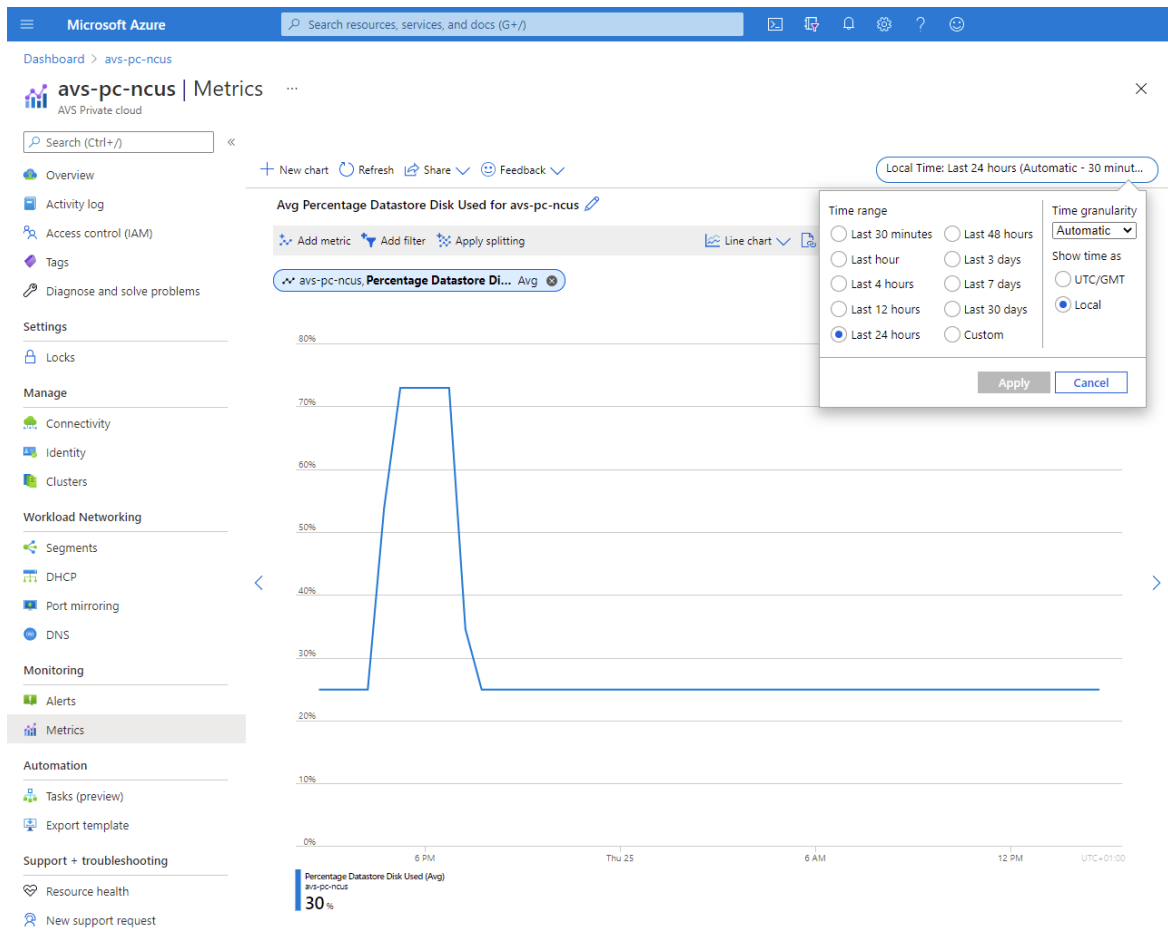
1. From your Azure VMware Solution private cloud, select **Monitoring > Metrics**. Then select the metric you want from the drop-down.

The screenshot shows the Microsoft Azure portal interface for the 'avs-pc-ncus' resource. The left-hand navigation pane includes sections for Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Settings (Locks, Manage, Connectivity, Identity, Clusters), Workload Networking (Segments, DHCP, Port mirroring, DNS), Monitoring (Alerts, Metrics), Automation (Tasks (preview), Export template), and Support + troubleshooting (Resource health, New support request). The 'Metrics' option in the Monitoring section is highlighted with a red box. The main content area displays a chart configuration interface. At the top, there are options for '+ New chart', 'Refresh', 'Share', and 'Feedback'. A search bar is present with the text 'Search (Ctrl+/)'. The chart configuration includes a 'Scope' dropdown set to 'avs-pc-ncus', a 'Metric Namespace' dropdown set to 'Standard metrics', and two dropdowns for 'Metric' and 'Aggregation'. The 'Metric' dropdown is open, showing a list of metrics including 'Percentage Datastore Disk Used', 'Average Effective Memory', 'Average Memory Overhead', 'Average Memory Usage', 'Average Total Memory', 'Datastore Disk Total Capacity', 'Datastore Disk Used', and 'Percentage CPU'. Below the configuration area, a message prompts the user to 'Select a metric above to see data appear on this chart or learn more below:'. Three buttons are provided: 'Filter + Split', 'Plot multiple metrics', and 'Build custom dashboards'. The chart area itself is currently empty, showing a y-axis from 0 to 100 and an x-axis with time markers (6 PM, Thu 25, 6 AM, 12 PM, 1/27-01:00). The 'Local Time: Last 24 hours (Automatic)' is displayed in the top right corner.

2. You can change the diagram's parameters, such as the **Time range** or the **Time granularity**.

Other options are:

- **Drill into Logs** and query the data in the related Log Analytics workspace
- **Pin this diagram** to an Azure Dashboard for convenience.



Next steps

Now that you've configured an alert rule for your Azure VMware Solution private cloud, you may want to learn even more about:

- [Azure Monitor Metrics](#)
- [Azure Monitor Alerts](#)
- [Azure Action Groups](#)

You can also continue with one of the other [Azure VMware Solution](#) how-to guides.

Configure customer-managed key encryption at rest in Azure VMware Solution

12/16/2022 • 8 minutes to read • [Edit Online](#)

This article illustrates how to encrypt VMware vSAN Key Encryption Keys (KEKs) with customer-managed keys (CMKs) managed by customer-owned Azure Key Vault.

When CMK encryptions are enabled on your Azure VMware Solution private cloud, Azure VMware Solution uses the CMK from your key vault to encrypt the vSAN KEKs. Each ESXi host that participates in the vSAN cluster uses randomly generated Disk Encryption Keys (DEKs) that ESXi uses to encrypt disk data at rest. vSAN encrypts all DEKs with a KEK provided by Azure VMware Solution key management system (KMS). Azure VMware Solution private cloud and Azure Key Vault don't need to be in the same subscription.

When managing your own encryption keys, you can do the following actions:

- Control Azure access to vSAN keys.
- Centrally manage the lifecycle of CMKs.
- Revoke Azure from accessing the KEK.

The Customer-managed keys (CMKs) feature supports the following key types. See the following key types, shown by key type and key size.

- RSA: 2048, 3072, 4096
- RSA-HSM: 2048, 3072, 4096

Topology

The following diagram shows how Azure VMware Solution uses Azure Active Directory (Azure AD) and a key vault to deliver the customer-managed key.



Prerequisites

Before you begin to enable customer-managed key (CMK) functionality, ensure the following listed requirements are met:

- You'll need an Azure Key Vault to use CMK functionality. If you don't have an Azure Key Vault, you can create one using [Quickstart: Create a key vault using the Azure portal](#).
- If you enabled restricted access to key vault, you'll need to allow Microsoft Trusted Services to bypass the Azure Key Vault firewall. Go to [Configure Azure Key Vault networking settings](#) to learn more.

NOTE

After firewall rules are in effect, users can only perform Key Vault [data plane](#) operations when their requests originate from allowed VMs or IPv4 address ranges. This also applies to accessing key vault from the Azure portal. This also affects the key vault Picker by Azure VMware Solution. Users may be able to see a list of key vaults, but not list keys, if firewall rules prevent their client machine or user does not have list permission in key vault.

- Enable **System Assigned identity** on your Azure VMware Solution private cloud if you didn't enable it during software-defined data center (SDDC) provisioning.
 - [Portal](#)
 - [Azure CLI](#)

Use the following steps to enable System Assigned identity:

1. Sign in to Azure portal.
2. Navigate to **Azure VMware Solution** and locate your SDDC.
3. From the left navigation, open **Manage** and select **Identity**.
4. In **System Assigned**, check **Enable** and select **Save**.
 - a. **System Assigned identity** should now be enabled.

Once System Assigned identity is enabled, you'll see the tab for **Object ID**. Make note of the Object ID for use later.

- Configure the key vault access policy to grant permissions to the managed identity. It will be used to authorize access to the key vault.
 - [Portal](#)
 - [Azure CLI](#)
 1. Sign in to Azure portal.
 2. Navigate to **Key vaults** and locate the key vault you want to use.
 3. From the left navigation, under **Settings**, select **Access policies**.
 4. In **Access policies**, select **Add Access Policy**.
 - a. From the Key Permissions drop-down, check **Select all**, **Unwrap Key**, and **Wrap Key**.
 - b. Under Select principal, select **None selected**. A new **Principal** window with a search box will open.
 - c. In the search box, paste the **Object ID** from the previous step, or search the private cloud name you want to use. Choose **Select** when you're done.
 - d. Select **ADD**.
 - e. Verify the new policy appears under the current policy's Application section.
 - f. Select **Save** to commit changes.

Customer-managed key version lifecycle

You can change the customer-managed key (CMK) by creating a new version of the key. The creation of a new version won't interrupt the virtual machine (VM) workflow.

In Azure VMware Solution, CMK key version rotation will depend on the key selection setting you've chosen during CMK setup.

Key selection setting 1

A customer enables CMK encryption without supplying a specific key version for CMK. Azure VMware Solution selects the latest key version for CMK from the customer's key vault to encrypt the vSAN Key Encryption Keys (KEKs). Azure VMware Solution tracks the CMK for version rotation. When a new version of the CMK key in Azure Key Vault is created, it's captured by Azure VMware Solution automatically to encrypt vSAN KEKs.

NOTE

Azure VMware Solution can take up to ten minutes to detect a new auto-rotated key version.

Key selection setting 2

A customer can enable CMK encryption for a specified CMK key version to supply the full key version URI under the **Enter Key from URI** option. When the customer's current key expires, they'll need to extend the CMK key expiration or disable CMK.

Enable CMK with system-assigned identity

System-assigned identity is restricted to one per resource and is tied to the lifecycle of the resource. You can grant permissions to the managed identity on Azure resource. The managed identity is authenticated with Azure AD, so you don't have to store any credentials in code.

IMPORTANT

Ensure that key vault is in the same region as the Azure VMware Solution private cloud.

- [Portal](#)
- [Azure CLI](#)

Navigate to your **Azure Key Vault** and provide access to the SDDC on Azure Key Vault using the Principal ID captured in the **Enable MSI** tab.

1. From your Azure VMware Solution private cloud, under **Manage**, select **Encryption**, then select **Customer-managed keys (CMK)**.
2. CMK provides two options for **Key Selection** from Azure Key Vault.

Option 1

- a. Under **Encryption key**, choose the **select from Key Vault** button.
- b. Select the encryption type, then the **Select Key Vault and key** option.
- c. Select the **Key Vault and key** from the drop-down, then choose **Select**.

Option 2

- a. Under **Encryption key**, choose the **Enter key from URI** button.
- b. Enter a specific Key URI in the **Key URI** box.

IMPORTANT

If you want to select a specific key version instead of the automatically selected latest version, you'll need to specify the key URI with key version. This will affect the CMK key version life cycle.

3. Select **Save** to grant access to the resource.

Change from customer-managed key to Microsoft managed key

When a customer wants to change from a customer-managed key (CMK) to a Microsoft managed key (MMK), it won't interrupt VM workload. To make the change from CMK to MMK, use the following steps.

1. Select **Encryption**, located under **Manage** from your Azure VMware Solution private cloud.
2. Select **Microsoft-managed keys (MMK)**.
3. Select **Save**.

Limitations

The Azure Key Vault must be configured as recoverable.

- Configure Azure Key Vault with the **Soft Delete** option.
- Turn on **Purge Protection** to guard against force deletion of the secret vault, even after soft delete.

Updating CMK settings won't work if the key is expired or the Azure VMware Solution access key has been revoked.

Troubleshooting and best practices

Accidental deletion of a key

If you accidentally delete your key in the Azure Key Vault, private cloud won't be able to perform some cluster modification operations. To avoid this scenario, we recommend that you keep soft deletes enabled on key vault. This option ensures that, if a key is deleted, it can be recovered within a 90-day period as part of the default soft-delete retention. If you are within the 90-day period, you can restore the key in order to resolve the issue.

Restore key vault permission

If you have a private cloud that lost access to the customer managed key, check if Managed System Identity (MSI) requires permissions in key vault. The error notification returned from Azure may not correctly indicate MSI requiring permissions in key vault as the root cause. Remember, the required permissions are: get, wrapKey, and unwrapKey. See step 4 in [Prerequisites](#).

Fix expired key

If you aren't using the auto-rotate function and the Customer Managed Key has expired in key vault, you can change the expiration date on key.

Restore key vault access

Ensure Managed System Identity (MSI) is used for providing private cloud access to key vault.

Deletion of MSI

If you accidentally delete the Managed System Identity (MSI) associated with private cloud, you'll need to disable CMK, then follow the steps to enable CMK from start.

Next steps

Learn about [Azure Key Vault backup and restore](#)

Learn about [Azure Key Vault recovery](#)

Attach Azure NetApp Files datastores to Azure VMware Solution hosts

12/16/2022 • 9 minutes to read • [Edit Online](#)

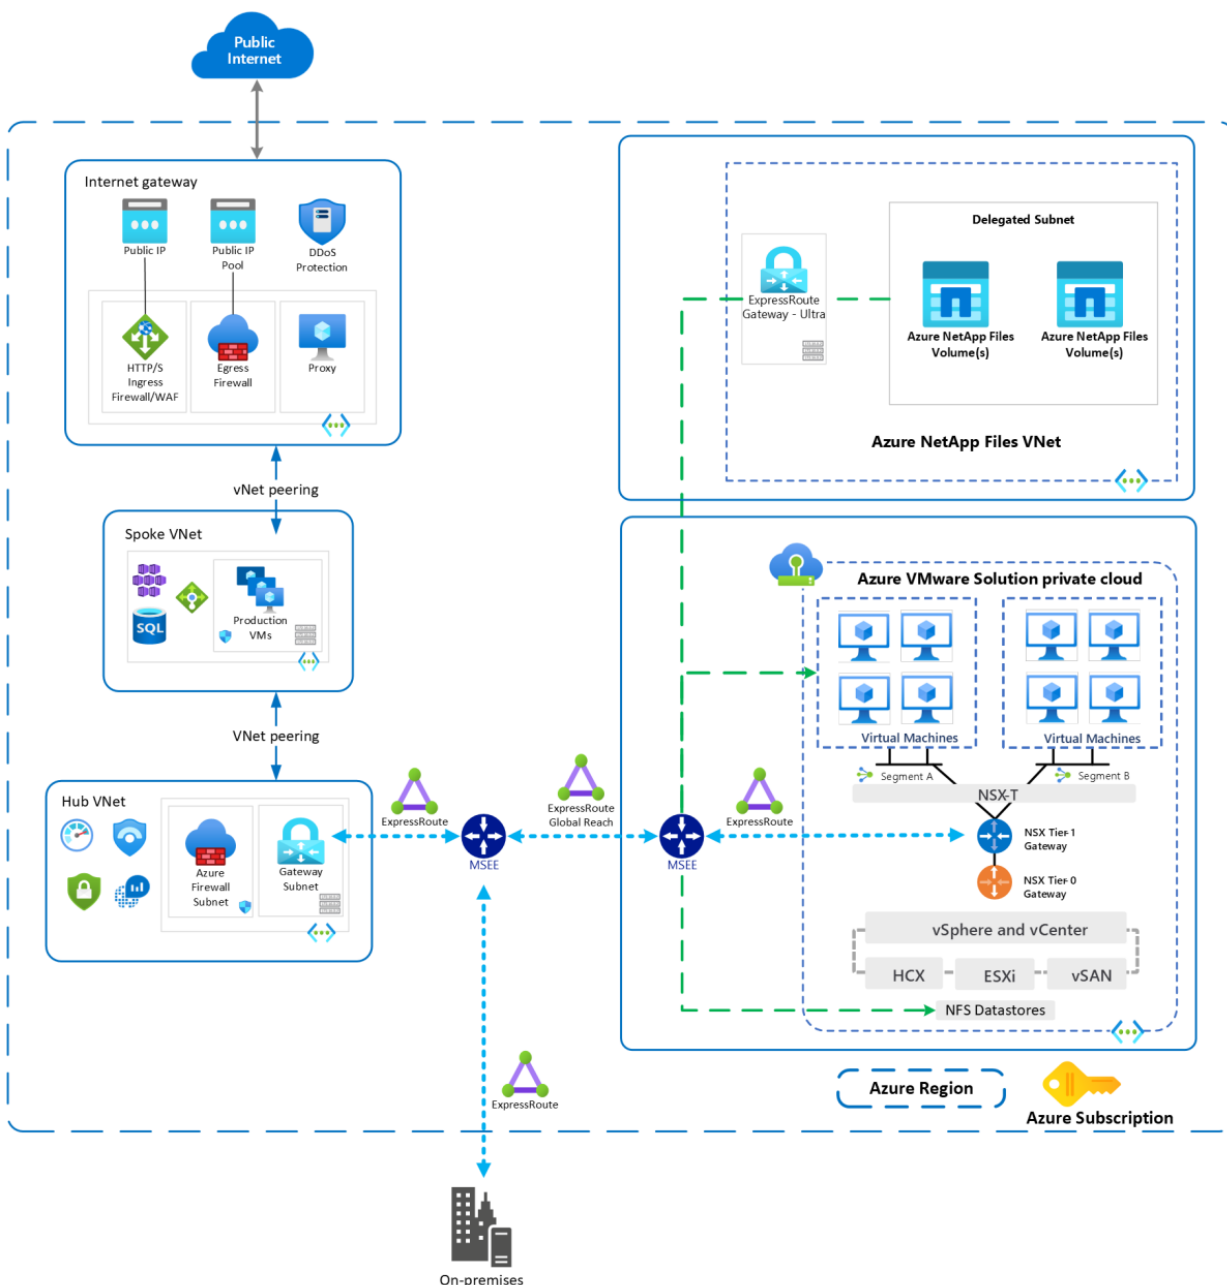
[Azure NetApp Files](#) is an enterprise-class, high-performance, metered file storage service. The service supports the most demanding enterprise file-workloads in the cloud: databases, SAP, and high-performance computing applications, with no code changes. For more information on Azure NetApp Files, see [Azure NetApp Files](#) documentation.

[Azure VMware Solution](#) supports attaching Network File System (NFS) datastores as a persistent storage option. You can create NFS datastores with Azure NetApp Files volumes and attach them to clusters of your choice. You can also create virtual machines (VMs) for optimal cost and performance.

By using NFS datastores backed by Azure NetApp Files, you can expand your storage instead of scaling the clusters. You can also use Azure NetApp Files volumes to replicate data from on-premises or primary VMware environments for the secondary site.

Create your Azure VMware Solution and create Azure NetApp Files NFS volumes in the virtual network connected to it using an ExpressRoute. Ensure there's connectivity from the private cloud to the NFS volumes created. Use those volumes to create NFS datastores and attach the datastores to clusters of your choice in a private cloud. As a native integration, no other permissions configured via vSphere are needed.

The following diagram demonstrates a typical architecture of Azure NetApp Files backed NFS datastores attached to an Azure VMware Solution private cloud via ExpressRoute.



Prerequisites

Before you begin the prerequisites, review the [Performance best practices](#) section to learn about optimal performance of NFS datastores on Azure NetApp Files volumes.

1. [Deploy Azure VMware Solution](#) private cloud and a dedicated virtual network connected via ExpressRoute gateway. The virtual network gateway should be configured with the Ultra performance SKU and have FastPath enabled. For more information, see [Configure networking for your VMware private cloud](#) and [Network planning checklist](#).
2. Create an [NFSv3 volume for Azure NetApp Files](#) in the same virtual network created in the previous step.
 - a. Verify connectivity from the private cloud to Azure NetApp Files volume by pinging the attached target IP. 2. Verify the subscription is registered to the `ANFAvsDataStore` feature in the `Microsoft.NetApp` namespace. If the subscription isn't registered, register it now.

```
az feature register --name "ANFAvsDataStore" --namespace "Microsoft.NetApp"
```

```
az feature show --name "ANFAvsDataStore" --namespace "Microsoft.NetApp" --query properties.state
```

- b. Based on your performance requirements, select the correct service level needed for the Azure

NetApp Files capacity pool. For optimal performance, it's recommended to use the Ultra tier. Select option **Azure VMware Solution Datastore** listed under the **Protocol** section.

- c. Create a volume with **Standard** [network features](#) if available for ExpressRoute FastPath connectivity.
- d. Under the **Protocol** section, select **Azure VMware Solution Datastore** to indicate the volume is created to use as a datastore for Azure VMware Solution private cloud.
- e. If you're using [export policies](#) to control access to Azure NetApp Files volumes, enable the Azure VMware private cloud IP range, not individual host IPs. Faulty hosts in a private cloud could get replaced so if the IP isn't enabled, connectivity to datastore will be impacted.

NOTE

Azure NetApp Files datastores for Azure VMware Solution are generally available. You must register Azure NetApp Files datastores for Azure VMware Solution before using it.

Supported regions

Azure VMware Solution currently supports the following regions:

Asia : East Asia, Japan East, Japan West, Southeast Asia.

Australia : Australia East, Australia Southeast.

Brazil : Brazil South.

Europe : France Central, Germany West Central, North Europe, Sweden Central, Sweden North, Switzerland West, UK South, UK West, West Europe

North America : Canada Central, Canada East, Central US, East US, East US 2, North Central US, South Central US, West US, West US 2.

Performance best practices

There are some important best practices to follow for optimal performance of NFS datastores on Azure NetApp Files volumes.

- Create Azure NetApp Files volumes using **Standard** network features to enable optimized connectivity from Azure VMware Solution private cloud via ExpressRoute FastPath connectivity.
- For optimized performance, choose **UltraPerformance** gateway and enable [ExpressRoute FastPath](#) from a private cloud to Azure NetApp Files volumes virtual network. View more detailed information on gateway SKUs at [About ExpressRoute virtual network gateways](#).
- Based on your performance requirements, select the correct service level needed for the Azure NetApp Files capacity pool. For best performance, it's recommended to use the Ultra tier.
- Create multiple datastores of 4-TB size for better performance. The default limit is 64 but it can be increased up to a maximum of 256 by submitting a support ticket. To submit a support ticket, go to [Create an Azure support request](#).
- Work with your Microsoft representative to ensure that the Azure VMware Solution private cloud and the Azure NetApp Files volumes are deployed within same [Availability Zone](#).

IMPORTANT

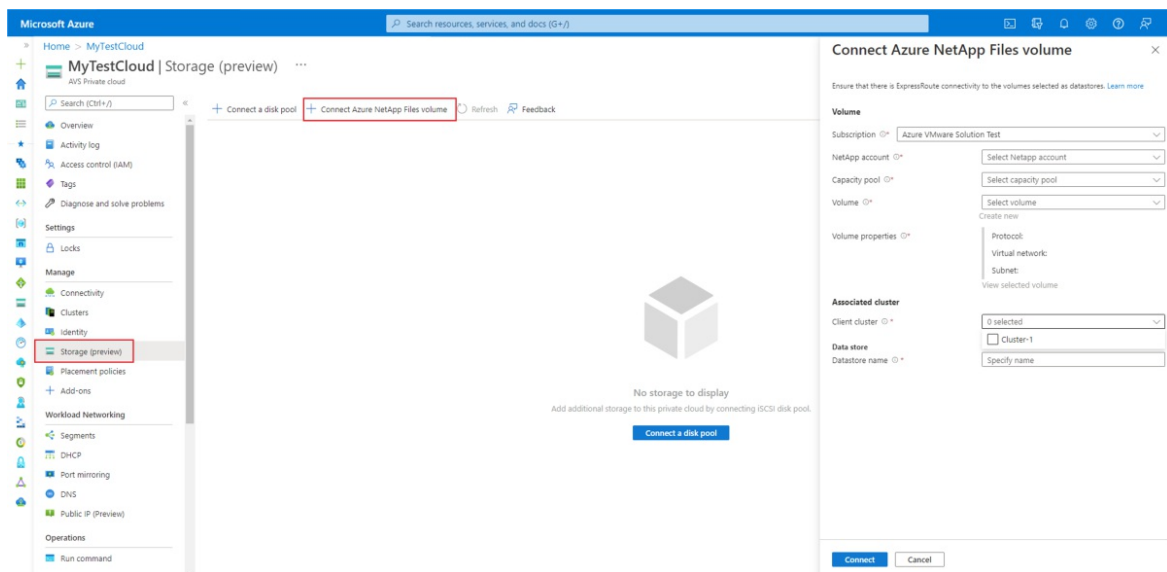
Changing the Azure NetApp Files volumes tier after creating the datastore will result in unexpected behavior in portal and API due to metadata mismatch. Set your performance tier of the Azure NetApp Files volume when creating the datastore. If you need to change tier during run time, detach the datastore, change the performance tier of the volume and attach the datastore. We are working on improvements to make this seamless.

Attach an Azure NetApp Files volume to your private cloud

- [Portal](#)
- [Azure CLI](#)

To attach an Azure NetApp Files volume to your private cloud using Portal, follow these steps:

1. Sign in to the Azure portal.
2. Select **Subscriptions** to see a list of subscriptions.
3. From the list, select the subscription you want to use.
4. Under Settings, select **Resource providers**.
5. Search for **Microsoft.AVS** and select it.
6. Select **Register**.
7. Under **Settings**, select **Preview features**.
 - a. Verify you're registered for both the `CCloudSanExperience` and `AnfDatstoreExperience` features.
8. Navigate to your Azure VMware Solution. Under **Manage**, select **Storage**.
9. Select **Connect Azure NetApp Files volume**.
10. In **Connect Azure NetApp Files volume**, select the **Subscription**, **NetApp account**, **Capacity pool**, and **Volume** to be attached as a datastore.



11. Verify the protocol is NFS. You'll need to verify the virtual network and subnet to ensure connectivity to the Azure VMware Solution private cloud.
12. Under **Associated cluster**, select the **Client cluster** to associate the NFS volume as a datastore
13. Under **Data store**, create a personalized name for your **Datastore name**.

- a. When the datastore is created, you should see all of your datastores in the **Storage**.
- b. You'll also notice that the NFS datastores are added in vCenter.

Disconnect an Azure NetApp Files-based datastore from your private cloud

You can use the instructions provided to disconnect an Azure NetApp Files-based (ANF) datastore using either Azure portal or Azure CLI. There's no maintenance window required for this operation. The disconnect action only disconnects the ANF volume as a datastore, it doesn't delete the data or the ANF volume.

Disconnect an ANF datastore using the Azure Portal

1. Select the datastore you want to disconnect from.
2. Right-click on the datastore and select **disconnect**.

Disconnect an ANF datastore using Azure CLI

```
az vmware datastore delete --name ANFDatastore1 --resource-group MyResourceGroup --cluster Cluster-1 --private-cloud MyPrivateCloud
```

Next steps

Now that you've attached a datastore on Azure NetApp Files-based NFS volume to your Azure VMware Solution hosts, you can create your VMs. Use the following resources to learn more.

- [Service levels for Azure NetApp Files](#)
- Datastore protection using [Azure NetApp Files snapshots](#)
- [About ExpressRoute virtual network gateways](#)
- [Understand Azure NetApp Files backup](#)
- [Guidelines for Azure NetApp Files network planning](#)

FAQs

- **Are there any special permissions required to create the datastore with the Azure NetApp Files volume and attach it onto the clusters in a private cloud?**

No other special permissions are needed. The datastore creation and attachment is implemented via Azure VMware Solution control plane.

- **Which NFS versions are supported?**

NFSv3 is supported for datastores on Azure NetApp Files.

- **Should Azure NetApp Files be in the same subscription as the private cloud?**

It's recommended to create the Azure NetApp Files volumes for the datastores in the same VNet that has connectivity to the private cloud.

- **How many datastores are we supporting with Azure VMware Solution?**

The default limit is 64 but it can be increased up to a maximum of 256 by submitting a support ticket. To submit a support ticket, go to [Create an Azure support request](#).

- **What latencies and bandwidth can be expected from the datastores backed by Azure NetApp Files?**

We're currently validating and working on benchmarking. For now, follow the [Performance best practices](#) outlined in this article.

- **What are my options for backup and recovery?**

Azure NetApp Files supports [snapshots](#) of datastores for quick checkpoints for near term recovery or quick clones. Azure NetApp Files backup lets you offload your Azure NetApp Files snapshots to Azure storage. With snapshots, copies and stores-changed blocks relative to previously offloaded snapshots are stored in an efficient format. This ability decreases Recovery Point Objective (RPO) and Recovery Time Objective (RTO) while lowering backup data transfer burden on the Azure VMware Solution service.

- **How do I monitor Storage Usage?**

Use [Metrics for Azure NetApp Files](#) to monitor storage and performance usage for the Datastore volume and to set alerts.

- **What metrics are available for monitoring?**

Usage and performance metrics are available for monitoring the Datastore volume. Replication metrics are also available for ANF datastore that can be replicated to another region using Cross Regional Replication. For more information about metrics, see [Metrics for Azure NetApp Files](#).

- **What happens if a new node is added to the cluster, or an existing node is removed from the cluster?**

When you add a new node to the cluster, it will automatically gain access to the datastore. Removing an existing node from the cluster won't affect the datastore.

- **How are the datastores charged, is there an additional charge?**

Azure NetApp Files NFS volumes that are used as datastores will be billed following the [capacity pool based billing model](#). Billing will depend on the service level. There's no extra charge for using Azure NetApp Files NFS volumes as datastores.

Attach Azure NetApp Files to Azure VMware Solution VMs

12/16/2022 • 3 minutes to read • [Edit Online](#)

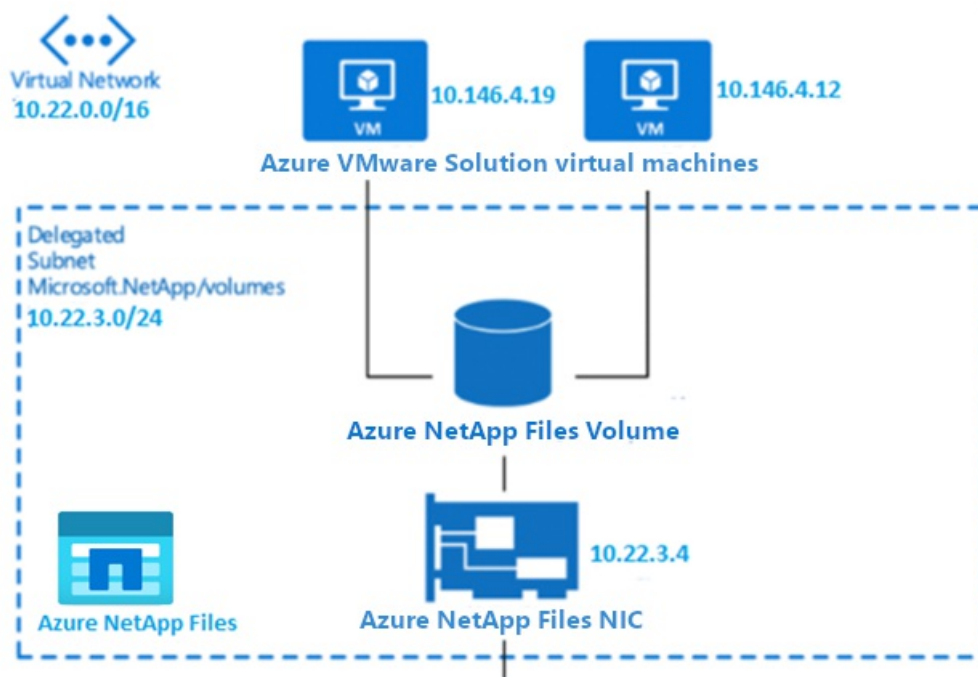
[Azure NetApp Files](#) is an Azure service for migration and running the most demanding enterprise file-workloads in the cloud: databases, SAP, and high-performance computing applications, with no code changes. In this article, you'll set up, test, and verify the Azure NetApp Files volume as a file share for Azure VMware Solution workloads using the Network File System (NFS) protocol. The guest operating system runs inside virtual machines (VMs) accessing Azure NetApp Files volumes.

Azure NetApp Files and Azure VMware Solution are created in the same Azure region. Azure NetApp Files is available in many [Azure regions](#) and supports cross-region replication. For information on Azure NetApp Files configuration methods, see [Storage hierarchy of Azure NetApp Files](#).

Services where Azure NetApp Files are used:

- **Active Directory connections:** Azure NetApp Files supports [Understand guidelines for Active Directory Domain Services site design and planning for Azure NetApp Files](#).
- **Share Protocol:** Azure NetApp Files supports Server Message Block (SMB) and Network File System (NFS) protocols. This support means the volumes can be mounted on the Linux client and can be mapped on Windows client.
- **Azure VMware Solution:** Azure NetApp Files shares can be mounted from VMs that are created in the Azure VMware Solution environment.

The diagram shows a connection through Azure ExpressRoute to an Azure VMware Solution private cloud. The Azure VMware Solution environment accesses the Azure NetApp Files share mounted on Azure VMware Solution VMs.



Prerequisites

- Azure subscription with Azure NetApp Files enabled
- Subnet for Azure NetApp Files
- Linux VM on Azure VMware Solution
- Windows VMs on Azure VMware Solution

Create and mount Azure NetApp Files volumes

You'll create and mount Azure NetApp Files volumes onto Azure VMware Solution VMs.

1. [Create a NetApp account.](#)
2. [Set up a capacity pool.](#)
3. [Create an SMB volume for Azure NetApp Files.](#)
4. [Create an NFS volume for Azure NetApp Files.](#)
5. [Delegate a subnet to Azure NetApp Files.](#)

Verify pre-configured Azure NetApp Files

You'll verify the pre-configured Azure NetApp Files created in Azure on Azure NetApp Files Premium service level.

1. In the Azure portal, under **STORAGE**, select **Azure NetApp Files**. A list of your configured Azure NetApp Files will show.

Home >

Azure NetApp Files




Microsoft

+ Add Edit columns Refresh Assign tags

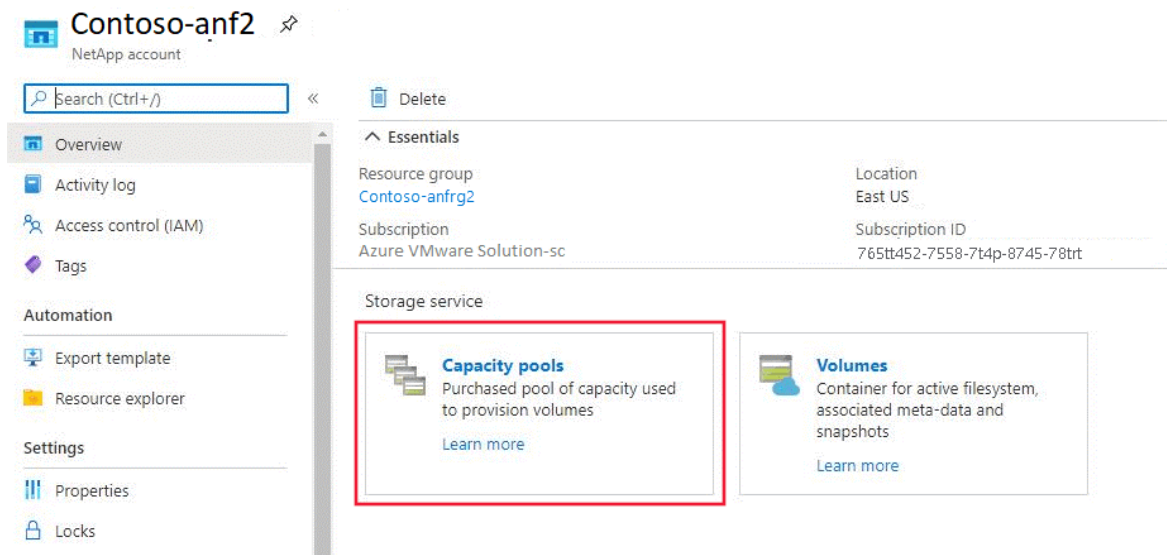
Subscriptions: All 3 selected – Don't see a subscription? Open Directory + Subscription settings

Filter by name... All subscriptions All resource groups All locations

3 items

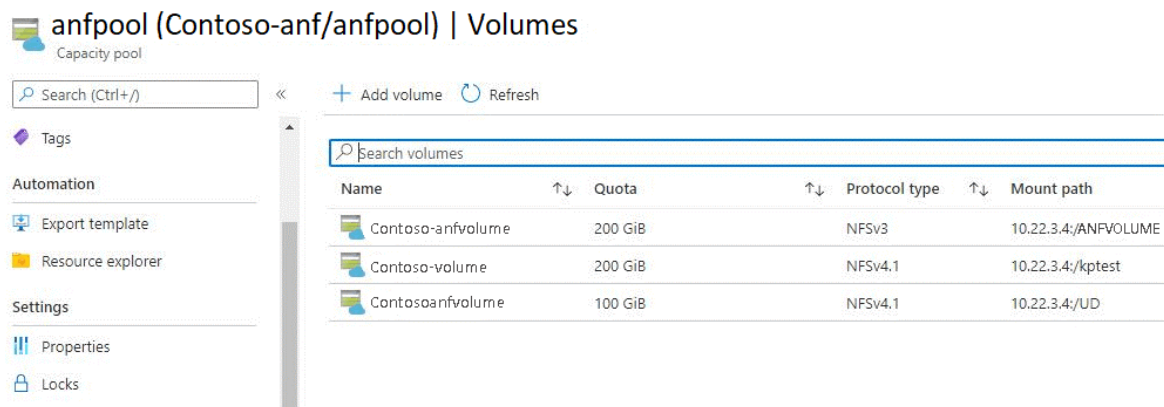
<input type="checkbox"/>	Name ↑↓	Type ↑↓	Resource group ↑↓	Location ↑↓
<input type="checkbox"/>	 Contoso-anf1	NetApp account	Contoso-anfrg1	West Europe
<input type="checkbox"/>	 Contoso-anf2	NetApp account	Contoso-anfrg2	East US
<input type="checkbox"/>	 Contoso-anf3	NetApp account	Contoso-anfrg3	East Europe

2. Select a configured NetApp Files account to view its settings. For example, select **Contoso-anf2**.
3. Select **Capacity pools** to verify the configured pool.

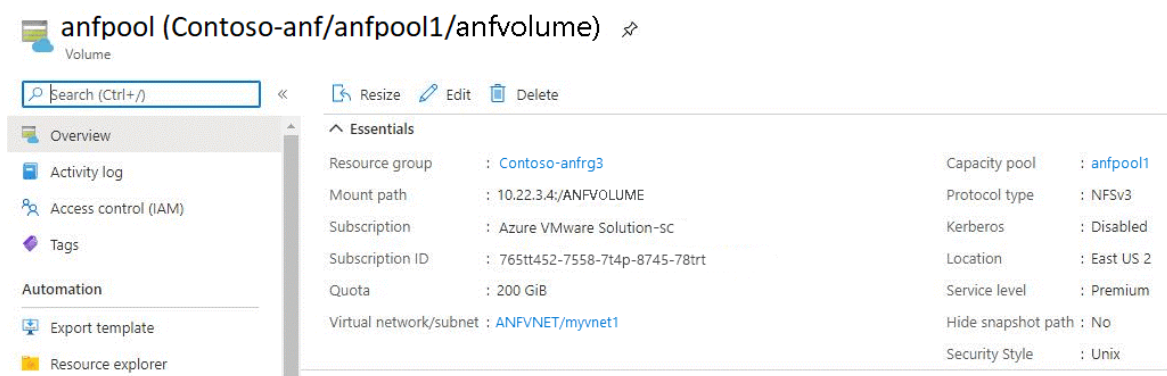


The Capacity pools page opens showing the capacity and service level. In this example, the storage pool is configured as 4 TiB with a Premium service level.

4. Select **Volumes** to view volumes created under the capacity pool. (See preceding screenshot.)
5. Select a volume to view its configuration.



A window opens showing the configuration details of the volume.



You can see that anfvolume has a size of 200 GiB and is in capacity pool anfpool1. It's exported as an NFS file share via 10.22.3.4:/ANFVOLUME. One private IP from the Azure Virtual Network (VNet) was created for Azure NetApp Files and the NFS path to mount on the VM.

To learn about Azure NetApp Files volume performance by size or "Quota," see [Performance considerations for Azure NetApp Files](#).

Verify pre-configured Azure VMware Solution VM share mapping

To make your Azure NetApp Files share accessible to your Azure VMware Solution VM, you'll need to

understand SMB and NFS share mapping. Only after configuring the SMB or NFS volumes, can you mount them as documented here.

- **SMB share:** Create an Active Directory connection before deploying an SMB volume. The specified domain controllers must be accessible by the delegated subnet of Azure NetApp Files for a successful connection. Once the Active Directory is configured within the Azure NetApp Files account, it will appear as a selectable item while creating SMB volumes.
- **NFS share:** Azure NetApp Files contributes to creating the volumes using NFS or dual protocol (NFS and SMB). A volume's capacity consumption counts against its pool's provisioned capacity. NFS can be mounted to the Linux server by using the command lines or /etc/fstab entries.

Next steps

Now that you've covered integrating Azure NetApp Files with your Azure VMware Solution workloads, you may want to learn about:

- [Resource limitations for Azure NetApp Files](#)
- [Guidelines for Azure NetApp Files network planning](#)
- [Cross-region replication of Azure NetApp Files volumes](#)
- [Azure NetApp Files NFS FAQs](#)
- [Azure NetApp Files SMB FAQs](#)

Set up Azure Backup Server for Azure VMware Solution

12/16/2022 • 15 minutes to read • [Edit Online](#)

Azure Backup Server contributes to your business continuity and disaster recovery (BCDR) strategy. With Azure VMware Solution, you can only configure a virtual machine (VM)-level backup using Azure Backup Server.

Azure Backup Server can store backup data to:

- **Disk:** For short-term storage, Azure Backup Server backs up data to disk pools.
- **Azure cloud:** For both short-term and long-term storage off-premises, Azure Backup Server data stored in disk pools can be backed up to the Microsoft Azure cloud by using Azure Backup.

Use Azure Backup Server to restore data to the source or an alternate location. That way, if the original data is unavailable because of planned or unexpected issues, you can restore data to an alternate location.

This article helps you prepare your Azure VMware Solution environment to back up VMs by using Azure Backup Server. We walk you through the steps to:

- Determine the recommended VM disk type and size to use.
- Create a Recovery Services vault that stores the recovery points.
- Set the storage replication for a Recovery Services vault.
- Add storage to Azure Backup Server.

Supported VMware features

- **Agentless backup:** Azure Backup Server doesn't require an agent to be installed on the vCenter Server or ESXi server to back up the VM. Instead, provide the IP address or fully qualified domain name (FQDN) and the sign in credentials used to authenticate the VMware vCenter Server with Azure Backup Server.
- **Cloud-integrated backup:** Azure Backup Server protects workloads to disk and the cloud. The backup and recovery workflow of Azure Backup Server helps you manage long-term retention and offsite backup.
- **Detect and protect VMs managed by vCenter Server:** Azure Backup Server detects and protects VMs deployed on a vCenter Server or ESXi hosts. Azure Backup Server also detects VMs managed by vCenter Server so that you can protect large deployments.
- **Folder-level auto protection:** vCenter Server lets you organize your VMs into Virtual Machine folders. Azure Backup Server detects these folders. You can use it to protect VMs at the folder level, including all subfolders. When protecting folders, Azure Backup Server protects the VMs in that folder and protects VMs added later. Azure Backup Server detects new VMs daily, protecting them automatically. As you organize your VMs in recursive folders, Azure Backup Server automatically detects and protects the new VMs deployed in the recursive folders.
- **Azure Backup Server continues to protect vMotioned VMs within the cluster:** As VMs are vMotioned for dynamic resource load balancing within the cluster, Azure Backup Server automatically detects and continues VM protection.
- **Recover necessary files faster:** Azure Backup Server can recover files or folders from a Windows VM without recovering the entire VM.
- **Application Consistent Backups:** If VMware Tools is not installed, a crash consistent backup will be executed. When VMware Tools is installed with Microsoft Windows virtual machines, all applications that support VSS freeze and thaw operations will support application consistent backups. When VMware Tools is installed with Linux virtual machines, application consistent snapshots are supported by calling the pre and

post scripts.

Limitations

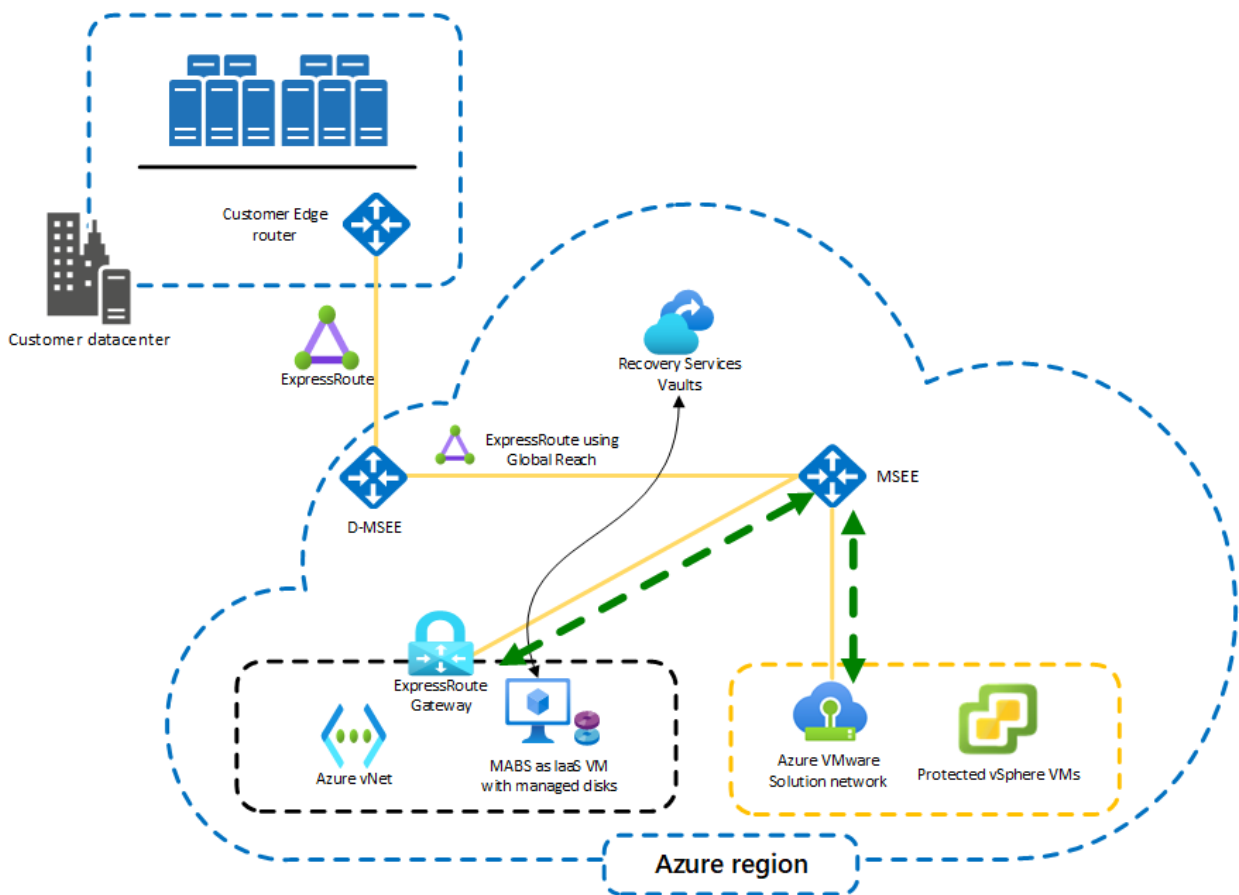
- Update Rollup 1 for Azure Backup Server v3 must be installed.
- You can't backup user snapshots before the first Azure Backup Server backup. After Azure Backup Server finishes the first backup, then you can back up user snapshots.
- Azure Backup Server can't protect VMware vSphere VMs with pass-through disks and physical raw device mappings (pRDMs).
- Azure Backup Server can't detect or protect VMware vSphere vApps.

To set up Azure Backup Server for Azure VMware Solution, you must finish the following steps:

- Set up the prerequisites and environment.
- Create a Recovery Services vault.
- Download and install Azure Backup Server.
- Add storage to Azure Backup Server.

Deployment architecture

Azure Backup Server is deployed as an Azure infrastructure as a service (IaaS) VM to protect Azure VMware Solution VMs.



Prerequisites for the Azure Backup Server environment

Consider the recommendations in this section when you install Azure Backup Server in your Azure environment.

Azure Virtual Network

Ensure that you [configure networking for your VMware private cloud in Azure](#).

Determine the size of the VM

Use the [MABS Capacity Planner](#) to determine the correct VM size. Based on your inputs, the capacity planner will give you the required memory size and CPU core count. Use this information to choose the appropriate Azure VM size. The capacity planner also provides total disk size required for the VM along with the required disk IOPS. We recommend using a standard SSD disk for the VM. By pooling more than one SSD, you can achieve the required IOPS.

Follow the instructions in the [Create your first Windows VM in the Azure portal](#) tutorial. You'll create the VM in the virtual network that you created in the previous step. Start with a gallery image of Windows Server 2019 Datacenter to run the Azure Backup Server.

NOTE

Azure Backup Server is designed to run on a dedicated, single-purpose server. You can't install Azure Backup Server on a computer that:

- Runs as a domain controller.
- Has the Application Server role installed.
- Is a System Center Operations Manager management server.
- Runs Exchange Server.
- Is a node of a cluster.

Disks and storage

Azure Backup Server requires disks for installation.

REQUIREMENT	RECOMMENDED SIZE
Azure Backup Server installation	Installation location: 3 GB Database files drive: 900 MB System drive: 1 GB for SQL Server installation You'll also need space for Azure Backup Server to copy the file catalog to a temporary installation location when you archive.
Disk for storage pool (Uses basic volumes, can't be on a dynamic disk)	Two to three times the protected data size. For detailed storage calculation, see DPM Capacity Planner .

To learn how to attach a new managed data disk to an existing Azure VM, see [Attach a managed data disk to a Windows VM by using the Azure portal](#).

NOTE

A single Azure Backup Server has a soft limit of 120 TB for the storage pool.

Store backup data on local disk and in Azure

Storing backup data in Azure reduces backup infrastructure on the Azure Backup Server VM. For operational recovery (backup), Azure Backup Server stores backup data on Azure disks attached to the VM. After the disks and storage space are attached to the VM, Azure Backup Server manages the storage for you. The amount of storage depends on the number and size of disks attached to each Azure VM. Each size of the Azure VM has a maximum number of disks that can be attached. For example, A2 is four disks, A3 is eight disks, and A4 is 16 disks. Again, the size and number of disks determine the total backup storage pool capacity.

IMPORTANT

You should *not* retain operational recovery data on Azure Backup Server-attached disks for more than five days. If data is more than five days old, store it in a Recovery Services vault.

To store backup data in Azure, create or use a Recovery Services vault. When you prepare to back up the Azure Backup Server workload, you [configure the Recovery Services vault](#). Once configured, each time an online backup job runs, a recovery point gets created in the vault. Each Recovery Services vault holds up to 9,999 recovery points. Depending on the number of recovery points created and how long kept, you can keep backup data for many years. For example, you could create monthly recovery points and keep them for five years.

IMPORTANT

Whether you send backup data to Azure or keep it locally, you must register Azure Backup Server with a Recovery Services vault.

Scale deployment

If you want to scale your deployment, you have the following options:

- **Scale up:** Increase the size of the Azure Backup Server VM from A series to DS3 series, and increase the local storage.
- **Offload data:** Send older data to Azure and keep only the newest data on the storage attached to the Azure Backup Server machine.
- **Scale out:** Add more Azure Backup Server machines to protect the workloads.

.NET Framework

The VM must have .NET Framework 3.5 SP1 or higher installed.

Join a domain

The Azure Backup Server VM must be joined to a domain. A domain user with administrator privileges on the VM must install Azure Backup Server.

Azure Backup Server deployed in an Azure VM can back up workloads on the VMs in Azure VMware Solution. The workloads should be in the same domain to enable the backup operation.

Create a Recovery Services vault

A Recovery Services vault is a storage entity that stores the recovery points created over time. It also contains backup policies that are associated with protected items.

1. Sign in to the [Azure portal](#), and on the left menu, select **All services**.
2. In the **All services** dialog box, enter **Recovery Services** and select **Recovery Services vaults** from the list.

The list of Recovery Services vaults in the subscription appears.

3. On the **Recovery Services vaults** dashboard, select **Add**.

The **Recovery Services vault** dialog box opens.

4. Enter values and then select **Create**.

- **Name:** Enter a friendly name to identify the vault. The name must be unique to the Azure subscription. Specify a name that has at least two but not more than 50 characters. The name must start with a letter and consist only of letters, numbers, and hyphens.

- **Subscription:** Choose the subscription to use. If you're a member of only one subscription, you'll see that name. If you're not sure which subscription to use, use the default (suggested) subscription. There are multiple choices only if your work or school account is associated with more than one Azure subscription.
- **Resource group:** Use an existing resource group or create a new one. To see the list of available resource groups in your subscription, select **Use existing**, and then select a resource from the drop-down list. To create a new resource group, select **Create new** and enter the name.
- **Location:** Select the geographic region for the vault. To create a vault to protect Azure VMware Solution virtual machines, the vault *must* be in the same region as the Azure VMware Solution private cloud.

It can take a while to create the Recovery Services vault. Monitor the status notifications in the **Notifications** area in the upper-right corner of the portal. After creating your vault, it's visible in the list of Recovery Services vaults. If you don't see your vault, select **Refresh**.

Set storage replication

The storage replication option lets you choose between geo-redundant storage (the default) and locally redundant storage. Geo-redundant storage copies the data in your storage account to a secondary region, making your data durable. Locally redundant storage is a cheaper option that isn't as durable. To learn more about geo-redundant and locally redundant storage options, see [Azure Storage redundancy](#).

IMPORTANT

Changing the setting of **Storage replication type** **Locally-redundant/Geo-redundant** for a Recovery Services vault must be done before you configure backups in the vault. After you configure backups, the option to modify it is disabled, and you can't change the storage replication type.

1. From **Recovery Services vaults**, select the new vault.
2. Under **Settings**, select **Properties**. Under **Backup Configuration**, select **Update**.
3. Select the storage replication type, and select **Save**.

Download and install the software package

Follow the steps in this section to download, extract, and install the software package.

Download the software package

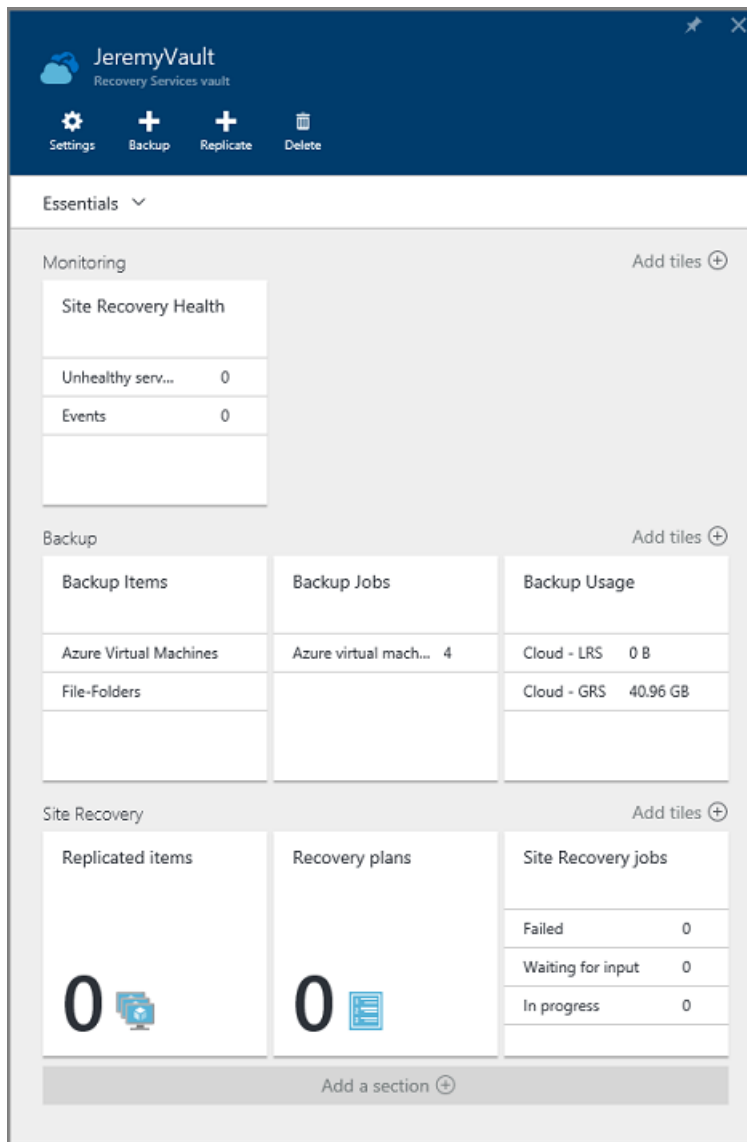
1. Sign in to the [Azure portal](#).
2. If you already have a Recovery Services vault open, continue to the next step.

TIP

If you don't have a Recovery Services vault open, and you're in the Azure portal, in the list of resources enter **Recovery Services > Recovery Services vaults**.

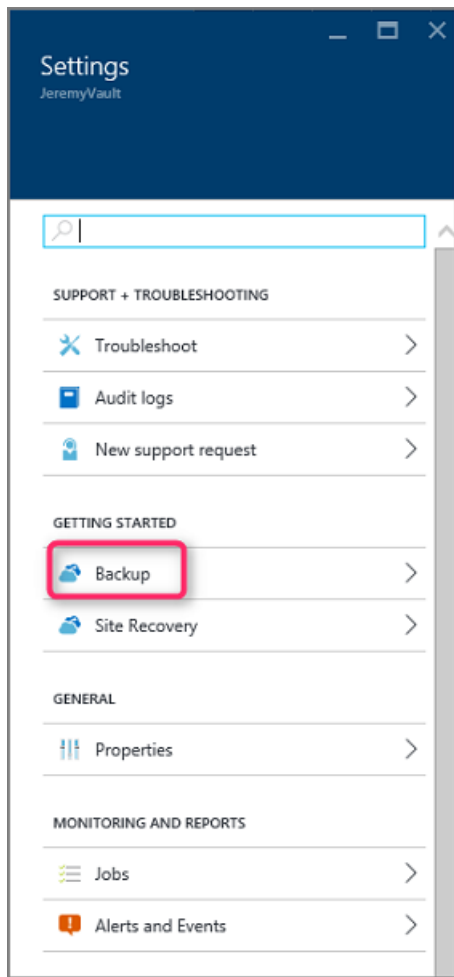
3. From the list of Recovery Services vaults, select a vault.

The selected vault dashboard opens.

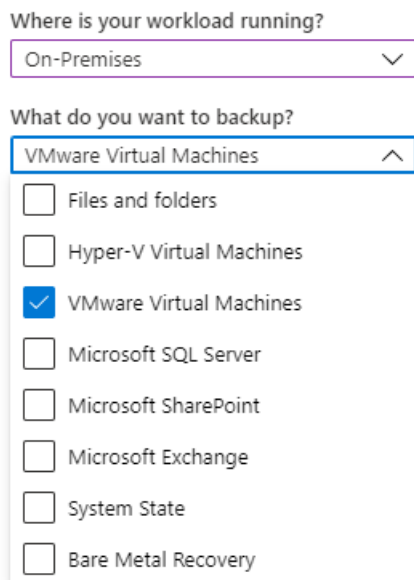


The **Settings** option opens by default. If closed, select **Settings** to open it.

4. Select **Backup** to open the **Getting Started** wizard.



5. In the window that opens:
- a. From the **Where is your workload running?** menu, select **On-Premises**.



- b. From the **What do you want to back up?** menu, select the workloads you want to protect by using Azure Backup Server.
- c. Select **Prepare Infrastructure** to download and install Azure Backup Server and the vault credentials.

Where is your workload running?

What do you want to backup?

Step: Prepare Infrastructure

[Prepare Infrastructu...](#)

6. In the **Prepare infrastructure** window that opens:
 - a. Select the **Download** link to install Azure Backup Server.
 - b. Select **Already downloaded or using the latest Azure Backup Server installation** and then **Download** to download the vault credentials. You'll use these credentials when you register the Azure Backup Server to the Recovery Services vault. The links take you to the Download Center, where you download the software package.

Prepare infrastructure ×

Already using [System Center Data Protection Manager](#) or any other [System Center Product](#)

Azure Backup Server
Please follow the steps mentioned below.

1. Install Microsoft Azure Backup Server
[Download](#)
2. Download vault credentials to register the server to the vault. Vault credentials will expire after 2 days.
 Already downloaded or using the latest Azure Backup Server installation

[Download](#)

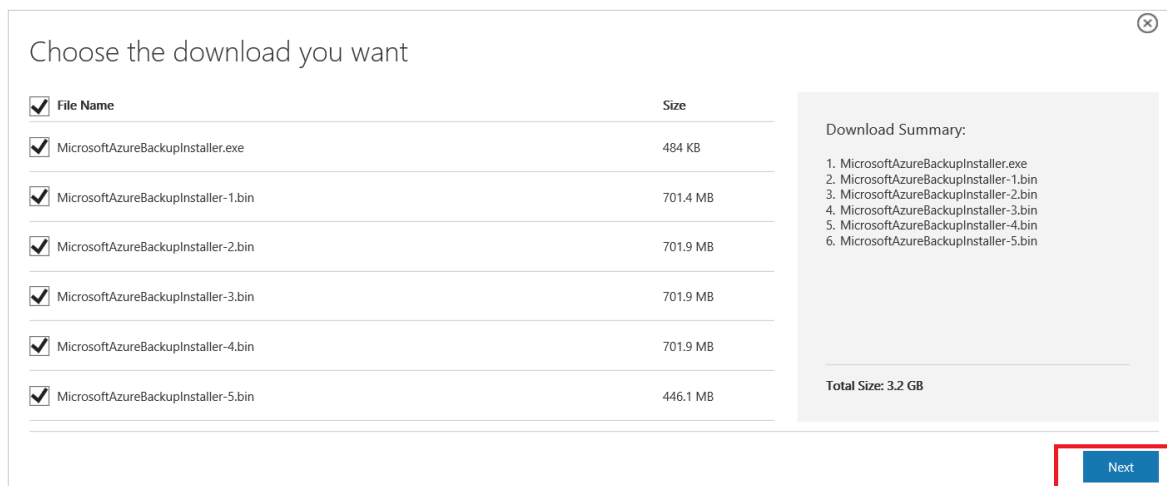
3. Post infrastructure preparation, please use Microsoft Azure Backup Server UI(on-premises) to configure backup.

[Learn More](#)

7. On the download page, select all the files and select **Next**.

NOTE

You must download all the files to the same folder. Because the download size of the files together is greater than 3 GB, it might take up to 60 minutes for the download to complete.



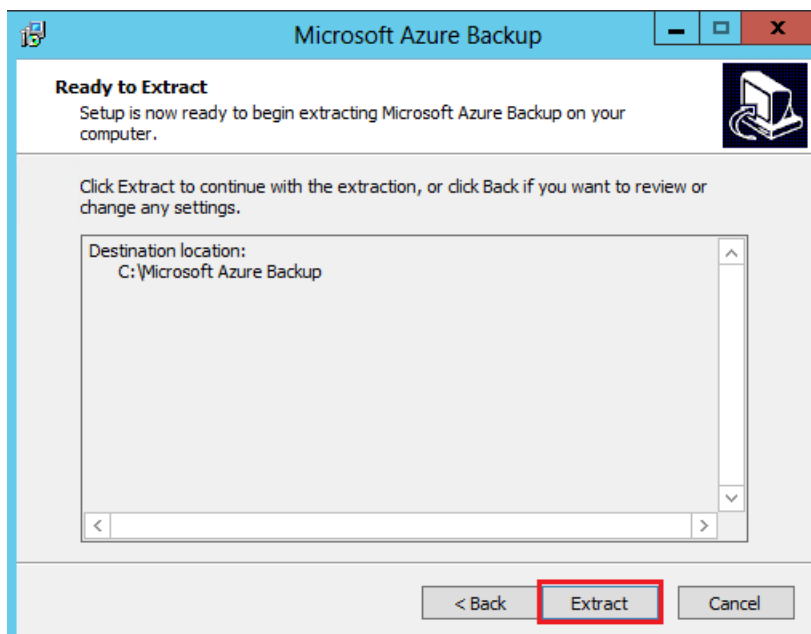
Extract the software package

If you downloaded the software package to a different server, copy the files to the VM you created to deploy Azure Backup Server.

WARNING

At least 4 GB of free space is required to extract the setup files.

1. After you've downloaded all the files, double-click **MicrosoftAzureBackupInstaller.exe** to open the **Microsoft Azure Backup** setup wizard, and then select **Next**.
2. Select the location to extract the files to and select **Next**.
3. Select **Extract** to begin the extraction process.



4. Once extracted, select the option to **Execute setup.exe** and then select **Finish**.


TIP

You can also locate the setup.exe file from the folder where you extracted the software package.

Install the software package

1. On the setup window under **Install**, select **Microsoft Azure Backup** to open the setup wizard.

2. On the **Welcome** screen, select **Next** to continue to the **Prerequisite Checks** page.
3. To determine if the hardware and software meet the prerequisites for Azure Backup Server, select **Check Again**. If met successfully, select **Next**.
4. The Azure Backup Server installation package comes bundled with the appropriate SQL Server binaries that are needed. When you start a new Azure Backup Server installation, select the **Install new Instance of SQL Server with this Setup** option. Then select **Check and Install**.



NOTE

If you want to use your own SQL Server instance, the supported SQL Server versions are SQL Server 2014 SP1 or higher, 2016, and 2017. All SQL Server versions should be Standard or Enterprise 64-bit. The instance used by Azure Backup Server must be local only; it can't be remote. If you use an existing SQL Server instance for Azure Backup Server, the setup only supports the use of *named instances* of SQL Server.

If a failure occurs with a recommendation to restart the machine, do so, and select **Check Again**. For any SQL Server configuration issues, reconfigure SQL Server according to the SQL Server guidelines. Then retry to install or upgrade Azure Backup Server using the existing instance of SQL Server.

Manual configuration

When you use your own SQL Server instance, make sure you add builtin\Administrators to the sysadmin role to the main database sysadmin role.

Configure reporting services with SQL Server 2017

If you use your instance of SQL Server 2017, you must configure SQL Server 2017 Reporting Services (SSRS) manually. After configuring SSRS, make sure to set the **IsInitialized** property of SSRS to **True**. When set to **True**, Azure Backup Server assumes that SSRS is already configured and skips the SSRS

configuration.

To check the SSRS configuration status, run:

```
$configset =Get-WmiObject -namespace  
"root\Microsoft\SqlServer\ReportServer\RS_SSRS\v14\Admin" -class  
MSReportServer_ConfigurationSetting -ComputerName localhost  
  
$configset.IsInitialized
```

Use the following values for SSRS configuration:

- **Service Account:** Use built-in account should be **Network Service**.
- **Web Service URL:** Virtual Directory should be **ReportServer_<SQLInstanceName>**.
- **Database:** DatabaseName should be **ReportServer\$<SQLInstanceName>**.
- **Web Portal URL:** Virtual Directory should be **Reports_<SQLInstanceName>**.

[Learn more](#) about SSRS configuration.

NOTE

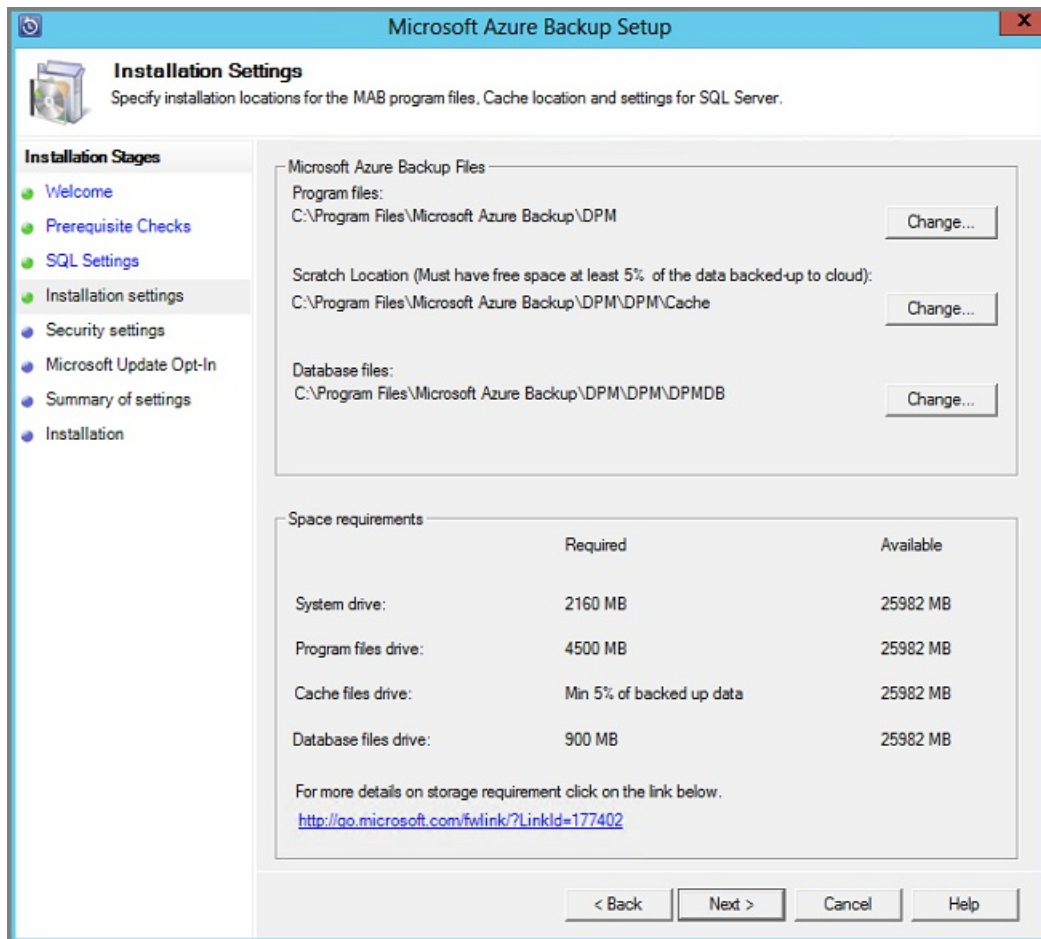
[Microsoft Online Services Terms](#) (OST) governs the licensing for SQL Server used as the database for Azure Backup Server. According to OST, only use SQL Server bundled with Azure Backup Server as the database for Azure Backup Server.

5. After the installation is successful, select **Next**.

6. Provide a location for installing Microsoft Azure Backup Server files, and select **Next**.

NOTE

The scratch location is required for backup to Azure. Ensure the scratch location is at least 5% of the data planned for backing up to the cloud. For disk protection, separate disks need configuring after the installation finishes. For more information about storage pools, see [Configure storage pools and disk storage](#).



7. Provide a strong password for restricted local user accounts, and select **Next**.
8. Select whether you want to use Microsoft Update to check for updates, and select **Next**.

NOTE

We recommend having Windows Update redirect to Microsoft Update, which offers security and important updates for Windows and other products like Azure Backup Server.

9. Review the **Summary of Settings**, and select **Install**.

The installation happens in phases.

- The first phase installs the Microsoft Azure Recovery Services Agent.
 - The second phase checks for internet connectivity. If available, you can continue with the installation. If not available, you must provide proxy details to connect to the internet.
 - The final phase checks the prerequisite software. If not installed, any missing software gets installed along with the Microsoft Azure Recovery Services Agent.
10. Select **Browse** to locate your vault credentials to register the machine to the Recovery Services vault, and then select **Next**.
 11. Select a passphrase to encrypt or decrypt the data sent between Azure and your premises.

TIP

You can automatically generate a passphrase or provide your minimum 16-character passphrase.

12. Enter the location to save the passphrase, and then select **Next** to register the server.

IMPORTANT

Save the passphrase to a safe location other than the local server. We strongly recommend using the Azure Key Vault to store the passphrase.

After the Microsoft Azure Recovery Services Agent setup finishes, the installation step moves on to the installation and configuration of SQL Server and the Azure Backup Server components.

13. After the installation step finishes, select **Close**.

Install Update Rollup 1

Installing the Update Rollup 1 for Azure Backup Server v3 is mandatory before you can protect the workloads. You can find the bug fixes and installation instructions in the [knowledge base article](#).

Add storage to Azure Backup Server

Azure Backup Server v3 supports Modern Backup Storage that offers:

- Storage savings of 50%.
- Backups that are three times faster.
- More efficient storage.
- Workload-aware storage.

Volumes in Azure Backup Server

Add the data disks with the Azure Backup Server VM's required storage capacity if not already added.

Azure Backup Server v3 only accepts storage volumes. When you add a volume, Azure Backup Server formats the volume to Resilient File System (ReFS), which Modern Backup Storage requires.

Add volumes to Azure Backup Server disk storage

1. In the **Management** pane, rescan the storage and then select **Add**.
2. Select from the available volumes to add to the storage pool.
3. After you add the available volumes, give them a friendly name to help you manage them.
4. Select **OK** to format these volumes to ReFS so that Azure Backup Server can use Modern Backup Storage benefits.

Next steps

Now that you've covered how to set up Azure Backup Server for Azure VMware Solution, you can use the following resources to learn more.

- [Configuring backups for your Azure VMware Solution VMs](#).
- [Protecting your Azure VMware Solution VMs with Microsoft Defender for Cloud integration](#).

Back up Azure VMware Solution VMs with Azure Backup Server

12/16/2022 • 12 minutes to read • [Edit Online](#)

This article shows you how to back up VMware virtual machines (VMs) running on Azure VMware Solution with Azure Backup Server. First, thoroughly go through [Set up Microsoft Azure Backup Server for Azure VMware Solution](#).

Then, we'll walk through all of the necessary procedures to:

- Set up a secure channel so that Azure Backup Server can communicate with VMware vCenter Server over HTTPS.
- Add the account credentials to Azure Backup Server.
- Add the vCenter Server to Azure Backup Server.
- Set up a protection group that contains the VMware vSphere VMs you want to back up, specify backup settings, and schedule the backup.

Create a secure connection to the vCenter Server

By default, Azure Backup Server communicates with VMware vCenter Server over HTTPS. To set up the HTTPS connection, download the VMware certificate authority (CA) certificate and import it on the Azure Backup Server.

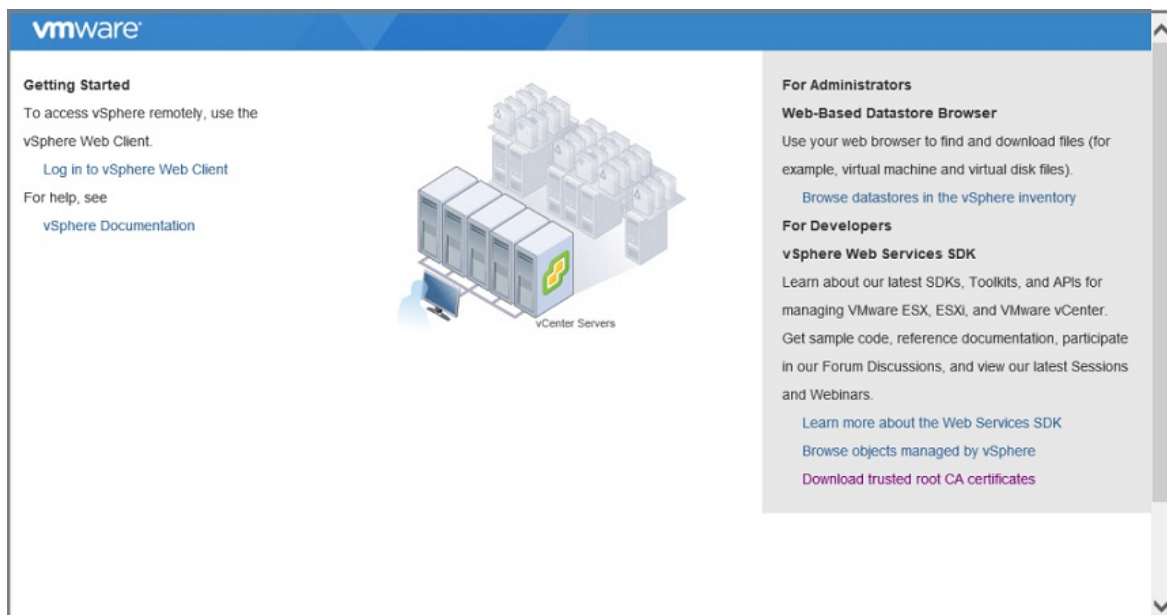
Set up the certificate

1. In the browser, on the Azure Backup Server machine, enter the vSphere Client URL.

NOTE

If the VMware vSphere Client **Getting Started** page doesn't appear, verify the connection and browser proxy settings and try again.

2. On the VMware vSphere Client **Getting Started** page, select **Download trusted root CA certificates**.



3. Save the **download.zip** file to the Azure Backup Server machine, and then extract its contents to the **certs** folder, which contains the:
 - Root certificate file with an extension that begins with a numbered sequence like .0 and .1.
 - CRL file with an extension that begins with a sequence like .r0 or .r1.
4. In the **certs** folder, right-click the root certificate file and select **Rename** to change the extension to **.crt**.

The file icon changes to one that represents a root certificate.

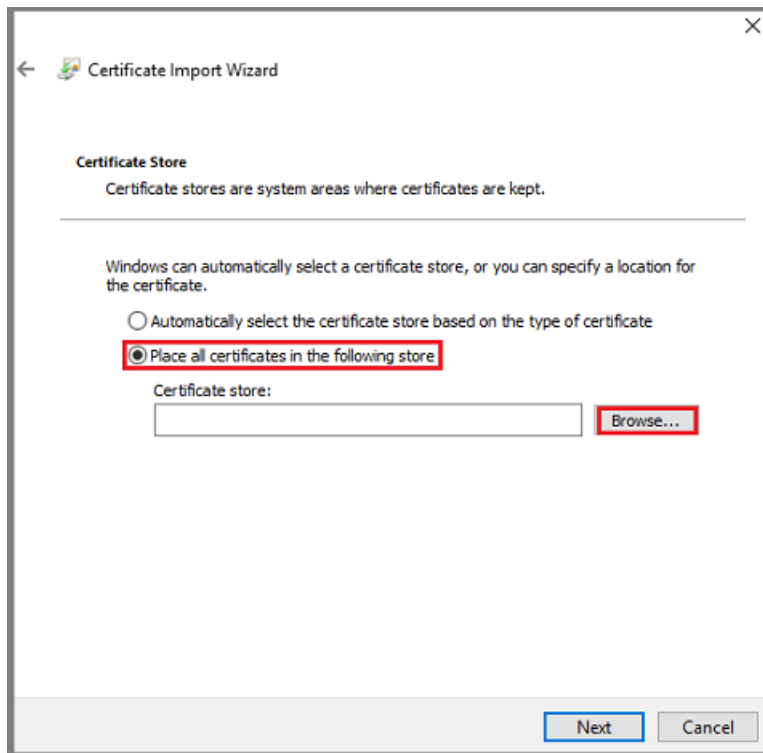
5. Right-click the root certificate, and select **Install Certificate**.
6. In the **Certificate Import Wizard**, select **Local Machine** as the destination for the certificate, and select **Next**.



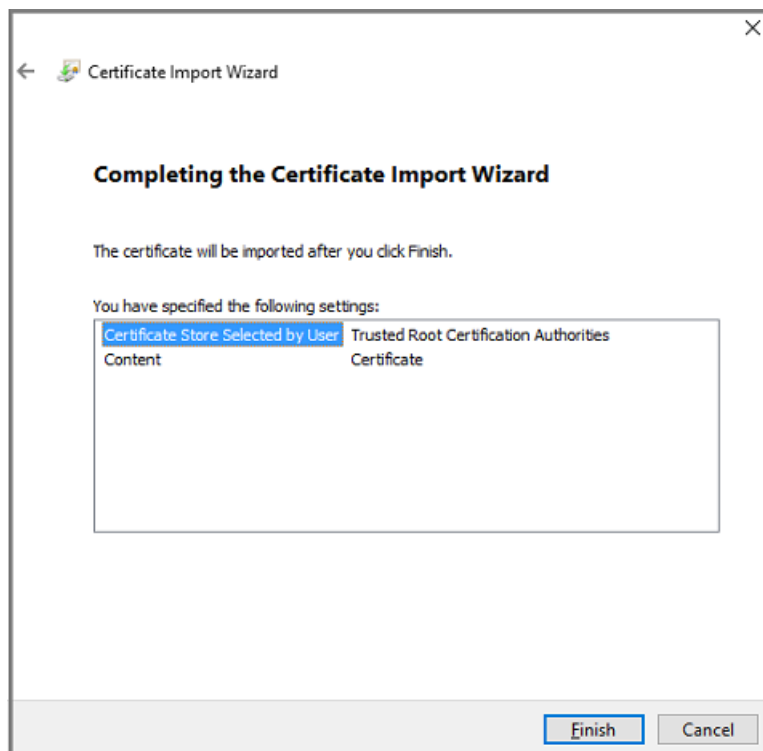
NOTE

If asked, confirm that you want to allow changes to the computer.

7. Select **Place all certificates in the following store**, and select **Browse** to choose the certificate store.



8. Select **Trusted Root Certification Authorities** as the destination folder, and select **OK**.
9. Review the settings, and select **Finish** to start importing the certificate.



10. After the certificate import is confirmed, sign in to the vCenter Server to confirm that your connection is secure.

Enable TLS 1.2 on Azure Backup Server

VMware vSphere 6.7 onwards has TLS enabled as the communication protocol.

1. Copy the following registry settings, and paste them into Notepad. Then save the file as TLS.REG without the .txt extension.

Windows Registry Editor Version 5.00

[HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\Microsoft\.NETFramework\v2.0.50727]

"SystemDefaultTlsVersions"=dword:00000001

"SchUseStrongCrypto"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\Microsoft\.NETFramework\v4.0.30319]

"SystemDefaultTlsVersions"=dword:00000001

"SchUseStrongCrypto"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\.NETFramework\v2.0.50727]

"SystemDefaultTlsVersions"=dword:00000001

"SchUseStrongCrypto"=dword:00000001

[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\.NETFramework\v4.0.30319]

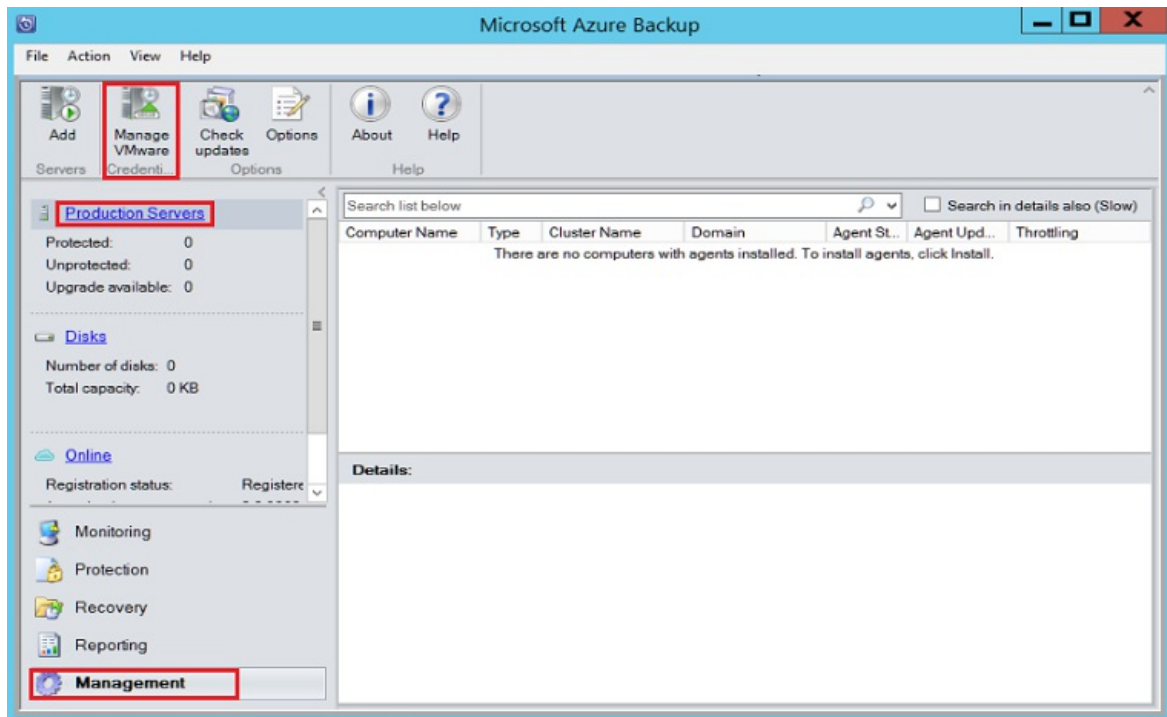
"SystemDefaultTlsVersions"=dword:00000001

"SchUseStrongCrypto"=dword:00000001

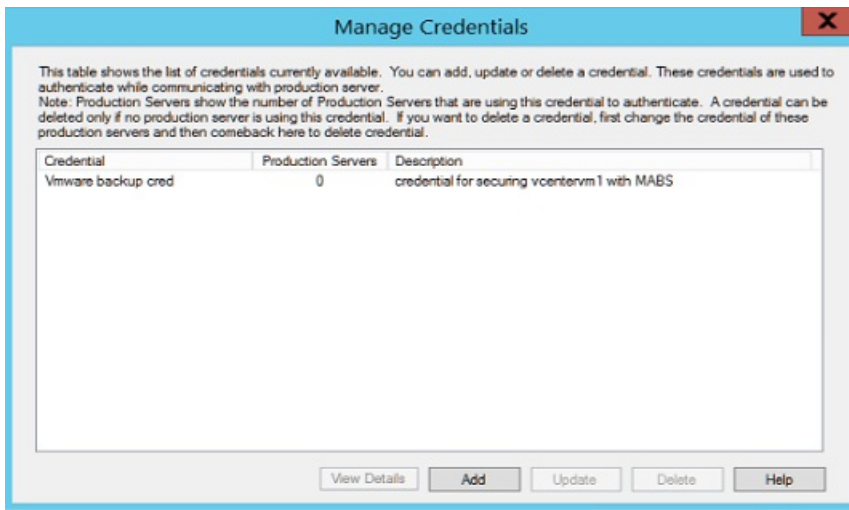
2. Right-click the TLS.REG file, and select **Merge** or **Open** to add the settings to the registry.

Add the account on Azure Backup Server

1. Open Azure Backup Server, and in the Azure Backup Server console, select **Management** > **Production Servers** > **Manage VMware**.

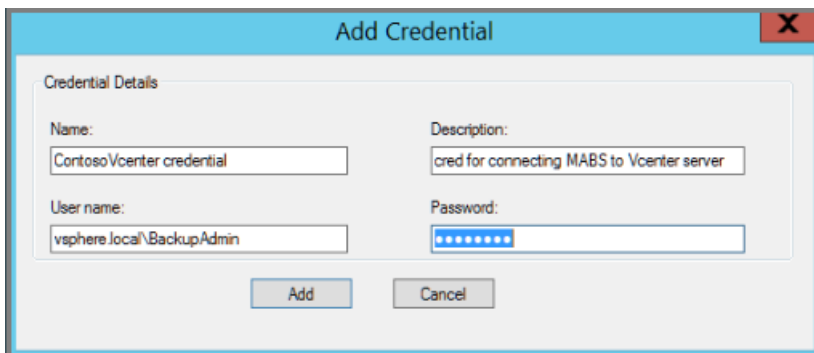


2. In the **Manage Credentials** dialog box, select **Add**.

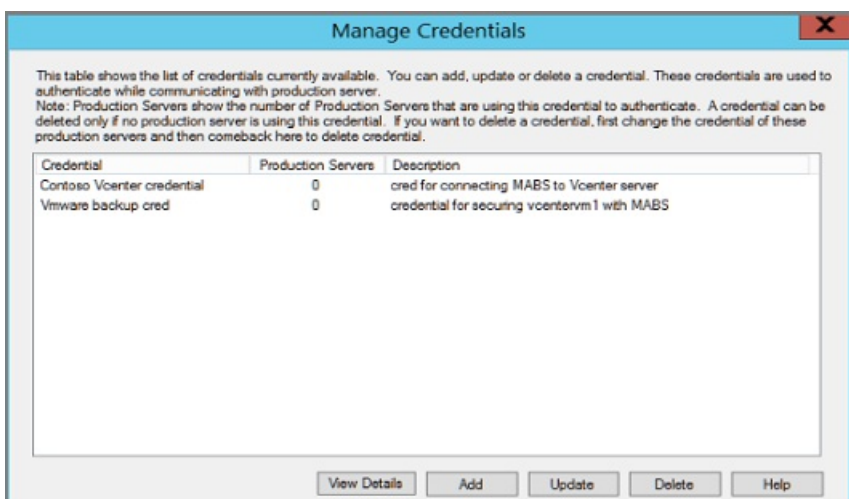


- In the **Add Credential** dialog box, enter a name and a description for the new credential. Specify the user name and password you defined on the VMware server.

NOTE
 If the VMware vSphere virtual machine and Azure Backup Server aren't in the same domain, specify the domain in the User name box.

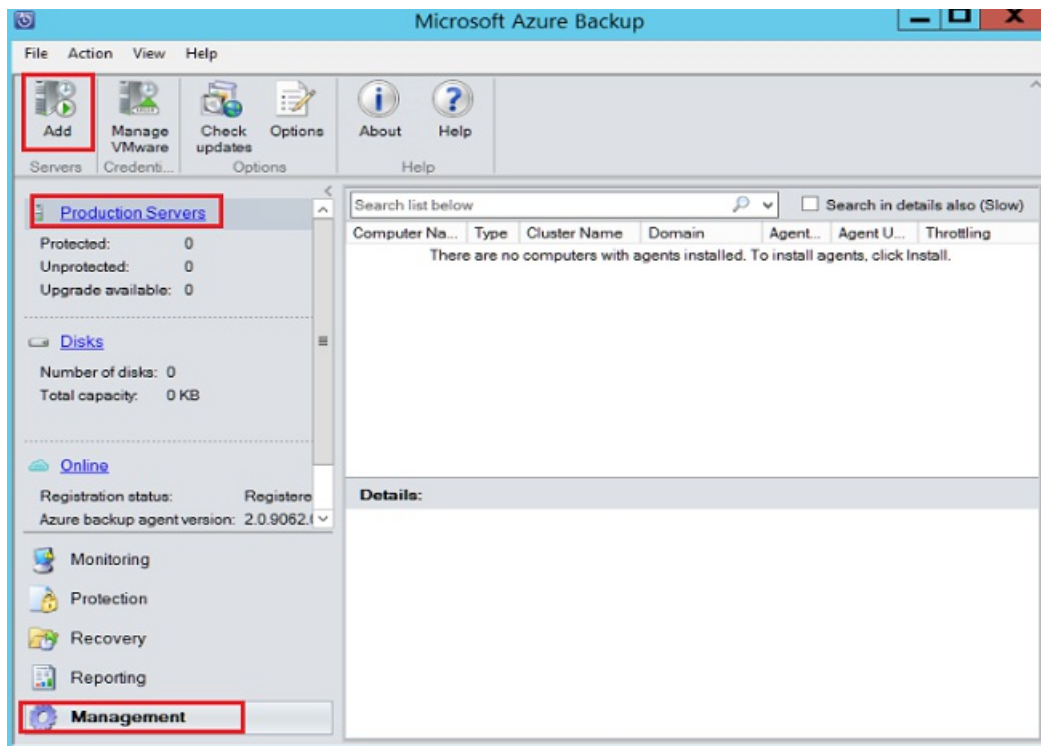


- Select **Add** to add the new credential.

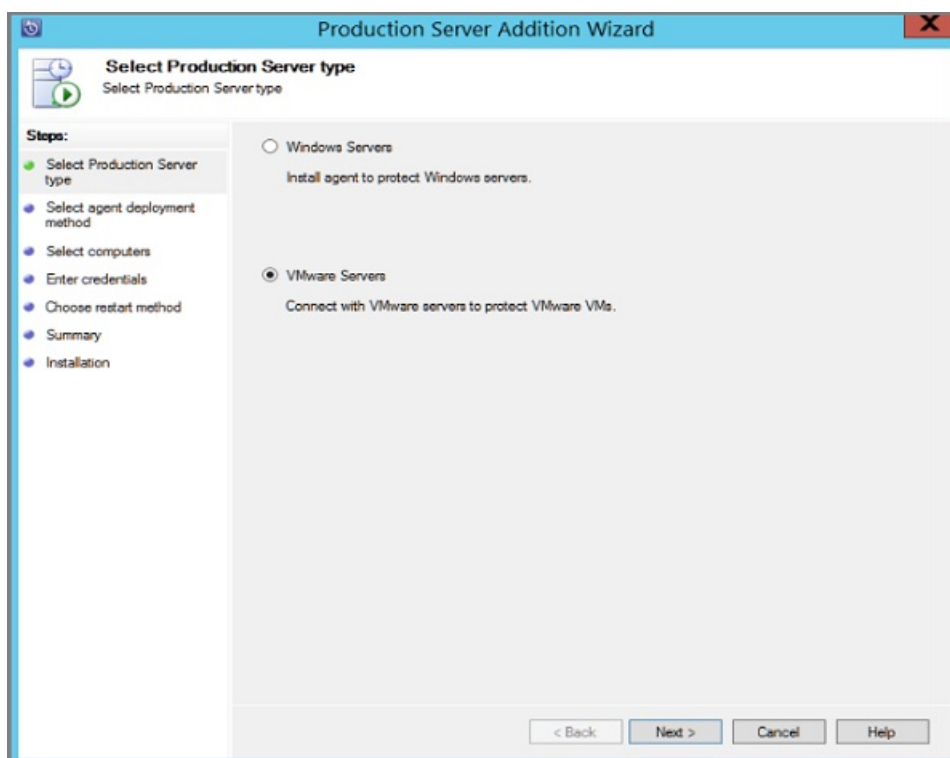


Add the vCenter Server to Azure Backup Server

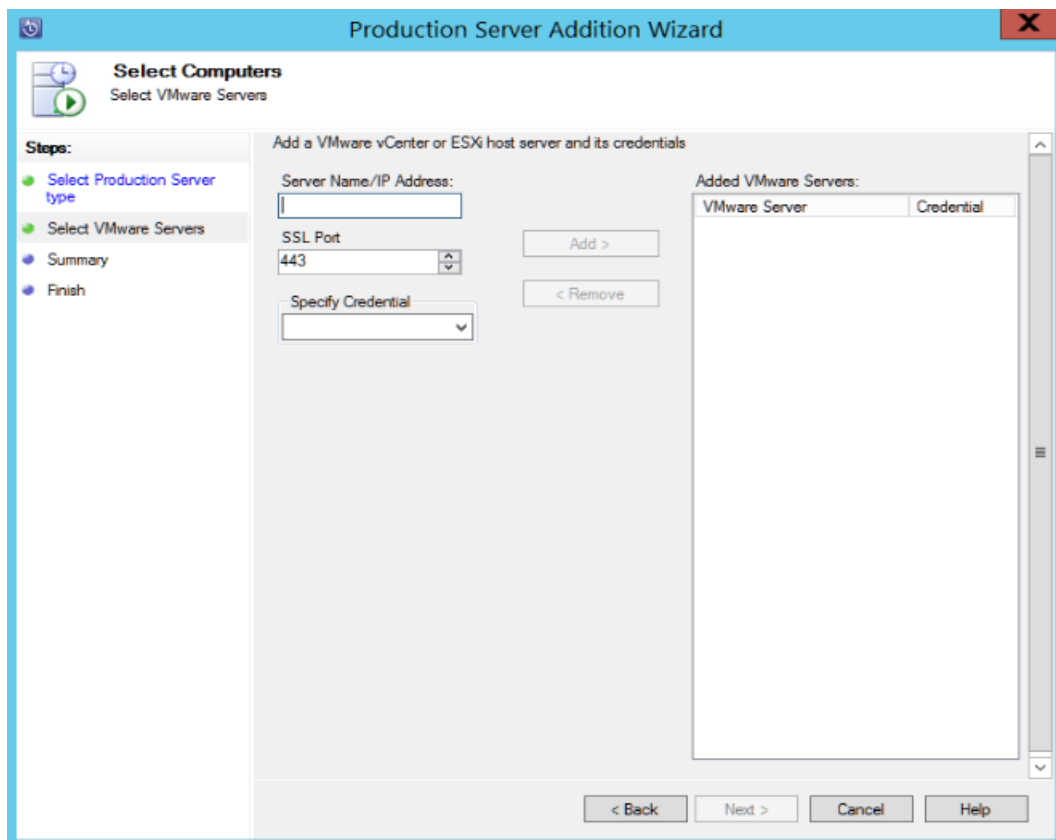
- In the Azure Backup Server console, select **Management > Production Servers > Add**.



2. Select VMware Servers, and select Next.



3. Specify the IP address of the vCenter Server.

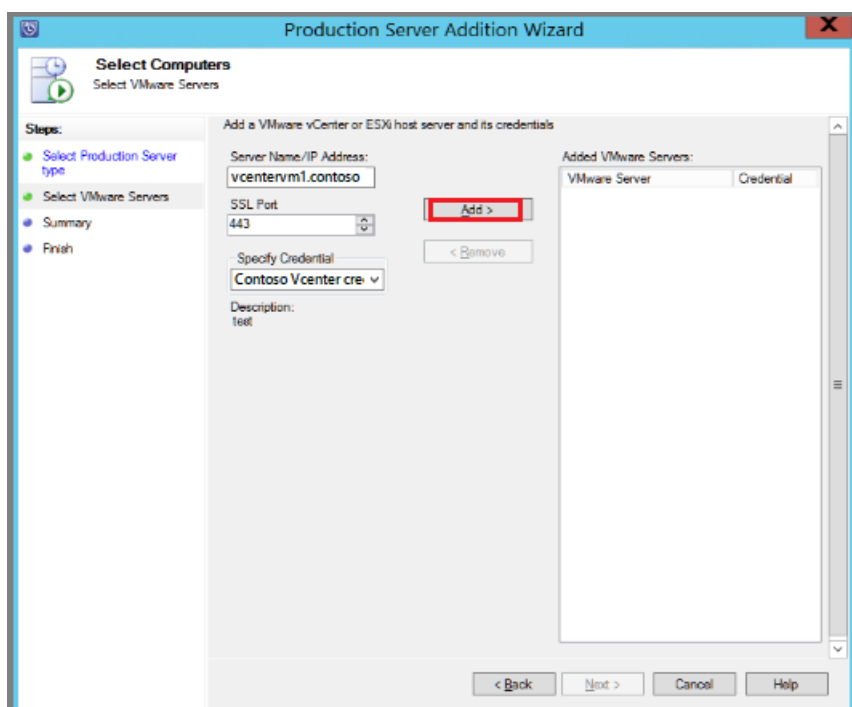


- In the **SSL Port** box, enter the port used to communicate with the vCenter Server.

TIP

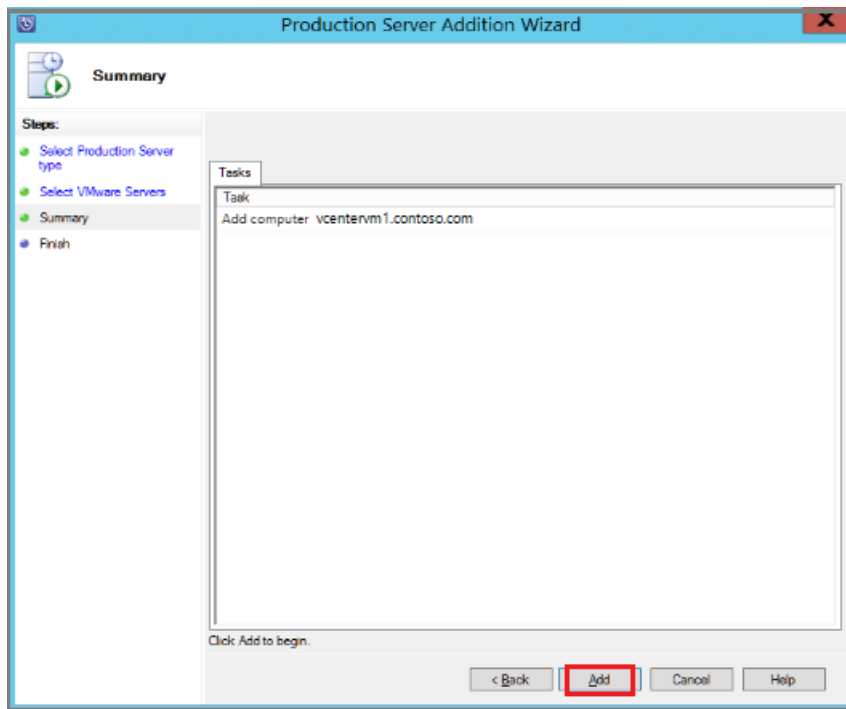
Port 443 is the default port, but you can change it if your vCenter Server listens on a different port.

- In the **Specify Credential** box, select the credential that you created in the previous section.
- Select **Add** to add the vCenter Server to the servers list, and select **Next**.

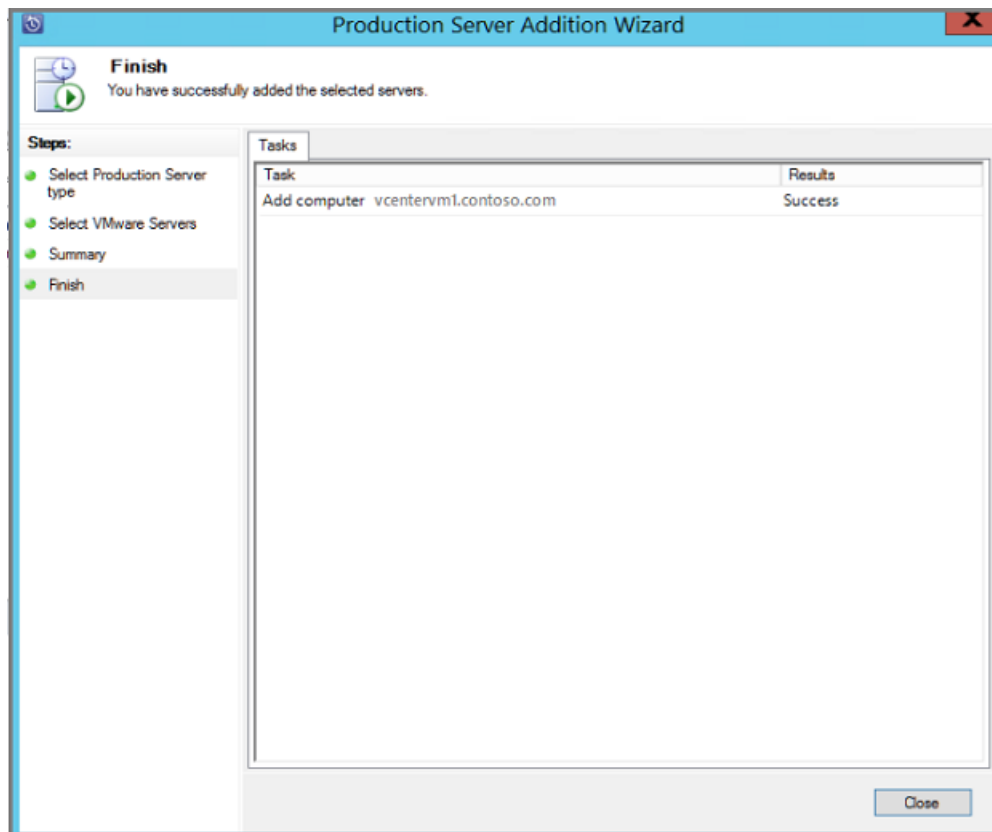


- On the **Summary** page, select **Add** to add the vCenter Server to Azure Backup Server.

The new vCenter Server gets added immediately. vCenter Server doesn't need an agent.



8. On the **Finish** page, review the settings, and then select **Close**.



You see the vCenter Server listed under **Production Server** with:

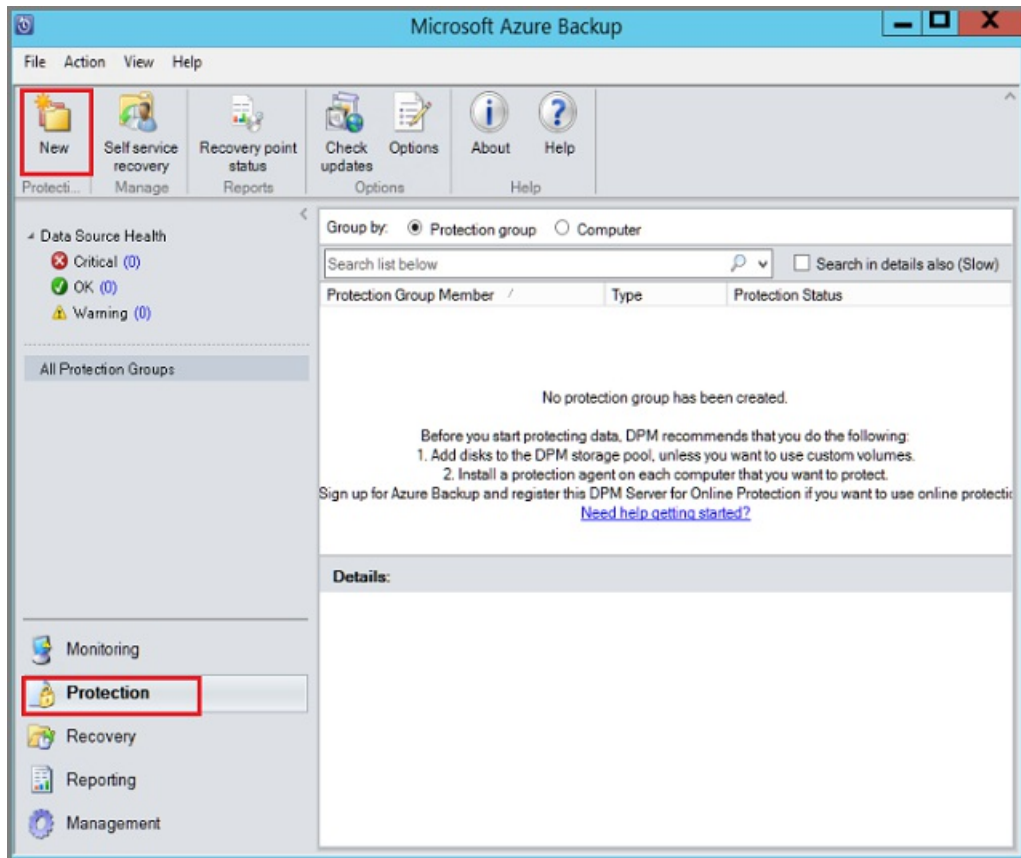
- Type as **VMware Server**
- Agent Status as **OK**

If you see **Agent Status** as **Unknown**, select **Refresh**.

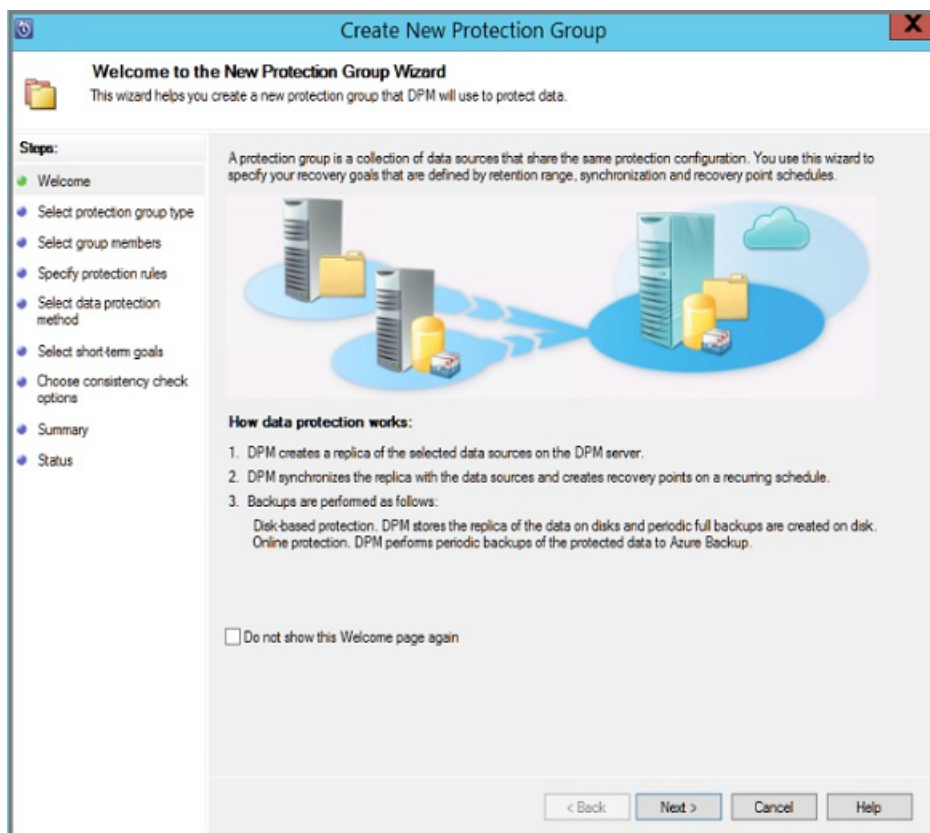
Configure a protection group

Protection groups gather multiple VMs and apply the same data retention and backup settings to all VMs in the group.

1. In the Azure Backup Server console, select **Protection** > **New**.



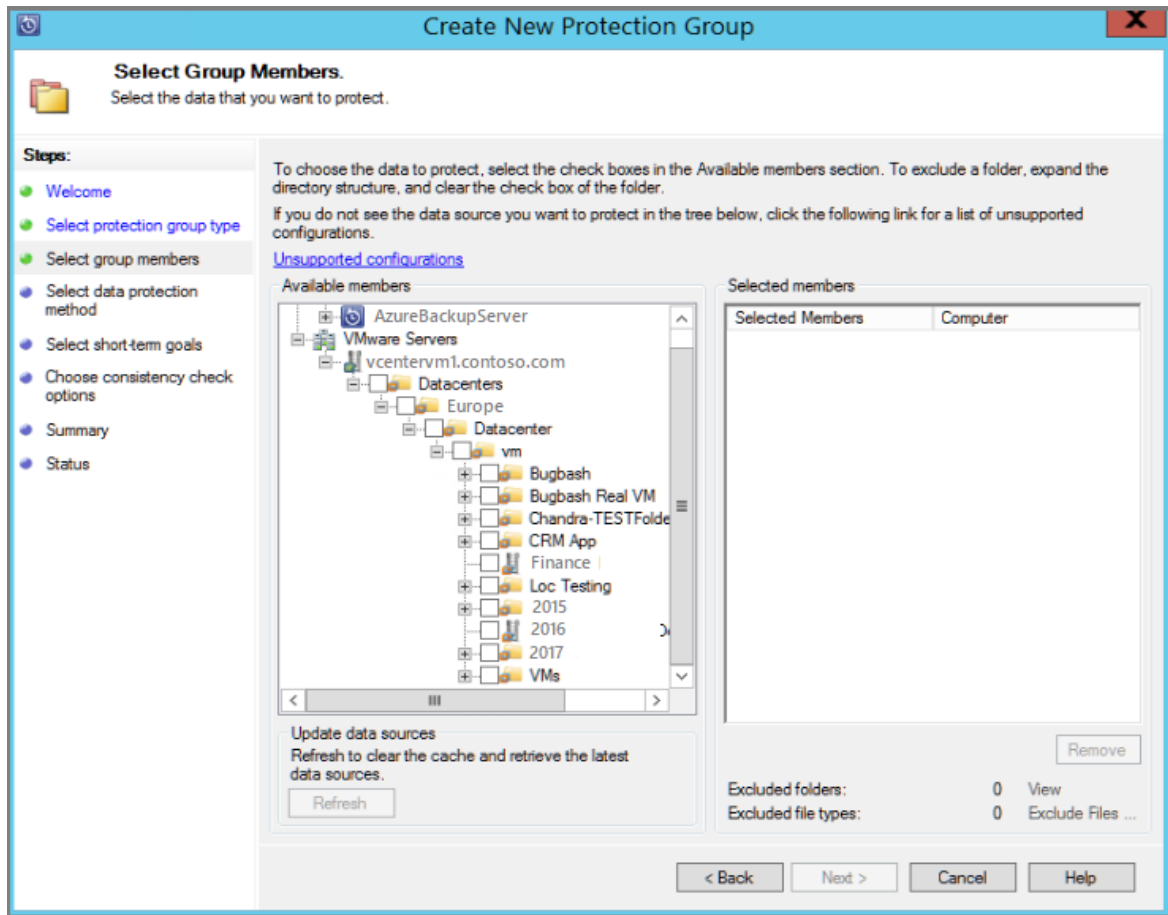
2. On the Create New Protection Group wizard welcome page, select **Next**.



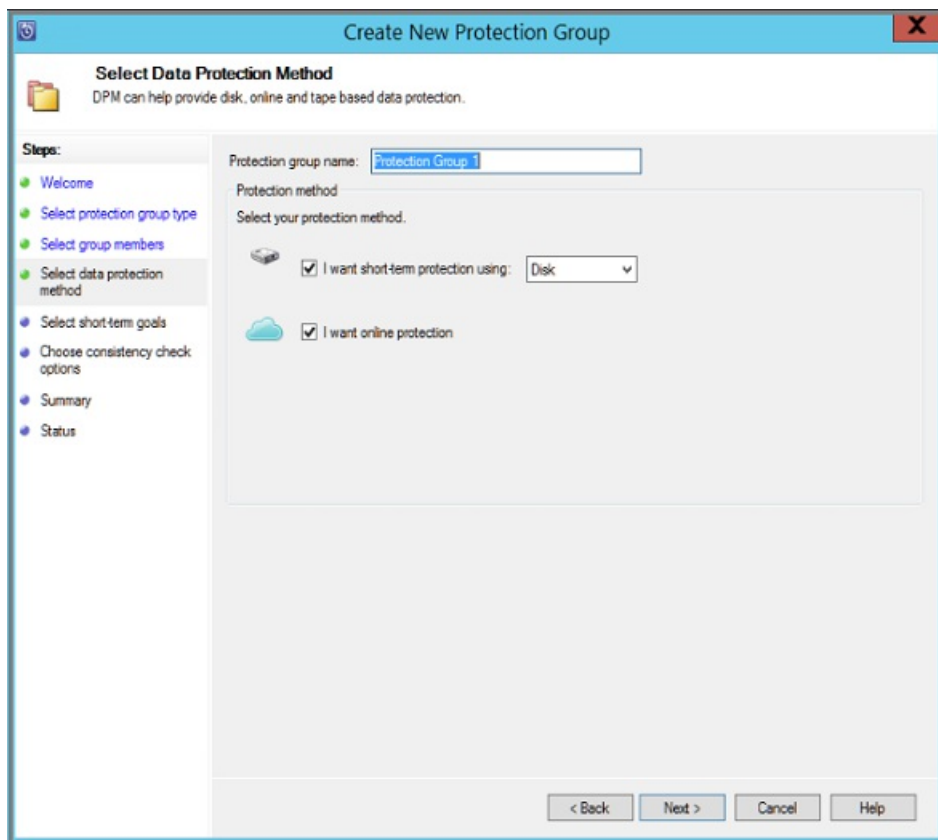
3. On the **Select Protection Group Type** page, select **Servers**, and then select **Next**. The **Select Group Members** page appears.
4. On the **Select Group Members** page, select the VMs (or VM folders) that you want to back up, and then select **Next**.

NOTE

When you select a folder or VMs, folders inside that folder are also selected for backup. You can uncheck folders or VMs you don't want to back up. If a VM or folder is already being backed up, you can't select it, which ensures duplicate recovery points aren't created for a VM.

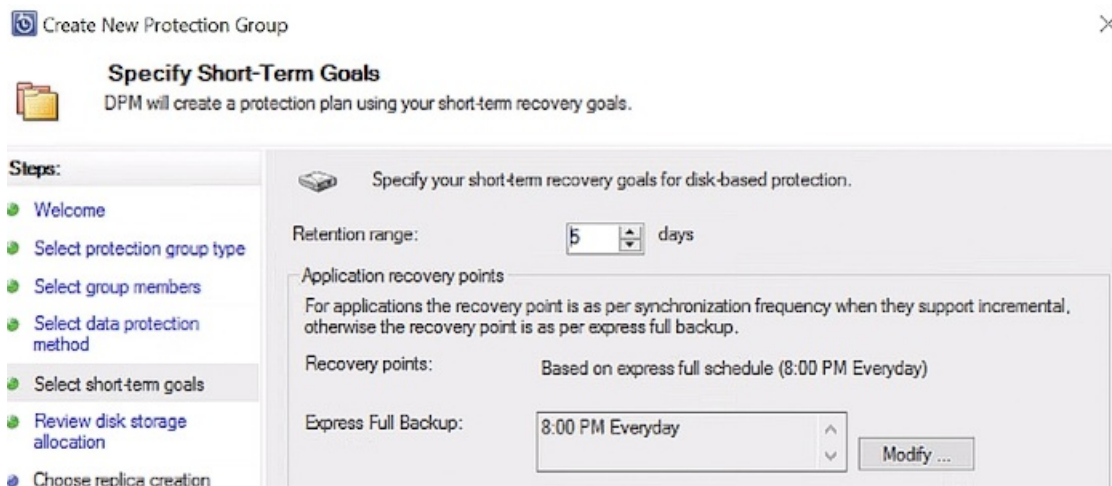


5. On the **Select Data Protection Method** page, enter a name for the protection group and protection settings.
6. Set the short-term protection to **Disk**, enable online protection, and then select **Next**.



7. Specify how long you want to keep data backed up to disk.

- **Retention range:** The number of days that disk recovery points are kept.
- **Express Full Backup:** How often disk recovery points are taken. To change the times or dates when short-term backups occur, select **Modify**.



8. On the **Review Disk Storage Allocation** page, review the disk space provided for the VM backups.

- The recommended disk allocations are based on the retention range you specified, the workload type, and the protected data size. Make any changes required, and then select **Next**.
- **Data size:** Size of the data in the protection group.
- **Disk space:** Recommended amount of disk space for the protection group. If you want to modify this setting, select space lightly larger than the amount you estimate each data source grows.
- **Storage pool details:** Shows the status of the storage pool, which includes total and remaining disk size.

Review Disk Storage Allocation

Review disk space allocated in the storage pool for this protection group.

Steps:

- Welcome
- Select protection group type
- Select group members
- Select data protection method
- Select short-term goals
- **Review disk storage allocation**
- Choose replica creation method
- Choose consistency check options
- Specify online protection data
- Specify online backup schedule
- Specify online retention policy
- Choose online replication
- Summary
- Status

Review target storage assigned for each data source and change if need be.

Disk storage allocation for new members

Total data size:	0.01 GB
Disk storage to be provisioned on DPM:	10.00 MB

Disk storage allocation details:

Data Source /	Data Size	Space To ...	Target Storage
FC-MABS-1MSDPMINSTANCE\master on...	0.01 GB	10.00 MB	Disk - 2,027.91 GB

Available target disk storage:

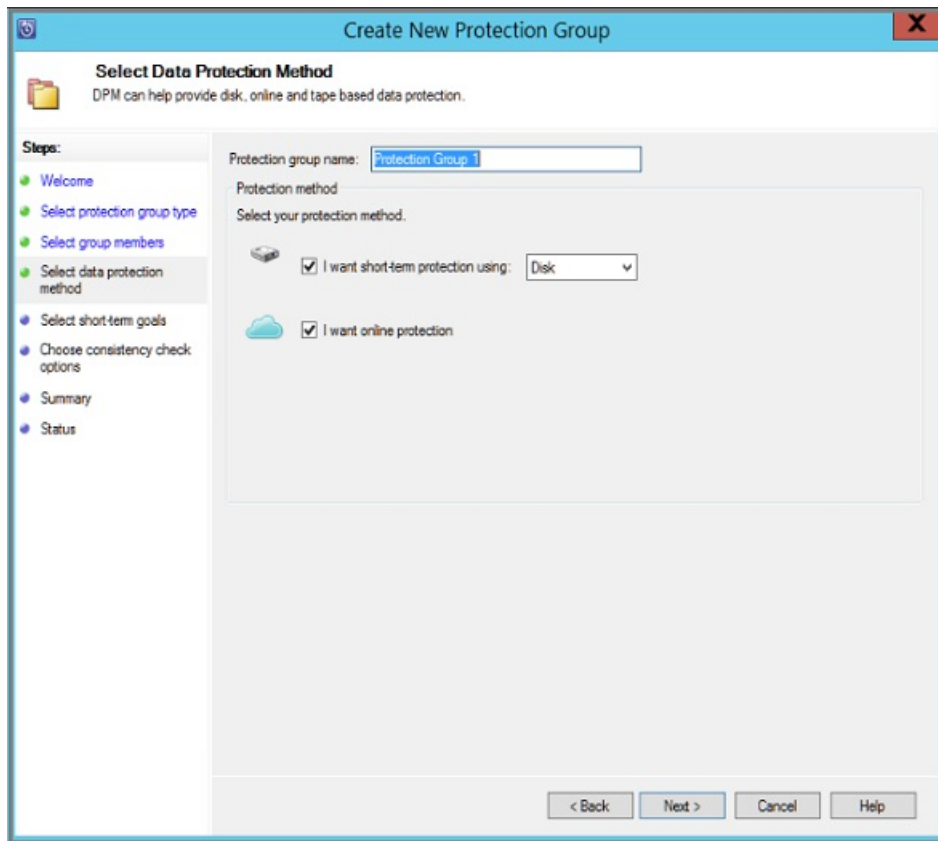
Name /	Friendly N...	Allowed Dataso...	Total Spa...	Free Space	Underpro...
E:\	Disk	All	2,047.93 GB	2,027.91 GB	0 KB

< Back **Next >** Cancel Help

NOTE

In some scenarios, the data size reported is higher than the actual VM size. We're aware of the issue and currently investigating it.

9. On the **Choose Replica Creation Method** page, indicate how you want to take the initial backup, and select **Next**.
 - The default is **Automatically over the network** and **Now**. If you use the default, specify an off-peak time. If you choose **Later**, specify a day and time.
 - For large amounts of data or less-than-optimal network conditions, consider replicating the data offline by using removable media.



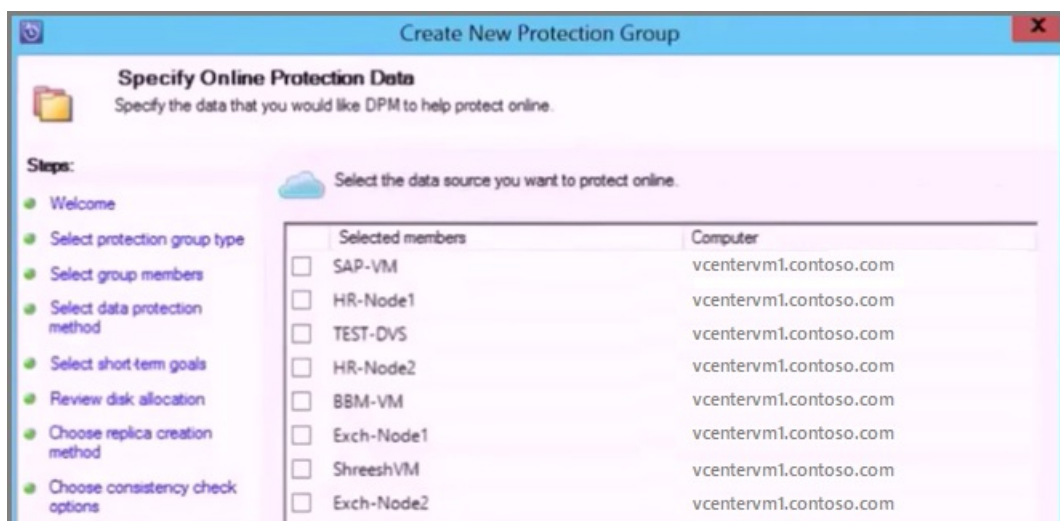
10. For **Consistency check options**, select how and when to automate the consistency checks and select **Next**.

- You can run consistency checks when replica data becomes inconsistent, or on a set schedule.
- If you don't want to configure automatic consistency checks, you can run a manual check by right-clicking the protection group **Perform Consistency Check**.

11. On the **Specify Online Protection Data** page, select the VMs or VM folders that you want to back up, and then select **Next**.

TIP

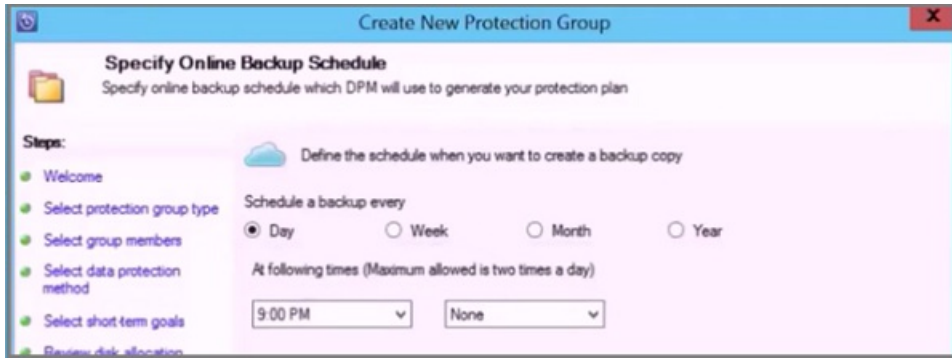
You can select the members individually or choose **Select All** to choose all members.



12. On the **Specify Online Backup Schedule** page, indicate how often you want to back up data from local storage to Azure.

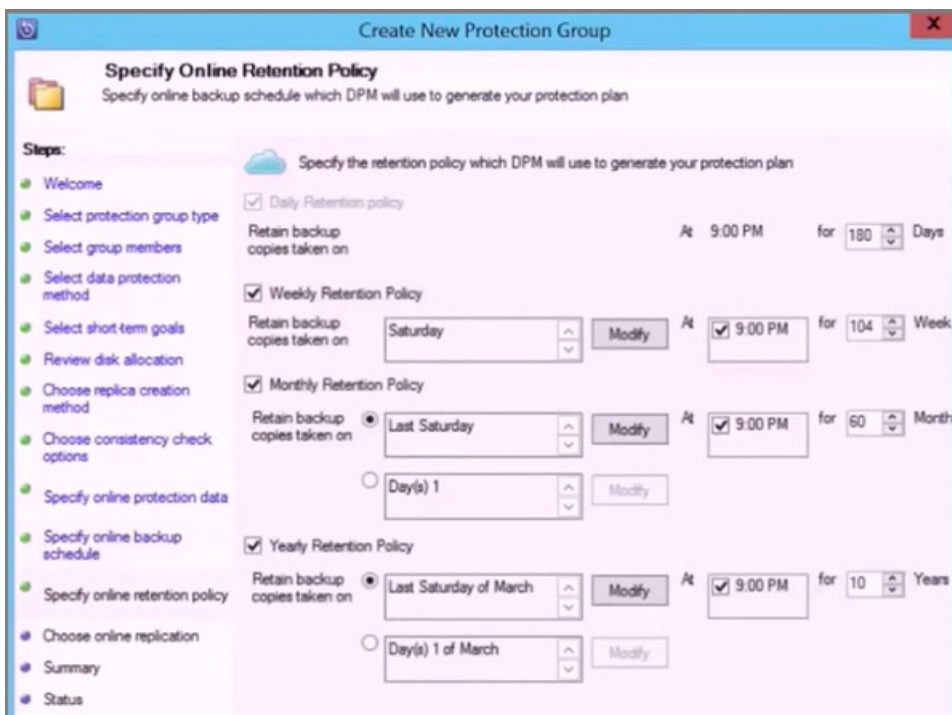
- Cloud recovery points for the data to get generated according to the schedule.

- After the recovery point gets generated, it's then transferred to the Recovery Services vault in Azure.

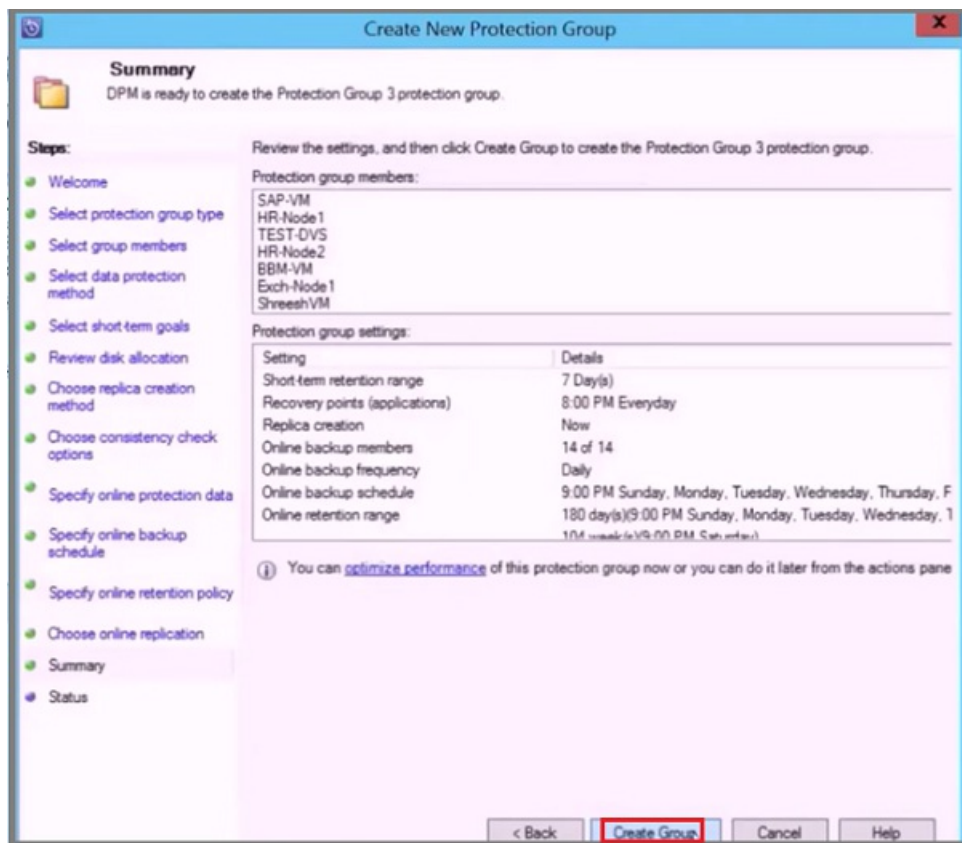


13. On the **Specify Online Retention Policy** page, indicate how long you want to keep the recovery points created from the backups to Azure.

- There's no time limit for how long you can keep data in Azure.
- The only limit is that you can't have more than 9,999 recovery points per protected instance. In this example, the protected instance is the VMware vCenter Server.



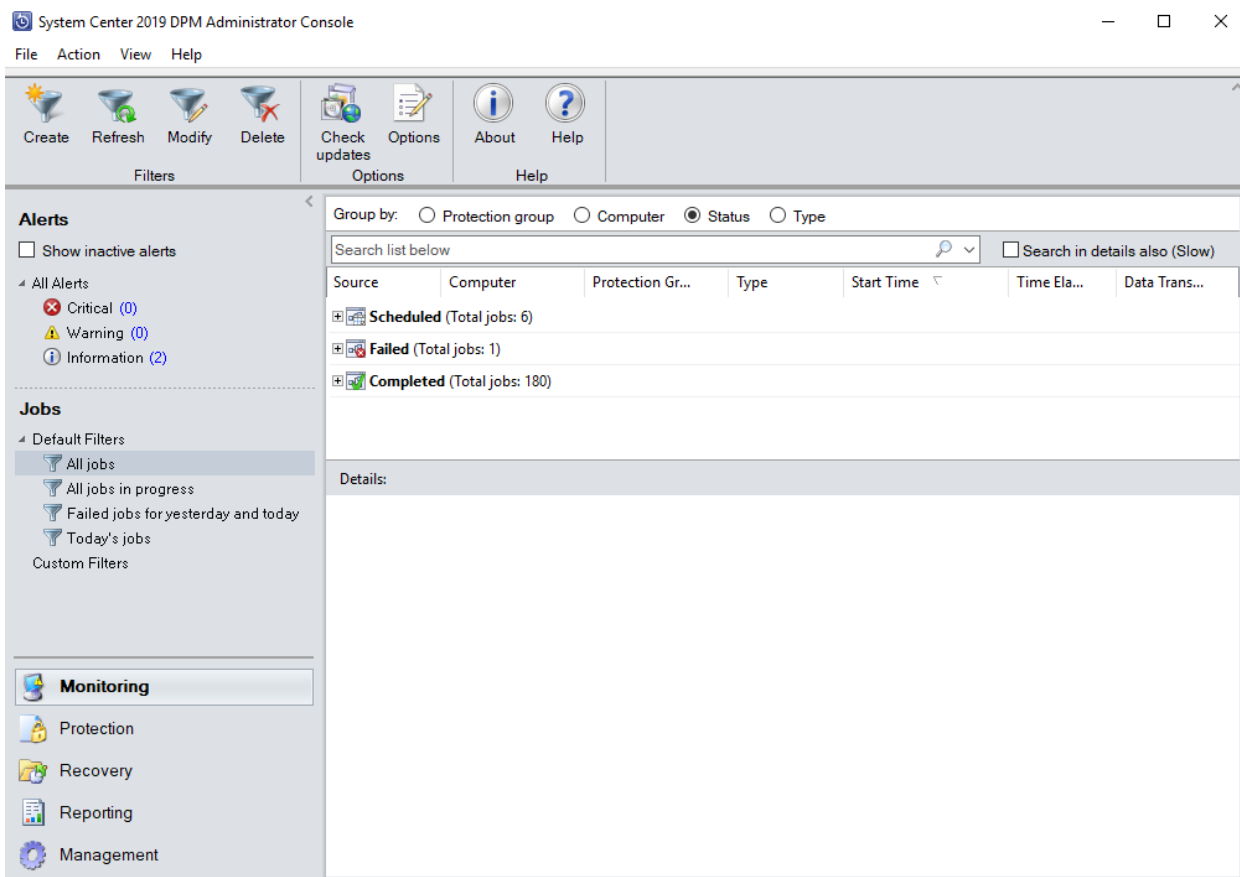
14. On the **Summary** page, review the settings and then select **Create Group**.



Monitor with the Azure Backup Server console

After you configure the protection group to back up Azure VMware Solution VMs, you can monitor the backup job status and alert by using the Azure Backup Server console. Here's what you can monitor.

- In the **Monitoring** task area:
 - Under **Alerts**, you can monitor errors, warnings, and general information. You can view active and inactive alerts and set up email notifications.
 - Under **Jobs**, you can view jobs started by Azure Backup Server for a specific protected data source or protection group. You can follow job progress or check resources consumed by jobs.
- In the **Protection** task area, you can check the status of volumes and shares in the protection group. You can also check configuration settings such as recovery settings, disk allocation, and the backup schedule.
- In the **Management** task area, you can view the **Disks**, **Online**, and **Agents** tabs to check the status of disks in the storage pool, registration to Azure, and deployed DPM agent status.



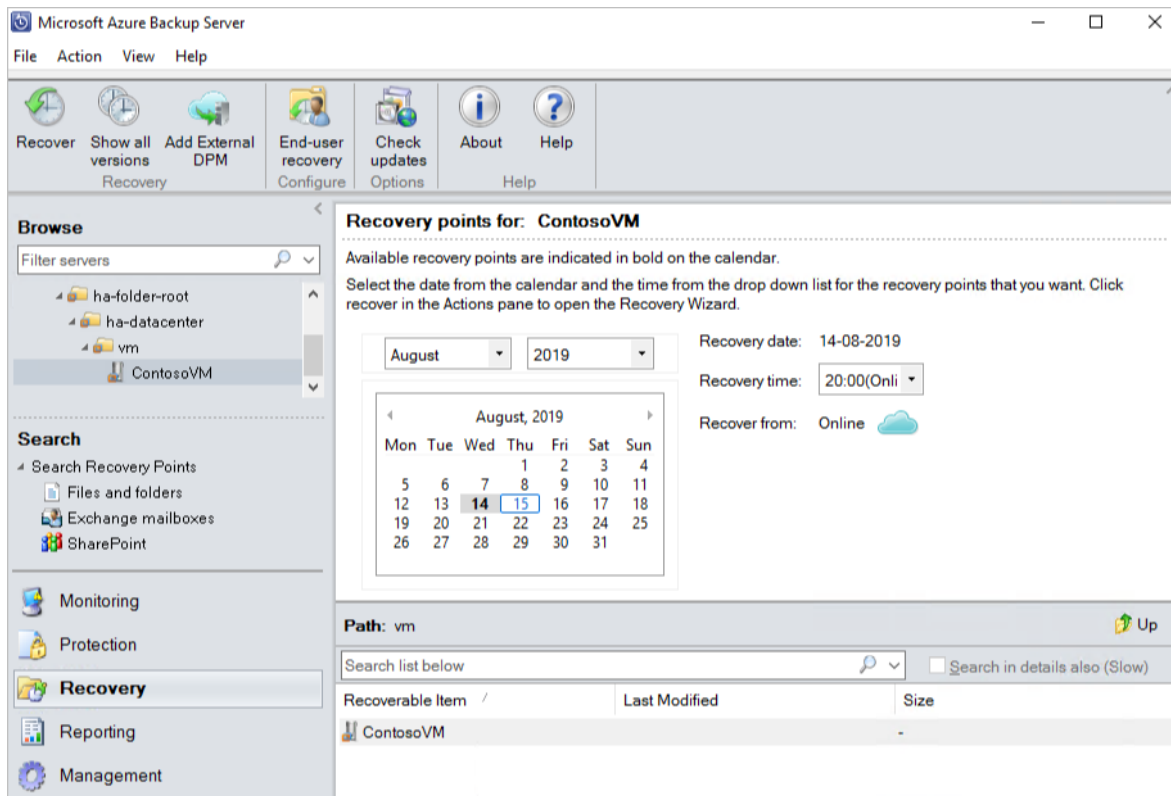
Restore VMware vSphere virtual machines

In the Azure Backup Server Administrator Console, there are two ways to find recoverable data. You can search or browse. When you recover data, you might or might not want to restore data or a VM to the same location. For this reason, Azure Backup Server supports three recovery options for VMware VM backups:

- **Original location recovery (OLR):** Use OLR to restore a protected VM to its original location. You can restore a VM to its original location only if no disks were added or deleted since the backup occurred. If disks were added or deleted, you must use alternate location recovery.
- **Alternate location recovery (ALR):** Use when the original VM is missing, or you don't want to disturb the original VM. Provide the location of an ESXi host, resource pool, folder, and the storage datastore and path. To help differentiate the restored VM from the original VM, Azure Backup Server appends *"-Recovered"* to the name of the VM.
- **Individual file location recovery (ILR):** If the protected VM is a Windows Server VM, individual files or folders inside the VM can be recovered by using the ILR capability of Azure Backup Server. To recover individual files, see the procedure later in this article. Restoring an individual file from a VM is available only for Windows VM and disk recovery points.

Restore a recovery point

1. In the Azure Backup Server Administrator Console, select the **Recovery** view.
2. Using the **Browse** pane, browse or filter to find the VM you want to recover. After you select a VM or folder, the **Recovery points for pane** display the available recovery points.

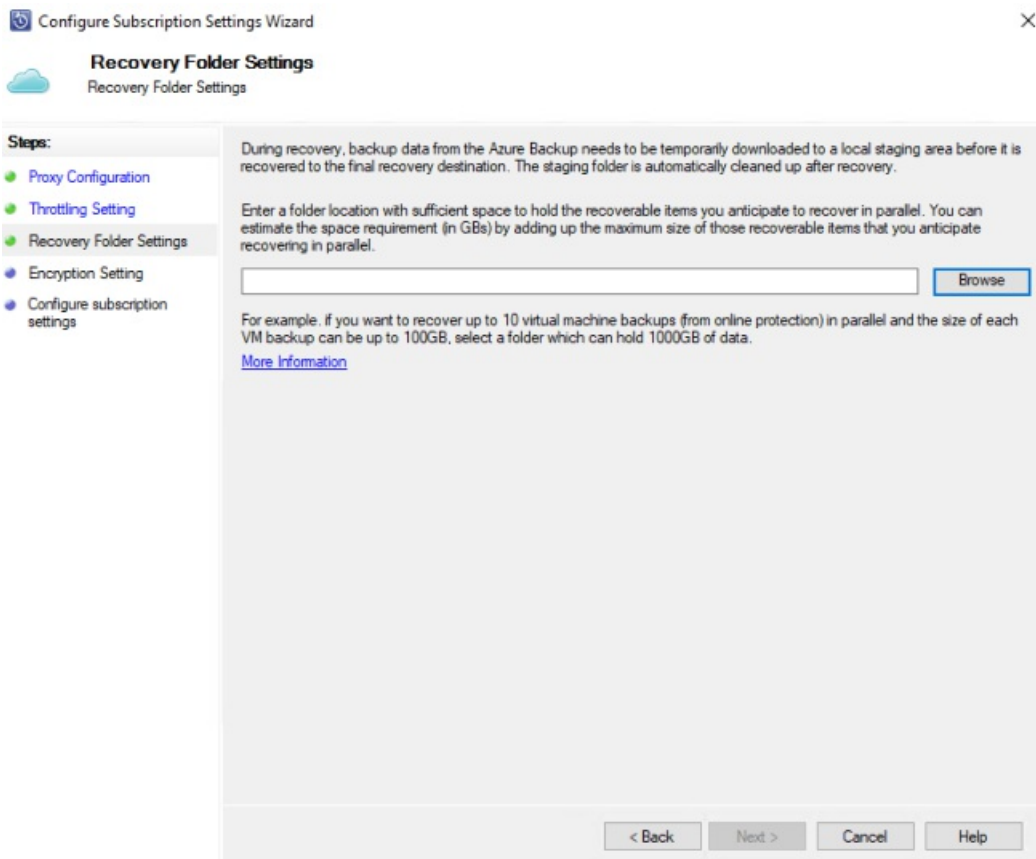


3. In the **Recovery points for** pane, select a date when a recovery point was taken. For example, calendar dates in bold have available recovery points. Alternately, you can right-click the VM, select **Show all recovery points**, and then select the recovery point from the list.

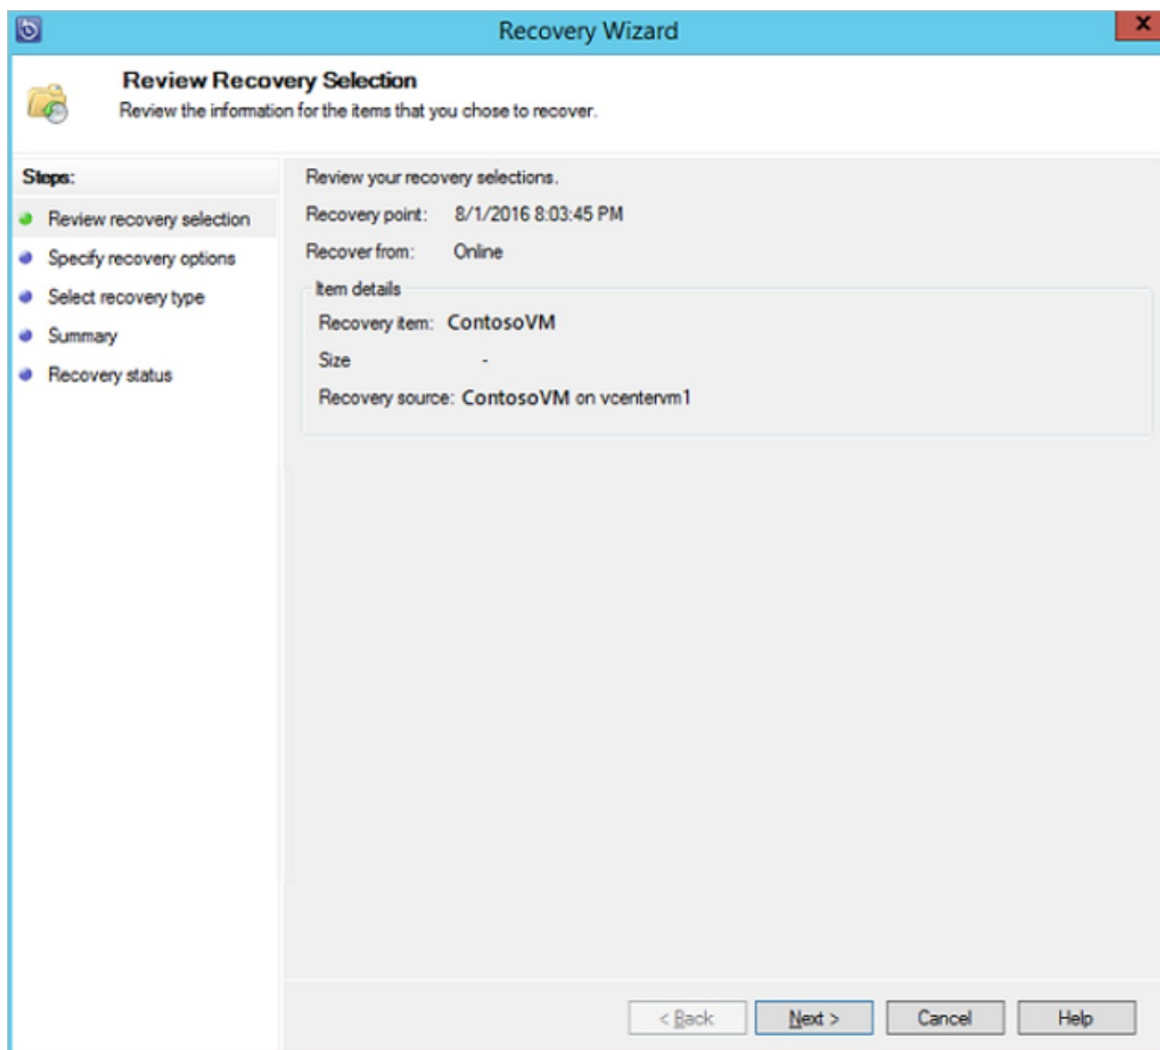
NOTE

For short-term protection, select a disk-based recovery point for faster recovery. After short-term recovery points expire, you see only **Online** recovery points to recover.

4. Before recovering from an online recovery point, ensure the staging location contains enough free space to house the full uncompressed size of the VM you want to recover. The staging location can be viewed or changed by running the **Configure Subscription Settings Wizard**.



5. Select **Recover** to open the **Recovery Wizard**.



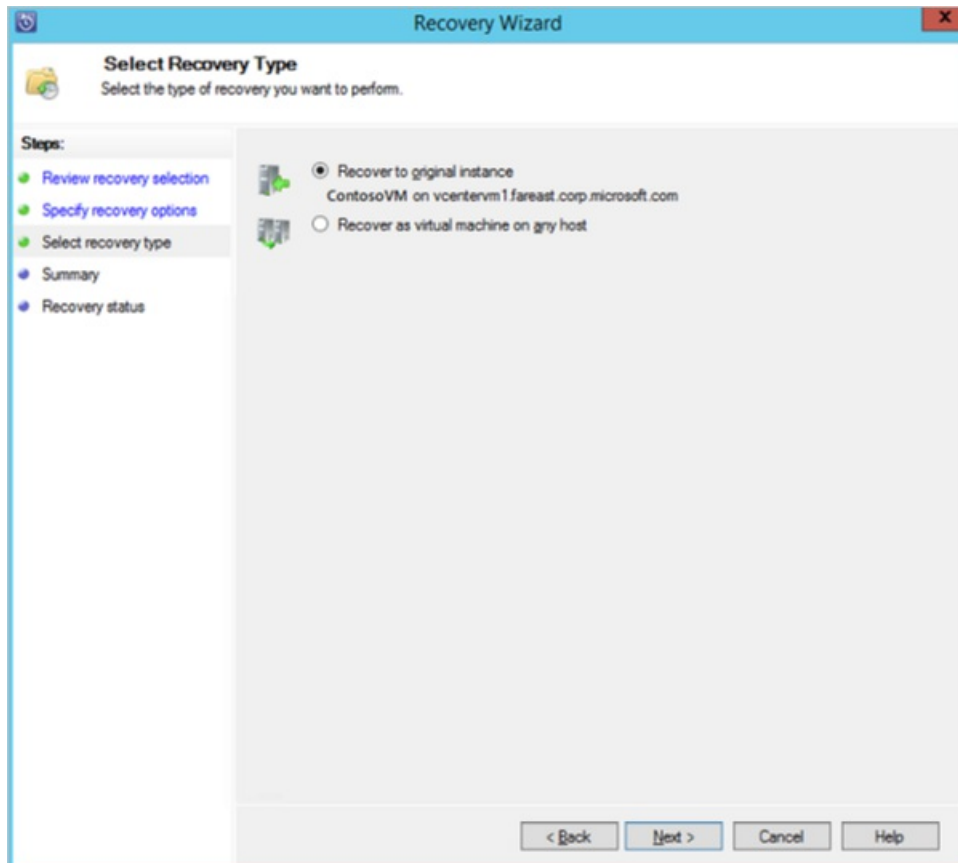
6. Select **Next** to go to the **Specify Recovery Options** screen. Select **Next** again to go to the **Select**

Recovery Type screen.

NOTE

VMware vSphere workloads don't support enabling network bandwidth throttling.

7. On the **Select Recovery Type** page, either recover to the original instance or a new location.
 - If you choose **Recover to original instance**, you don't need to make any more choices in the wizard. The data for the original instance is used.
 - If you choose **Recover as virtual machine on any host**, then on the **Specify Destination** screen, provide the information for **ESXi Host, Resource Pool, Folder,** and **Path**.



8. On the **Summary** page, review your settings and select **Recover** to start the recovery process.

The **Recovery status** screen shows the progression of the recovery operation.

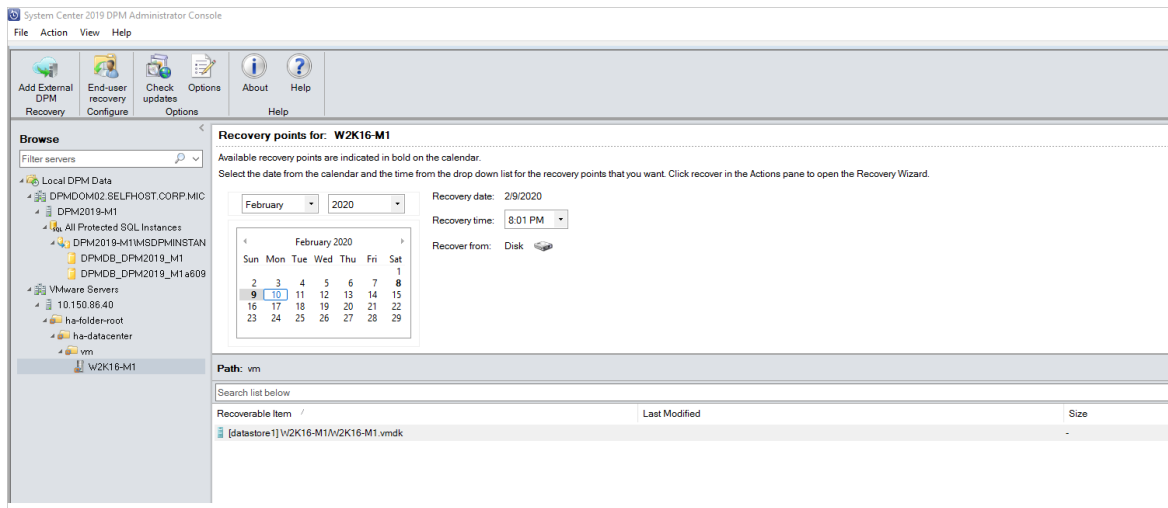
Restore an individual file from a VM

You can restore individual files from a protected VM recovery point. This feature is only available for Windows Server VMs. Restoring individual files is similar to restoring the entire VM, except you browse into the VMDK and find the files you want before you start the recovery process.

NOTE

Restoring an individual file from a VM is available only for Windows VM and disk recovery points.

1. In the Azure Backup Server Administrator Console, select the **Recovery** view.
2. Using the **Browse** pane, browse or filter to find the VM you want to recover. After you select a VM or folder, the ****Recovery points** for pane display the available recovery points.



- In the **Recovery points** for pane, use the calendar to select the wanted recovery points' date. Depending on how the backup policy was configured, dates can have more than one recovery point.
- After you select the day when the recovery point was taken, make sure you choose the correct **Recovery time**.

NOTE


If the selected date has multiple recovery points, choose your recovery point by selecting it in the **Recovery time** drop-down menu.

After you choose the recovery point, the list of recoverable items appears in the **Path** pane.

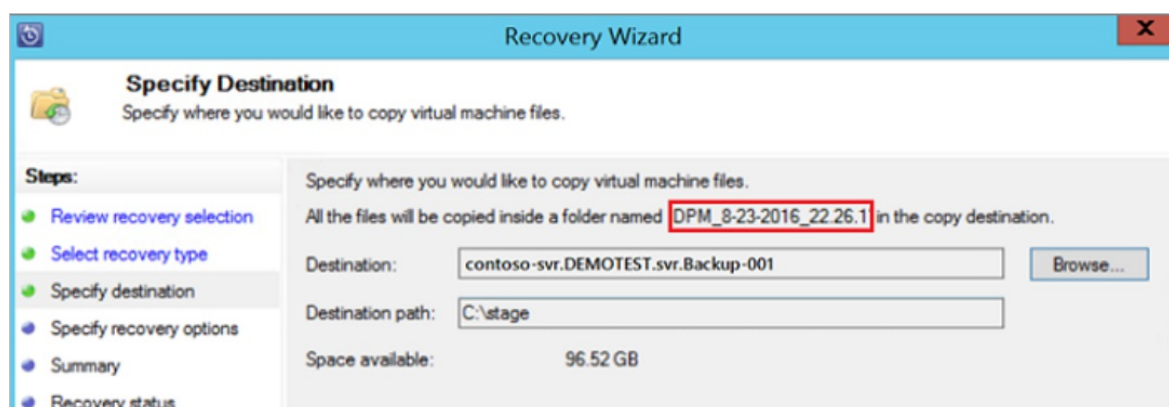
- To find the files you want to recover, in the **Path** pane, double-click the item in the **Recoverable Item** column to open it. Then select the file or folders you want to recover. To select multiple items, select the **Ctrl** key while you select each item. Use the **Path** pane to search the list of files or folders that appear in the **Recoverable Item** column.

NOTE

Search list below doesn't search into subfolders. To search through subfolders, double-click the folder. Use the **Up** button to move from a child folder into the parent folder. You can select multiple items (files and folders), but they must be in the same parent folder. You can't recover items from multiple folders in the same recovery job.



6. When you've selected the items for recovery, in the Administrator Console tool ribbon, select **Recover** to open the **Recovery Wizard**. In the **Recovery Wizard**, the **Review Recovery Selection** screen shows the selected items to be recovered.
7. On the **Specify Recovery Options** screen, do one of the following steps:
 - Select **Modify** to enable network bandwidth throttling. In the **Throttle** dialog box, select **Enable network bandwidth usage throttling** to turn it on. Once enabled, configure the **Settings** and **Work Schedule**.
 - Select **Next** to leave network throttling disabled.
8. On the **Select Recovery Type** screen, select **Next**. You can only recover your files or folders to a network folder.
9. On the **Specify Destination** screen, select **Browse** to find a network location for your files or folders. Azure Backup Server creates a folder where all recovered items are copied. The folder name has the prefix MABS_day-month-year. When you select a location for the recovered files or folder, the details for that location are provided.



10. On the **Specify Recovery Options** screen, choose which security setting to apply. You can opt to modify the network bandwidth usage throttling, but throttling is disabled by default. Also, **SAN Recovery** and **Notification** aren't enabled.
11. On the **Summary** screen, review your settings and select **Recover** to start the recovery process. The **Recovery status** screen shows the progression of the recovery operation.

Next steps

Now that you've covered backing up your Azure VMware Solution VMs with Azure Backup Server, you may want to learn about:

- [Troubleshooting when setting up backups in Azure Backup Server.](#)
- [Lifecycle management of Azure VMware Solution VMs.](#)

Prepare Azure Site Recovery resources for disaster recovery of Azure VMware Solution VMs

12/16/2022 • 3 minutes to read • [Edit Online](#)

This article describes how to prepare Azure resources and components so that you can set up disaster recovery of Azure VMware Solution VMs using [Azure Site Recovery](#) service. [Azure VMware Solution](#) provides private clouds in Azure. These private clouds contain vSphere clusters, built from dedicated bare-metal Azure infrastructure.

This article is the first tutorial in a series that shows you how to set up disaster recovery for Azure VMware Solution VMs.

In this tutorial, you learn how to:

- Verify that the Azure account has replication permissions.
- Create a Recovery Services vault. A vault holds metadata and configuration information for VMs, and other replication components.
- Set up an Azure virtual network (VNet). When Azure VMs are created after failover, they're joined to this network.

NOTE

Tutorials show you the simplest deployment path for a scenario. They use default options where possible, and don't show all possible settings and paths. For detailed instructions, review the article in the How To section of the Site Recovery Table of Contents.

NOTE

Some of the concepts of using Azure Site Recovery for Azure VMware Solution overlap with disaster recovery of on-prem VMware VMs and hence documentation will be cross-referenced accordingly.

Before you start

- [Deploy](#) an Azure VMware Solution private cloud in Azure
- Review the architecture for [VMware](#) disaster recovery
- Read common questions for [VMware](#)

If you don't have an Azure subscription, create a [free account](#) before you begin. Then sign in to the [Azure portal](#).

Verify account permissions

If you just created your free Azure account, you're the administrator of your subscription and you have the permissions you need. If you're not the subscription administrator, work with the administrator to assign the permissions you need. To enable replication for a new virtual machine, you must have permission to:

- Create a VM in the selected resource group.
- Create a VM in the selected virtual network.
- Write to an Azure storage account.

- Write to an Azure managed disk.

To complete these tasks your account should be assigned the Virtual Machine Contributor built-in role. In addition, to manage Site Recovery operations in a vault, your account should be assigned the Site Recovery Contributor built-in role.

Create a Recovery Services vault

1. From the Azure portal menu, select **Create a resource**, and search the Marketplace for **Recovery**.
2. Select **Backup and Site Recovery** from the search results, and in the Backup and Site Recovery page, click **Create**.
3. In the **Create Recovery Services vault** page, select the **Subscription**. We're using **Contoso Subscription**.
4. In **Resource group**, select an existing resource group or create a new one. For this tutorial we're using **contosoRG**.
5. In **Vault name**, enter a friendly name to identify the vault. For this set of tutorials we're using **ContosoVMVault**.
6. In **Region**, select the region in which the vault should be located. We're using **West Europe**.
7. Select **Review + create**.

The screenshot shows the 'Create Recovery Services vault' page in the Microsoft Azure portal. The breadcrumb navigation is 'Home > New > Backup and Site Recovery > Create Recovery Services vault'. The page title is 'Create Recovery Services vault'. There are three tabs: 'Basics *', 'Tags', and 'Review + create'. The 'Basics' tab is active. Under 'Project Details', there is a description: 'Select the subscription and the resource group in which you want to create the vault.' There are two dropdown menus: 'Subscription *' with 'Contoso Subscription' selected, and 'Resource group *' with 'contosoRG' selected. A 'Create new' link is visible next to the Resource group dropdown. Under 'Instance Details', there are two dropdown menus: 'Vault name *' with 'ConstosoVMVault' selected and a green checkmark, and 'Region *' with 'West Europe' selected. At the bottom, there are two buttons: 'Review + create' (highlighted in blue) and 'Next: Tags'.

The new vault will now be listed in **Dashboard > All resources**, and on the main **Recovery Services vaults** page.

Set up an Azure network

Azure VMware Solution VMs are replicated to Azure managed disks. When failover occurs, Azure VMs are

created from these managed disks, and joined to the Azure network you specify in this procedure.

1. In the [Azure portal](#), select **Create a resource** > **Networking** > **Virtual network**.
2. Keep **Resource Manager** selected as the deployment model.
3. In **Name**, enter a network name. The name must be unique within the Azure resource group. We're using **ContosoASRnet** in this tutorial.
4. In **Address space**, enter the virtual network's address range in CDR notation. We're using **10.1.0.0/24**.
5. In **Subscription**, select the subscription in which to create the network.
6. Specify the **Resource group** in which the network will be created. We're using the existing resource group **contosoRG**.
7. In **Location**, select the same region as that in which the Recovery Services vault was created. In our tutorial it's **West Europe**. The network must be in the same region as the vault.
8. In **Address range**, enter the range for the network. We're using **10.1.0.0/24**, and not using a subnet.
9. We're leaving the default options of basic DDoS protection, with no service endpoint, or firewall on the network.
10. Select **Create**.

Create virtual network □ ×

Name *
 ✓

Address space * ⓘ
 ✓
10.1.0.0 - 10.1.0.255 (256 addresses)

Add an IPv6 address space ⓘ

Subscription *
 ▼

Resource group *
 ▼
[Create new](#)

Location *
 ▼

Subnet

Name *

Address range * ⓘ
 ✓
10.1.0.0 - 10.1.0.255 (256 addresses)

DDoS protection ⓘ
 Basic Standard

Service endpoints ⓘ
 Disabled Enabled

Firewall ⓘ
 Disabled Enabled

[Automation options](#)

The virtual network takes a few seconds to create. After it's created, you'll see it in the Azure portal dashboard.

Next steps

[Prepare infrastructure](#)

- [Learn about](#) Azure networks.
- [Learn about](#) managed disks.

Prepare Azure VMware Solution for disaster recovery to Azure Site Recovery

12/16/2022 • 5 minutes to read • [Edit Online](#)

This article describes how to prepare Azure VMware Solution servers for disaster recovery to Azure using the [Azure Site Recovery](#) services.

This is the second tutorial in a series that shows you how to set up disaster recovery to Azure for Azure VMware Solution VMs. In the first tutorial, we [set up the Azure components](#) needed for Azure VMware Solution disaster recovery.

In this article, you learn how to:

- Prepare an account on the vCenter Server to automate VM discovery.
- Prepare an account for automatic installation of the Mobility service on VMware vSphere VMs.
- Review VMware vCenter Server and VM requirements and support.
- Prepare to connect to Azure VMs after failover.

NOTE

Tutorials show you the simplest deployment path for a scenario. They use default options where possible, and don't show all possible settings and paths. For detailed instructions, review the article in the How To section of the Site Recovery Table of Contents.

Before you start

Make sure you've prepared Azure as described in the [first tutorial in this series](#).

Prepare an account for automatic discovery

Site Recovery needs access to Azure VMware Solution servers to:

- Automatically discover VMs. At least a read-only account is required.
- Orchestrate replication, failover, and failback. You need an account that can run operations such as creating and removing disks, and powering on VMs.

Create the account as follows:

1. To use a dedicated account, create a role at the vCenter Server level. Give the role a name such as **Azure_Site_Recovery**.
2. Assign the role the permissions summarized in the table below.
3. Create a user on the vCenter Server. Assign the role to the user.

VMware account permissions

TASK	ROLE/PERMISSIONS	DETAILS
------	------------------	---------

TASK	ROLE/PERMISSIONS	DETAILS
VM discovery	<p>At least a read-only user</p> <p>Data Center object -> Propagate to Child Object, role=Read-only</p>	<p>User assigned at datacenter level, and has access to all the objects in the datacenter.</p> <p>To restrict access, assign the No access role with the Propagate to child object, to the child objects (vSphere hosts, datastores, VMs and networks).</p>
Full replication, failover, failback	<p>Create a role (Azure_Site_Recovery) with the required permissions, and then assign the role to a VMware user or group</p> <p>Data Center object -> Propagate to Child Object, role=Azure_Site_Recovery</p> <p>Datastore -> Allocate space, browse datastore, low-level file operations, remove file, update virtual machine files</p> <p>Network -> Network assign</p> <p>Resource -> Assign VM to resource pool, migrate powered off VM, migrate powered on VM</p> <p>Tasks -> Create task, update task</p> <p>Virtual machine -> Configuration</p> <p>Virtual machine -> Interact -> answer question, device connection, configure CD media, configure floppy media, power off, power on, VMware tools install</p> <p>Virtual machine -> Inventory -> Create, register, unregister</p> <p>Virtual machine -> Provisioning -> Allow virtual machine download, allow virtual machine files upload</p> <p>Virtual machine -> Snapshots -> Remove snapshots</p>	<p>User assigned at datacenter level, and has access to all the objects in the datacenter.</p> <p>To restrict access, assign the No access role with the Propagate to child object, to the child objects (vSphere hosts, datastores, VMs and networks).</p>

Prepare an account for Mobility service installation

The Mobility service must be installed on machines you want to replicate. Azure Site Recovery can do a push installation of this service when you enable replication for a machine, or you can install it manually, or using installation tools.

- In this tutorial, we're going to install the Mobility service with the push installation.
- For this push installation, you need to prepare an account that Azure Site Recovery can use to access the VM. You specify this account when you set up disaster recovery in the Azure console.

Prepare the account as follows:

- Prepare a domain or local account with permissions to install on the VM.
- **Windows VMs:** To install on Windows VMs if you're not using a domain account, disable Remote User Access control on the local machine. To do this, in the registry > `HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System`, add the DWORD entry `LocalAccountTokenFilterPolicy`, with a value of 1.
- **Linux VMs:** To install on Linux VMs, prepare a root account on the source Linux server.

Check Azure VMware Solution requirements

Make sure VMware vCenter Server and VMs comply with requirements.

1. Verify Azure VMware Solution [software versions](#).
2. Verify [VMware vCenter Server requirements](#).
3. For Linux VMs, [check](#) file system and storage requirements.
4. Check [network](#) and [storage](#) support.
5. Check what's supported for [Azure networking](#), [storage](#), and [compute](#), after failover.
6. Your Azure VMware Solution VMs you replicate to Azure must comply with [Azure VM requirements](#).
7. In Linux virtual machines, device name or mount point name should be unique. Ensure that no two devices/mount points have the same names. Note that name aren't case-sensitive. For example, naming two devices for the same VM as `device1` and `Device1` isn't allowed.

Prepare to connect to Azure VMs after failover

After failover, you might want to connect to the Azure VMs from your Azure VMware Solution network.

To connect to Windows VMs using RDP after failover, do the following:

- **Internet access.** Before failover, enable RDP on the Azure VMware Solution VM before failover. Make sure that TCP, and UDP rules are added for the **Public** profile, and that RDP is allowed in **Windows Firewall > Allowed Apps**, for all profiles.
- **Site-to-site VPN access:**
 - Before failover, enable RDP on the Azure VMware Solution VM.
 - RDP should be allowed in the **Windows Firewall -> Allowed apps and features for Domain and Private** networks.
 - Check that the operating system's SAN policy is set to **OnlineAll**. [Learn more](#).
- There should be no Windows updates pending on the VM when you trigger a failover. If there are, you won't be able to sign in to the virtual machine until the update completes.
- On the Windows Azure VM after failover, check **Boot diagnostics** to view a screenshot of the VM. If you can't connect, check that the VM is running and review these [troubleshooting tips](#).

To connect to Linux VMs using SSH after failover, do the following:

- On the Azure VMware Solution VM before failover, check that the Secure Shell service is set to start automatically on system boot.
- Check that firewall rules allow an SSH connection.
- On the Azure VM after failover, allow incoming connections to the SSH port for the network security group rules on the failed over VM, and for the Azure subnet to which it's connected.
- [Add a public IP address](#) for the VM.
- You can check **Boot diagnostics** to view a screenshot of the VM.

Failback requirements

If you plan to fail back to your Azure VMware Solution cloud, there are a number of [prerequisites for failback](#). You can prepare these now, but you don't need to. You can prepare after you fail over to Azure.

Next steps

[Setup disaster recovery](#)

- If you're replicating multiple VMs, [perform capacity planning](#).

Setup Azure Site Recovery for Azure VMware Solution VMs

12/16/2022 • 10 minutes to read • [Edit Online](#)

This article describes how to enable replication for Azure VMware Solution VMs, for disaster recovery to Azure using the [Azure Site Recovery](#) service.

This is the third tutorial in a series that shows you how to set up disaster recovery to Azure for Azure VMware Solution VMs. In the previous tutorial, we [prepared Azure VMware Solution environment](#) for disaster recovery to Azure.

In this tutorial, you learn how to:

- Set up the source replication settings, and an Azure Site Recovery configuration server in Azure VMware Solution private cloud
- Set up the replication target settings.
- Create a replication policy.
- Enable replication for a VMware vSphere VM.

NOTE

Tutorials show you the simplest deployment path for a scenario. They use default options where possible, and don't show all possible settings and paths. For detailed instructions, review the article in the How To section of the Site Recovery Table of Contents.

Before you start

Complete the previous tutorials:

1. Make sure you've [set up Azure](#) for disaster recovery to Azure.
2. Follow [these steps](#) to prepare your Azure VMware Solution deployment for disaster recovery to Azure.
3. In this tutorial we show you how to replicate a single VM. If you're deploying multiple VMs you should use the [Deployment Planner Tool](#). [Learn more](#) about this tool.
4. This tutorial uses a number of options you might want to do differently:
 - The tutorial uses an OVA template to create the configuration server VMware vSphere VM. If you can't do this for some reason, follow [these instructions](#) to set up the configuration server manually.
 - In this tutorial, Site Recovery automatically downloads and installs MySQL to the configuration server. If you prefer, you can set it up manually instead. [Learn more](#).

Select a protection goal

1. In **Recovery Services vaults**, select the vault name. We're using **ContosoVMVault** for this scenario.
2. In **Getting Started**, select Site Recovery. Then select **Prepare Infrastructure**.
3. In **Protection goal > Where are your machines located**, select **On-premises**.
4. In **Where do you want to replicate your machines**, select **To Azure**.
5. In **Are your machines virtualized**, select **Yes, with VMware vSphere Hypervisor**. Then select **OK**.

Set up the source environment

In your source environment, you need a single, highly available, on-premises machine to host these on-premises Site Recovery components:

- **Configuration server:** The configuration server coordinates communications between Azure VMware Solution private cloud and Azure, and manages data replication.
- **Process server:** The process server acts as a replication gateway. It receives replication data; optimizes it with caching, compression, and encryption, and sends it to a cache storage account in Azure. The process server also installs the Mobility Service agent on VMs you want to replicate, and performs automatic discovery of Azure VMware Solution VMs.
- **Master target server:** The master target server handles replication data during failback from Azure.

All of these components are installed together on a single Azure VMware Solution machine that's known as the *configuration server*. By default, for Azure VMware Solution disaster recovery, we set up the configuration server as a highly available VMware vSphere VM. To do this, you download a prepared Open Virtualization Application (OVA) template, and import the template into VMware vCenter Server to create the VM.

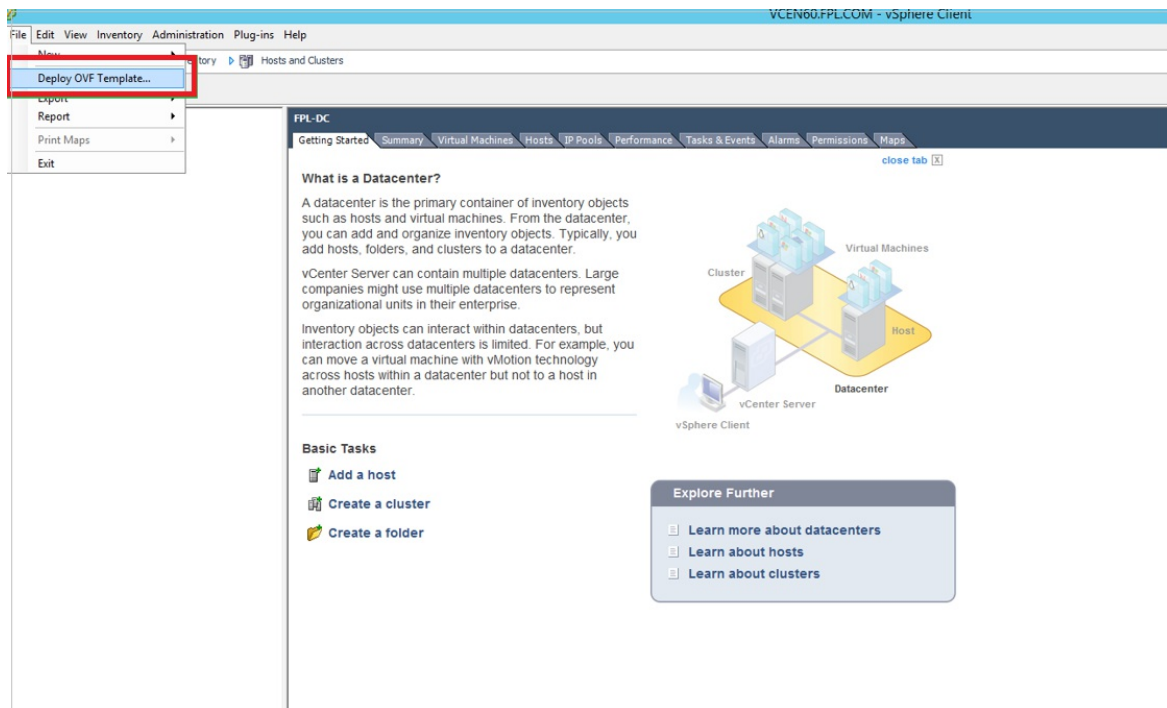
- The latest version of the configuration server is available in the portal. You can also download it directly from the [Microsoft Download Center](#).
- If for some reason you can't use an OVA template to set up a VM, follow [these instructions](#) to set up the configuration server manually.
- The license provided with OVF template is an evaluation license valid for 180 days. Windows running on the VM must be activated with the required license.

Download the VM template

1. In the vault, go to **Prepare Infrastructure > Source**.
2. In **Prepare source**, select **+Configuration server**.
3. In **Add Server**, check that **Configuration server for VMware** appears in **Server type**.
4. Download the OVA template for the configuration server.

Import the template in VMware

1. Sign in to the VMware vCenter Server with the VMware vSphere Client.
2. On the **File** menu, select **Deploy OVF Template** to start the **Deploy OVF Template Wizard**.



3. On **Select source**, enter the location of the downloaded OVF.
4. On **Review details**, select **Next**.
5. On **Select name and folder** and **Select configuration**, accept the default settings.
6. On **Select storage**, for best performance select **Thick Provision Eager Zeroed** in **Select virtual disk format**.
7. On the rest of the wizard pages, accept the default settings.
8. On **Ready to complete**, to set up the VM with the default settings, select **Power on after deployment** > **Finish**.

TIP

If you want to add an additional NIC, clear **Power on after deployment** > **Finish**. By default, the template contains a single NIC. You can add additional NICs after deployment.

Add an additional adapter

If you want to add an additional NIC to the configuration server, add it before you register the server in the vault. Adding additional adapters isn't supported after registration.

1. In the vSphere Client inventory, right-click the VM and select **Edit Settings**.
2. In **Hardware**, select **Add** > **Ethernet Adapter**. Then select **Next**.
3. Select an adapter type and a network.
4. To connect the virtual NIC when the VM is turned on, select **Connect at power on**. Select **Next** > **Finish**. Then select **OK**.

Register the configuration server

After the configuration server is setup, you register it in the vault.

1. From the VMware vSphere Client console, turn on the VM.
2. The VM boots up into a Windows Server 2016 installation experience. Accept the license agreement, and

enter an administrator password.

3. After the installation finishes, sign in to the VM as the administrator.
4. The first time you sign in, the Azure Site Recovery Configuration Tool starts within a few seconds.
5. Enter a name that's used to register the configuration server with Azure Site Recovery. Then select **Next**.
6. The tool checks that the VM can connect to Azure. After the connection is established, select **Sign in** to sign in to your Azure subscription. The credentials must have access to the vault in which you want to register the configuration server. Ensure that necessary [roles](#) are assigned to this user.
7. The tool performs some configuration tasks and then reboots.
8. Sign in to the machine again. In a few seconds, the Configuration Server Management Wizard starts automatically.

Configure settings and add the VMware vCenter Server

Finish setting up and registering the configuration server. Before proceeding, ensure all [pre-requisites](#) are met for successful set up of configuration server.

1. In the configuration server management wizard, select **Setup connectivity**. From the dropdowns, first select the NIC that the in-built process server uses for discovery and push installation of mobility service on source machines, and then select the NIC that Configuration Server uses for connectivity with Azure. Then select **Save**. You cannot change this setting after it's configured.
2. In **Select Recovery Services vault**, select your Azure subscription and the relevant resource group and vault.
3. In **Install third-party software**, accept the license agreement. Select **Download and Install** to install MySQL Server. If you placed MySQL in the path, this step can be skipped. [Learn more](#)
4. In **Validate appliance configuration**, prerequisites are verified before you continue.
5. In **Configure vCenter Server/vSphere ESXi server**, enter the FQDN or IP address of the vCenter Server, where the VMs you want to replicate are located. Enter the port on which the server is listening. Enter a friendly name to be used for the VMware vCenter Server in the vault.
6. Enter user credentials to be used by the configuration server to connect to the VMware vCenter Server. Ensure that the user name and password are correct and is a part of the Administrators group of the virtual machine to be protected. Azure Site Recovery uses these credentials to automatically discover VMware vSphere VMs that are available for replication. Select **Add**, and then select **Continue**.
7. In **Configure virtual machine credentials**, enter the user name and password that will be used to automatically install Mobility Service on VMs when replication is enabled.
 - For Windows machines, the account needs local administrator privileges on the machines you want to replicate.
 - For Linux, provide details for the root account.
8. Select **Finalize configuration** to complete registration.
9. After registration finishes, open the Azure portal and verify that the configuration server and VMware server are listed on **Recovery Services Vault > Manage > Site Recovery Infrastructure > Configuration Servers**.

After the configuration server is registered, Site Recovery connects to VMware vCenter Server by using the specified settings, and discovers VMs.

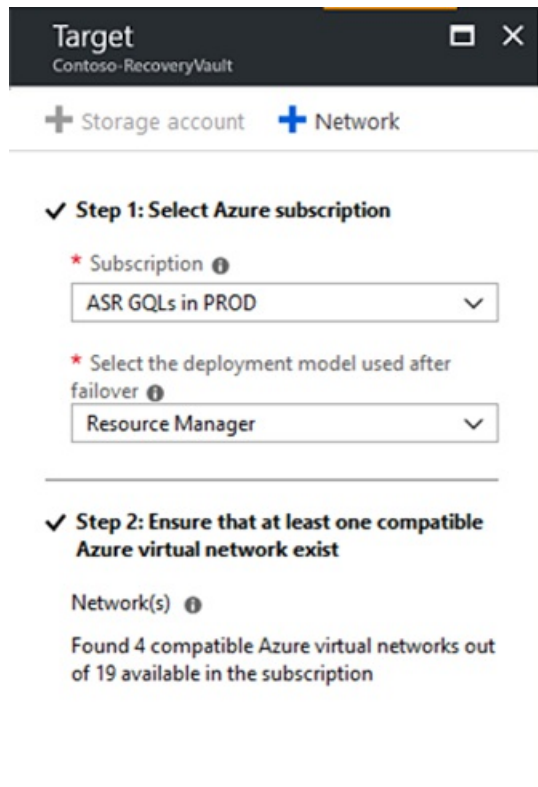
NOTE

It can take 15 minutes or more for the account name to appear in the portal. To update immediately, select **Configuration Servers > *server name* > Refresh Server**.

Set up the target environment

Select and verify target resources.

1. Select **Prepare infrastructure** > **Target**. Select the Azure subscription you want to use. We're using a Resource Manager model.
2. Azure Site Recovery checks that you have one or more virtual networks. You should have these when you set up the Azure components in the [first tutorial](#) in this tutorial series.



Create a replication policy

1. Open the [Azure portal](#). Search for and select **Recovery Services vaults**.
2. Select the Recovery Services vault (**ContosoVMVault** in this tutorial).
3. To create a replication policy, select **Site Recovery infrastructure** > **Replication Policies** > **+Replication Policy**.
4. In **Create replication policy**, enter the policy name. We're using **VMwareRepPolicy**.
5. In **RPO threshold**, use the default of 60 minutes. This value defines how often recovery points are created. An alert is generated if continuous replication exceeds this limit.
6. In **Recovery point retention**, specify how longer each recovery point is retained. For this tutorial we're using 72 hours. Replicated VMs can be recovered to any point in a retention window.
7. In **App-consistent snapshot frequency**, specify how often app-consistent snapshots are created. We're using the default of 60 minutes. Select **OK** to create the policy.

Create replication policy
ContosoVMVault

* Name ⓘ
VMwareRepPolicy ✓

Source type ⓘ
VMware / Physical machines ▼


Target type ⓘ
Azure ▼

* RPO threshold in mins ⓘ
60

* Recovery point retention in hours ⓘ
24

* App-consistent snapshot frequency in mins ⓘ
60

Failback replication policy name ⓘ
VMwareRepPolicy-failback

 A replication policy for failback from Azure to on-premises will be automatically created with the same settings.

OK

- The policy is automatically associated with the configuration server.
- A matching policy is automatically created for failback by default. For example, if the replication policy is **rep-policy**, then the failback policy is **rep-policy-failback**. This policy isn't used until you initiate a failback from Azure.

Note: In VMware vSphere-to-Azure scenario the crash-consistent snapshot is taken at 5 min interval.

Enable replication

Enable replication for VMs as follows:

1. Select **Replicate application > Source**.
2. In **Source**, select **On-premises**, and select the configuration server in **Source location**.
3. In **Machine type**, select **Virtual Machines**.
4. In **vCenter/vSphere Hypervisor**, select the vCenter Server that manages the host.
5. Select the process server (installed by default on the configuration server VM). Then select **OK**. Health status of each process server is indicated as per recommended limits and other parameters. Choose a healthy process server. A **critical** process server cannot be chosen. You can either **troubleshoot and resolve** the errors or set up a **scale-out process server**.
6. In **Target**, select the subscription and the resource group in which you want to create the failed-over VMs. We're using the Resource Manager deployment model.
7. Select the Azure network and subnet to which Azure VMs connect when they're created after failover.

8. Select **Configure now for selected machines** to apply the network setting to all VMs on which you enable replication. Select **Configure later** to select the Azure network per machine.
9. In **Virtual Machines > Select virtual machines**, select each machine you want to replicate. You can only select machines for which replication can be enabled. Then select **OK**. If you are not able to view/select any particular virtual machine, [learn more](#) about resolving the issue.
10. In **Properties > Configure properties**, select the account to be used by the process server to automatically install Mobility Service on the machine.
11. In **Replication settings > Configure replication settings**, verify that the correct replication policy is selected.
12. Select **Enable Replication**. Site Recovery installs the Mobility Service when replication is enabled for a VM.
13. You can track progress of the **Enable Protection** job in **Settings > Jobs > Site Recovery Jobs**. After the **Finalize Protection** job runs and a recovery point generation is complete, the machine is ready for failover.
14. It can take 15 minutes or longer for changes to take effect and appear in the portal.
15. To monitor VMs you add, check the last discovered time for VMs in **Configuration Servers > Last Contact At**. To add VMs without waiting for the scheduled discovery, highlight the configuration server (don't select it) and select **Refresh**.

Next steps

After enabling replication, run a drill to make sure everything's working as expected.

[Run a disaster recovery drill](#)

Run a disaster recovery drill from Azure VMware Solution to Azure

12/16/2022 • 3 minutes to read • [Edit Online](#)

This article describes how to run a disaster recovery drill for an Azure VMware Solution VM to Azure using the [Azure Site Recovery](#) service. A drill validates your replication strategy without data loss.

This is the fourth tutorial in a series that shows you how to set up disaster recovery to Azure for Azure VMware Solution machines.

In this tutorial, learn how to:

- Set up an isolated network for the test failover
- Prepare to connect to the Azure VM after failover
- Run a test failover for a single machine.

NOTE

Tutorials show you the simplest deployment path for a scenario. They use default options where possible, and don't show all possible settings and paths. If you want to learn about the disaster recovery drill steps in more detail, [review this article](#).

Before you start

Complete the previous tutorials:

1. Make sure you've [set up Azure](#) for disaster recovery to Azure.
2. Follow [these steps](#) to prepare your Azure VMware Solution deployment for disaster recovery to Azure.
3. [Set up](#) disaster recovery for Azure VMware Solution VMs.

Verify VM properties

Before you run a test failover, verify the VM properties, and make sure that the [VMware vSphere VM](#) complies with Azure requirements.

1. In **Protected Items**, click **Replicated Items** > and the VM.
2. In the **Replicated item** pane, there's a summary of VM information, health status, and the latest available recovery points. Click **Properties** to view more details.
3. In **Compute and Network**, you can modify the Azure name, resource group, target size, availability set, and managed disk settings.
4. You can view and modify network settings, including the network/subnet in which the Azure VM will be located after failover, and the IP address that will be assigned to it.
5. In **Disks**, you can see information about the operating system and data disks on the VM.

Create a network for test failover

We recommended that for test failover, you choose a network that's isolated from the production recovery site network specific in the **Compute and Network** settings for each VM. By default, when you create an Azure

virtual network, it is isolated from other networks. The test network should mimic your production network:

- The test network should have same number of subnets as your production network. Subnets should have the same names.
- The test network should use same IP address class and subnet range.
- Update the DNS of the test network with the IP address specified for the DNS VM in **Compute and Network** settings. Read [test failover considerations for Active Directory](#) for more details.

Run a test failover for a single VM

When you run a test failover, the following happens:

1. A prerequisites check runs to make sure all of the conditions required for failover are in place.
2. Failover processes the data, so that an Azure VM can be created. If you select the latest recovery point, a recovery point is created from the data.
3. An Azure VM is created using the data processed in the previous step.

Run the test failover as follows:

1. In **Settings > Replicated Items**, click the VM > **+Test Failover**.
2. Select the **Latest processed** recovery point for this tutorial. This fails over the VM to the latest available point in time. The time stamp is shown. With this option, no time is spent processing data, so it provides a low RTO (recovery time objective).
3. In **Test Failover**, select the target Azure network to which Azure VMs will be connected after failover occurs.
4. Click **OK** to begin the failover. You can track progress by clicking on the VM to open its properties. Or you can click the **Test Failover** job in vault name > **Settings > Jobs > Site Recovery jobs**.
5. After the failover finishes, the replica Azure VM appears in the Azure portal > **Virtual Machines**. Check that the VM is the appropriate size, that it's connected to the right network, and that it's running.
6. You should now be able to connect to the replicated VM in Azure.
7. To delete Azure VMs created during the test failover, click **Cleanup test failover** on the VM. In **Notes**, record and save any observations associated with the test failover.

In some scenarios, failover requires additional processing that takes around eight to ten minutes to complete. You might notice longer test failover times for VMware Linux machines, VMware VMs that don't have the DHCP service enables, and VMware VMs that don't have the following boot drivers: storvsc, vmbus, storflt, intelide, atapi.

Connect after failover

If you want to connect to Azure VMs using RDP/SSH after failover, [prepare to connect](#). If you encounter any connectivity issues after failover, follow the [troubleshooting](#) guide.

Next steps

[Run a failover](#)

Fail over Azure VMware Solution VMs

12/16/2022 • 3 minutes to read • [Edit Online](#)

This article describes how to fail over an Azure VMware Solution VM to Azure with [Azure Site Recovery](#).

This is the fifth tutorial in a series that shows you how to set up disaster recovery to Azure for Azure VMware Solution VMs.

In this tutorial, you learn how to:

- Verify that the Azure VMware Solution VM properties conform with Azure requirements.
- Fail over specific VMs to Azure.

NOTE

Tutorials show you the simplest deployment path for a scenario. They use default options where possible and don't show all possible settings and paths. If you want to learn about failover in detail, see [Fail over VMs](#).

[Learn about](#) different types of failover. If you want to fail over multiple VMs in a recovery plan, review [this article](#).

Before you start

Complete the previous tutorials:

1. Make sure you've [set up Azure](#) for disaster recovery to Azure.
2. Follow [these steps](#) to prepare your Azure VMware Solution deployment for disaster recovery to Azure.
3. [Set up](#) disaster recovery for Azure VMware Solution VMs.
4. Run a [disaster recovery drill](#) to make sure that everything's working as expected.

Verify VM properties

Before you run a failover, check the VM properties to make sure that the VMs meet [Azure requirements](#).

Verify properties as follows:

1. In **Protected Items**, select **Replicated Items**, and then select the VM you want to verify.
2. In the **Replicated item** pane, there's a summary of VM information, health status, and the latest available recovery points. Select **Properties** to view more details.
3. In **Compute and Network**, you can modify these properties as needed:
 - Azure name
 - Resource group
 - Target size
 - [Availability set](#)
 - Managed disk settings
4. You can view and modify network settings, including:
 - The network and subnet in which the Azure VM will be located after failover.

- The IP address that will be assigned to it.
5. In **Disks**, you can see information about the operating system and data disks on the VM.

Run a failover to Azure

1. In **Settings** > **Replicated items**, select the VM you want to fail over, and then select **Failover**.
2. In **Failover**, select a **Recovery Point** to fail over to. You can use one of the following options:
 - **Latest**: This option first processes all the data sent to Site Recovery. It provides the lowest Recovery Point Objective (RPO) because the Azure VM that's created after failover has all the data that was replicated to Site Recovery when the failover was triggered.
 - **Latest processed**: This option fails the VM over to the latest recovery point processed by Site Recovery. This option provides a low RTO (Recovery Time Objective) because no time is spent processing unprocessed data.
 - **Latest app-consistent**: This option fails the VM over to the latest app-consistent recovery point processed by Site Recovery.
 - **Custom**: This option lets you specify a recovery point.
3. Select **Shut down machine before beginning failover** to attempt to shut down source VMs before triggering the failover. Failover continues even if the shutdown fails. You can follow the failover progress on the **Jobs** page.

In some scenarios, failover requires additional processing that takes around 8 to 10 minutes to complete. You might notice longer test failover times for:

- VMware vSphere VMs running a Mobility service version older than 9.8.
- VMware vSphere Linux VMs.
- VMware vSphere VMs that don't have the DHCP service enabled.
- VMware vSphere VMs that don't have the following boot drivers: storvsc, vmbus, storflt, intelide, atapi.

WARNING

Don't cancel a failover in progress. Before failover is started, VM replication is stopped. If you cancel a failover in progress, failover stops, but the VM won't replicate again.

Connect to failed-over VM

1. If you want to connect to Azure VMs after failover by using Remote Desktop Protocol (RDP) and Secure Shell (SSH), [verify that the requirements have been met](#).
2. After failover, go to the VM and validate by [connecting](#) to it.
3. Use **Change recovery point** if you want to use a different recovery point after failover. After you commit the failover in the next step, this option will no longer be available.
4. After validation, select **Commit** to finalize the recovery point of the VM after failover.
5. After you commit, all the other available recovery points are deleted. This step completes the failover.

TIP

If you encounter any connectivity issues after failover, follow the [troubleshooting guide](#).

Next steps

After failover, reprotect the Azure VMs to Azure VMware Solution private cloud. Then, after the VMs are

reprotected and replicating to the Azure VMware Solution private cloud, fail back from Azure when you're ready.

[Reprotect Azure VMs Failback from Azure](#)

Reprotect from Azure to Azure VMware Solution private cloud

12/16/2022 • 4 minutes to read • [Edit Online](#)

After [failover](#) of Azure VMware Solution VMs to Azure, the first step in failing back to your Azure VMware Solution private cloud is to reprotect the Azure VMs that were created during failover. This article describes how to do this.

Before you begin

1. Follow the steps in [this article](#) to prepare for reprotection and failback, including setting up a process server in Azure, and an Azure VMware Solution private cloud master target server, and configuring a site-to-site VPN, or ExpressRoute private peering, for failback.
2. Make sure that the Azure VMware Solution private cloud configuration server is running and connected to Azure. During failback, the VM must exist in the configuration server database. Otherwise, failback is unsuccessful.
3. Delete any snapshots on the Azure VMware Solution private cloud master target server. Reprotection won't work if there are snapshots. The snapshots on the VM are automatically merged during a reprotect job.
4. If you're reprotecting VMs gathered into a replication group for multi-VM consistency, make sure they all have the same operating system (Windows or Linux) and make sure that the master target server you deploy has the same type of operating system. All VMs in a replication group must use the same master target server.
5. Open [the required ports](#) for failback.
6. Ensure that the vCenter Server is connected before failback. Otherwise, disconnecting disks and attaching them back to the virtual machine fails.
7. If a vCenter Server manages the VMs to which you'll fail back, make sure that you have the required permissions. If you perform a read-only user vCenter Server discovery and protect virtual machines, protection succeeds, and failover works. However, during reprotection, failover is unsuccessful because the datastores can't be discovered, and aren't listed during reprotection. To resolve this problem, you can update the vCenter Server credentials with an [appropriate account/permissions](#), and then retry the job.
8. If you used a template to create your virtual machines, ensure that each VM has its own UUID for the disks. If the Azure VMware Solution VM UUID clashes with the UUID of the master target server because both were created from the same template, reprotection fails. Deploy from a different template.
9. If you're failing back to an alternate vCenter Server, make sure that the new vCenter Server and the master target server are discovered. Typically if they're not the datastores aren't accessible, or aren't visible in **Reprotect**.
10. Verify the following scenarios in which you can't fail back:
 - If you're using either the ESXi 5.5 free edition or the vSphere 6 Hypervisor free edition. Upgrade to a different version.
 - If you have a Windows Server 2008 R2 SP1 physical server.
 - VMware vSphere VMs can't fail back to Hyper-V.
 - VMs that have been migrated.
 - A VM that's been moved to another resource group.
 - A replica Azure VM that's been deleted.
 - A replica Azure VM that isn't protected.
11. [Review the types of failback](#) you can use - original location recovery and alternate location recovery.

Enable re protection

Enable replication. You can reprotect specific VMs, or a recovery plan:

- If you reprotect a recovery plan, you must provide the values for every protected machine.
- If VMs belong to a replication group for multi-VM consistency, they can only be reprotected using a recovery plan. VMs in a replication group must use the same master target server

NOTE

The amount of data sent from Azure to erstwhile source during reprotect, can be anything between 0 bytes and sum of disk size for all protected machines, and can't be calculated.

Before you start

- After a VM boots in Azure after failover, it takes some time for the agent to register back to the configuration server (up to 15 minutes). During this time, you won't be able to reprotect and an error message indicates that the agent isn't installed. If this happens, wait for a few minutes, and then reprotect.
- If you want to fail back the Azure VM to an existing Azure VMware Solution VM, mount the VM datastores with read/write access on the master target server's ESXi host.
- If you want to fail back to an alternate location, for example if the Azure VMware Solution VM doesn't exist, select the retention drive and datastore that are configured for the master target server. When you fail back to the Azure VMware Solution private cloud, the virtual machines in the failback protection plan use the same datastore as the master target server. A new VM is then created in vCenter.

Enable re protection as follows:

1. Select **Vault > Replicated items**. Right-click the virtual machine that failed over, and then select **Re-Protect**. Or, from the command buttons, select the machine, and then select **Re-Protect**.
2. Verify that the **Azure to On-premises** direction of protection is selected.
3. In **Master Target Server** and **Process Server**, select the on-premises master target server and the process server.
4. For **Datastore**, select the datastore to which you want to recover the disks in Azure VMware Solution. This option is used when the Azure VMware Solution VM is deleted, and you need to create new disks. This option is ignored if the disks already exist. You still need to specify a value.
5. Select the retention drive.
6. The failback policy is automatically selected.
7. Select **OK** to begin re protection.

NAME	PROCESS SERVER	MASTER TARGET	DATA STORE	RETENTION DRIVE	FAILBACK POLICY
-0427-1	v2a-demo	PMraturajd-cs...	datastore62	D	V2A-failback

8. A job begins to replicate the Azure VM to the Azure VMware Solution private cloud. You can track the progress on the **Jobs** tab.
 - When the reprotection succeeds, the VM enters a protected state.
 - The Azure VMware Solution VM is turned off during reprotection. This helps ensure data consistency during replication.
 - Don't turn on the Azure VMware Solution VM after reprotection finishes.

Next steps

- If you encounter any issues, review the [troubleshooting article](#).
- After the Azure VMs are protected, you can [run a failback](#). Failback shuts down the Azure VM and boots the Azure VMware Solution VM. Expect some downtime for the application, and choose a failback time accordingly.

Fail back VMs to Azure VMware Solution private cloud

12/16/2022 • 2 minutes to read • [Edit Online](#)

This article describes how to failback Azure VMs to an Azure VMware Solution private cloud, following [failover](#) of Azure VMware Solution VMs to Azure with [Azure Site Recovery](#). After failback, you enable replication so that the Azure VMware Solution VMs start replicating to Azure.

Before you start

1. Learn about [VMware vSphere failback](#).
2. Make sure you've reviewed and completed the steps to [prepare for failback](#), and that all the required components are deployed. Components include a process server in Azure, a master target server, and a VPN site-to-site connection (or ExpressRoute private peering) for failback.
3. Make sure you've completed the [requirements](#) for reprotection and failback, and that you've [enabled reprotection](#) of Azure VMs, so that they're replicating from Azure to the Azure VMware Solution private cloud. VMs must be in a replicated state in order to fail back.

Run a failover to fail back

1. Make sure that Azure VMs are reprotected and replicating to the Azure VMware Solution private cloud.
 - A VM needs at least one recovery point in order to fail back.
 - If you fail back a recovery plan, then all machines in the plan should have at least one recovery point.
2. In the vault > **Replicated items**, select the VM. Right-click the VM > **Unplanned Failover**.
3. In **Confirm Failover**, verify the failover direction (from Azure).
4. Select the recovery point that you want to use for the failover.
 - We recommend that you use the **Latest** recovery point. The app-consistent point is behind the latest point in time, and causes some data loss.
 - **Latest** is a crash-consistent recovery point.
 - With **Latest**, a VM fails over to its latest available point in time. If you have a replication group for multi-VM consistency within a recovery plan, each VM in the group fails over to its independent latest point in time.
 - If you use an app-consistent recovery point, each VM fails back to its latest available point. If a recovery plan has a replication group, each group recovers to its common available recovery point.
5. Failover begins. Azure Site Recovery shuts down the Azure VMs.
6. After failover completes, check everything's working as expected. Check that the Azure VMs are shut down.
7. With everything verified, right-click the VM > **Commit**, to finish the failover process. Commit removes the failed-over Azure VM.

NOTE

For Windows VMs, Azure Site Recovery disables the VMware tools during failover. During failback of the Windows VM, the VMware tools are enabled again.

Reprotect from Azure VMware Solution to Azure

After committing the failback, the Azure VMs are deleted. The VM is back in the Azure VMware Solution private cloud, but it isn't protected. To start replicating VMs to Azure again,as follows:

1. In the vault > **Replicated items**, select failed back VMs, and then select **Re-Protect**.
2. Specify the process server that's used to send data back to Azure.
3. Select **OK** to begin the reprotect job.

NOTE

After an Azure VMware Solution VM starts, it takes up to 15 minutes for the agent to register back to the configuration server. During this time, reprotect fails and returns an error message stating that the agent isn't installed. If this occurs, wait for a few minutes, and reprotect.

Next steps

After the reprotect job finishes, the Azure VMware Solution VM is replicating to Azure. As needed, you can [run another failover](#) to Azure.

Create an Azure VMware Solution assessment

12/16/2022 • 9 minutes to read • [Edit Online](#)

This article describes how to create an Azure VMware Solution assessment for on-premises VMs in a VMware vSphere environment with Azure Migrate: Discovery and assessment.

[Azure Migrate](#) helps you to migrate to Azure. Azure Migrate provides a centralized hub to track discovery, assessment, and migration of on-premises infrastructure, applications, and data to Azure. The hub provides Azure tools for assessment and migration, as well as third-party independent software vendor (ISV) offerings.

Before you start

- Make sure you've [created](#) an Azure Migrate project.
- If you've already created a project, make sure you've [added](#) the Azure Migrate: Discovery and assessment tool.
- To create an assessment, you need to set up an Azure Migrate appliance for [VMware vSphere](#), which discovers the on-premises servers, and sends metadata and performance data to Azure Migrate: Discovery and assessment. [Learn more](#).
- You could also [import the server metadata](#) in comma-separated values (CSV) format.

Azure VMware Solution (AVS) Assessment overview

There are three types of assessments you can create using Azure Migrate: Discovery and assessment.

*ASSESSMENT TYPE	DETAILS
Azure VM	Assessments to migrate your on-premises servers to Azure virtual machines. You can assess your on-premises VMs in VMware vSphere and Hyper-V environment, and physical servers for migration to Azure VMs using this assessment type.
Azure SQL	Assessments to migrate your on-premises SQL servers from your VMware environment to Azure SQL Database or Azure SQL Managed Instance.
Azure App Service	Assessments to migrate your on-premises ASP.NET web apps, running on IIS web servers, from your VMware vSphere environment to Azure App Service.
Azure VMware Solution (AVS)	Assessments to migrate your on-premises servers to Azure VMware Solution (AVS) . You can assess your on-premises VMs in VMware vSphere environment for migration to Azure VMware Solution (AVS) using this assessment type. Learn more

NOTE

Azure VMware Solution (AVS) assessment can be created for virtual machines in VMware vSphere environment only.

There are two types of sizing criteria that you can use to create Azure VMware Solution (AVS) assessments:

ASSESSMENT	DETAILS	DATA
Performance-based	Assessments based on collected performance data of on-premises servers.	Recommended Node size: Based on CPU and memory utilization data along with node type, storage type, and FTT setting that you select for the assessment.
As on-premises	Assessments based on on-premises sizing.	Recommended Node size: Based on the on-premises server size along with the node type, storage type, and FTT setting that you select for the assessment.

Run an Azure VMware Solution (AVS) assessment

1. On the **Overview** page > **Servers, databases and web apps**, click **Assess and migrate servers**.
2. In **Azure Migrate: Discovery and assessment**, click **Assess**.
3. In **Assess servers** > **Assessment type**, select **Azure VMware Solution (AVS)**.
4. In **Discovery source**:
 - If you discovered servers using the appliance, select **Servers discovered from Azure Migrate appliance**.
 - If you discovered servers using an imported CSV file, select **Imported servers**.
5. Click **Edit** to review the assessment properties.

[Home](#) > [Azure Migrate](#) >

Create assessment

×

Basics Select servers to assess Review + create assessment

An assessment is created on a group of servers that you migrate together. Assessment helps you determine the Azure readiness of your Windows, Linux and SQL servers running on-premises or on any cloud. You can assess the servers discovered via the Azure Migrate appliance as well as the servers imported into Azure Migrate. [Learn more](#).

Assessment details

Assessment type * ⓘ Azure VMware Solution (AVS) ⓘ [Help me choose](#)

i Creates assessment for migrating on-premises VMware machines to Azure VMware Solution (AVS). Click here to learn more about AVS.

Discovery source * ⓘ Servers discovered from Azure Migrate appliance ⓘ

Assessment properties (Showing 4 of 13) [Edit](#)

Sizing criteria	As on-premises
Target location	East US
Reserved capacity (compute)	No reserved instances
Azure Hybrid Benefit	AHUB benefits do apply to Microsoft based guest OS's running in AVS. For non-Microsoft guest OS please consult your vendor for details.

< Previous [Next](#) >

6. In **Assessment properties** > **Target Properties**:

- In **Target location**, specify the Azure region to which you want to migrate.
 - Size and cost recommendations are based on the location that you specify.
- The **Storage type** is defaulted to **vSAN**. This is the default storage type for an Azure VMware Solution private cloud.

- In **Reserved Instances**, specify whether you want to use reserve instances for Azure VMware Solution nodes when you migrate your VMs.
- If you select to use a reserved instance, you can't specify **'Discount (%)'**
- [Learn more](#)

7. In **VM Size**:

- The **Node type** is defaulted to **AV36**. Azure Migrate recommends the node of nodes needed to migrate the servers to Azure VMware Solution.
- In **FTT setting, RAID level**, select the Failure to Tolerate and RAID combination. The selected FTT option, combined with the on-premises server disk requirement, determines the total vSAN storage required in AVS.
- In **CPU Oversubscription**, specify the ratio of virtual cores associated with one physical core in the AVS node. Oversubscription of greater than 4:1 might cause performance degradation, but can be used for web server type workloads.
- In **Memory overcommit factor**, specify the ratio of memory over commit on the cluster. A value of 1 represents 100% memory use, 0.5 for example is 50%, and 2 would be using 200% of available memory. You can only add values from 0.5 to 10 up to one decimal place.
- In **Dedupe and compression factor**, specify the anticipated deduplication and compression factor for your workloads. Actual value can be obtained from on-premises vSAN or storage config and this may vary by workload. A value of 3 would mean 3x so for 300GB disk only 100GB storage would be used. A value of 1 would mean no dedupe or compression. You can only add values from 1 to 10 up to one decimal place.

8. In **Node Size**:

- In **Sizing criterion**, select if you want to base the assessment on static metadata, or on performance-based data. If you use performance data:
 - In **Performance history**, indicate the data duration on which you want to base the assessment
 - In **Percentile utilization**, specify the percentile value you want to use for the performance sample.
- In **Comfort factor**, indicate the buffer you want to use during assessment. This accounts for issues like seasonal usage, short performance history, and likely increases in future usage. For example, if you use a comfort factor of two:

COMPONENT	EFFECTIVE UTILIZATION	ADD COMFORT FACTOR (2.0)
Cores	2	4
Memory	8 GB	16 GB

9. In **Pricing**:

- In **Offer**, [Azure offer](#) you're enrolled in is displayed. The Assessment estimates the cost for that offer.
- In **Currency**, select the billing currency for your account.
- In **Discount (%)**, add any subscription-specific discounts you receive on top of the Azure offer. The default setting is 0%.

10. Click **Save** if you make changes.

Edit properties ...

Assess2202

TARGET PROPERTIES

Target location ⓘ

East US

Storage type ⓘ

vSAN

Reserved instances ⓘ

No reserved instances

VM SIZE

Node type ⓘ

AV36

FTT setting, RAID level ⓘ

1, RAID-1

CPU Oversubscription ⓘ

4:1

Memory overcommit factor ⓘ

1

Dedupe and compression factor ⓘ

1.5

NODE SIZE

Sizing criterion ⓘ

Performance-based

Performance history ⓘ

Not applicable

Percentile utilization ⓘ

Not applicable

Comfort factor ⓘ

1

PRICING

Offer ⓘ

Pay-As-You-Go

Currency ⓘ

US Dollar (\$)

Discount (%) ⓘ

0

Performance history and Percentile utilization are not applicable for assessments with certain inventory sources. [Learn more.](#)

i AHUB benefits do apply to Microsoft based guest OS's running in AVS. For non-Microsoft guest OS please consult your vendor for details

Save

Discard

11. In **Assess Servers**, click **Next**.
12. In **Select servers to assess** > **Assessment name** > specify a name for the assessment.
13. In **Select or create a group** > select **Create New** and specify a group name.

Basics **Select servers to assess** Review + create assessment

Assessment name

Select or create a group

Create New Use Existing

Add machines to the group [How to create groups using dependency visualization?](#)

Appliance name

Select all [Clear selection](#) [< Previous](#) Page 1 of 7 [Next >](#)

Name	IP address	Operating system	Machine type
<input checked="" type="checkbox"/> SQLTestDBVM1		Microsoft Windows Server 2016 or later (64-...	VMware
<input checked="" type="checkbox"/> SQLTestDBVM55	2404:f801:4800:1c:5d08:de79:d48b:1db8,169...	Microsoft Windows Server 2016 or later (64-...	VMware
<input checked="" type="checkbox"/> SQLTestDBVM46	2404:f801:4800:1c:51acc7a2:46cc:9860,10.15...	Microsoft Windows Server 2016 or later (64-...	VMware
<input checked="" type="checkbox"/> SQLTestDBVM52	2404:f801:4800:1c:d9ffa646:b672:7fae,10.150...	Microsoft Windows Server 2016 or later (64-...	VMware
<input checked="" type="checkbox"/> SQLTestDBVM33	2404:f801:4800:1c:f040:2af7:3f04:2f54,10.150...	Microsoft Windows Server 2016 or later (64-...	VMware
<input checked="" type="checkbox"/> SQLTestDBVM59	2404:f801:4800:1c:dcf3:e2e8:8f4d:477d,10.15...	Microsoft Windows Server 2016 or later (64-...	VMware
<input checked="" type="checkbox"/> SQLTestDBVM35	2404:f801:4800:1c:b059:ec37:89fcaa7a,10.15...	Microsoft Windows Server 2016 or later (64-...	VMware
<input checked="" type="checkbox"/> SQLTestDBVM51	2404:f801:4800:1c:c86d:1797:f9ad:6e47,10.15...	Microsoft Windows Server 2016 or later (64-...	VMware

[< Previous](#) [Next >](#)

14. Select the appliance, and select the servers you want to add to the group. Then click **Next**.

15. In **Review + create assessment**, review the assessment details, and click **Create Assessment** to create the group and run the assessment.

NOTE

For performance-based assessments, we recommend that you wait at least a day after starting discovery before you create an assessment. This provides time to collect performance data with higher confidence. Ideally, after you start discovery, wait for the performance duration you specify (day/week/month) for a high-confidence rating.

Review an Azure VMware Solution (AVS) assessment

An Azure VMware Solution (AVS) assessment describes:

- **Azure VMware Solution (AVS) readiness:** Whether the on-premises VMs are suitable for migration to Azure VMware Solution (AVS).
- **Number of Azure VMware Solution nodes:** Estimated number of Azure VMware Solution nodes required to run the servers.
- **Utilization across AVS nodes:** Projected CPU, memory, and storage utilization across all nodes.
 - Utilization includes up front factoring in the following cluster management overheads such as the vCenter Server, NSX Manager (large), NSX Edge, if HCX is deployed also the HCX Manager and IX appliance consuming ~ 44vCPU (11 CPU), 75GB of RAM and 722GB of storage before compression and deduplication.
 - Limiting factor determines the number of hosts/nodes required to accommodate the resources.
- **Monthly cost estimation:** The estimated monthly costs for all Azure VMware Solution (AVS) nodes running the on-premises VMs.

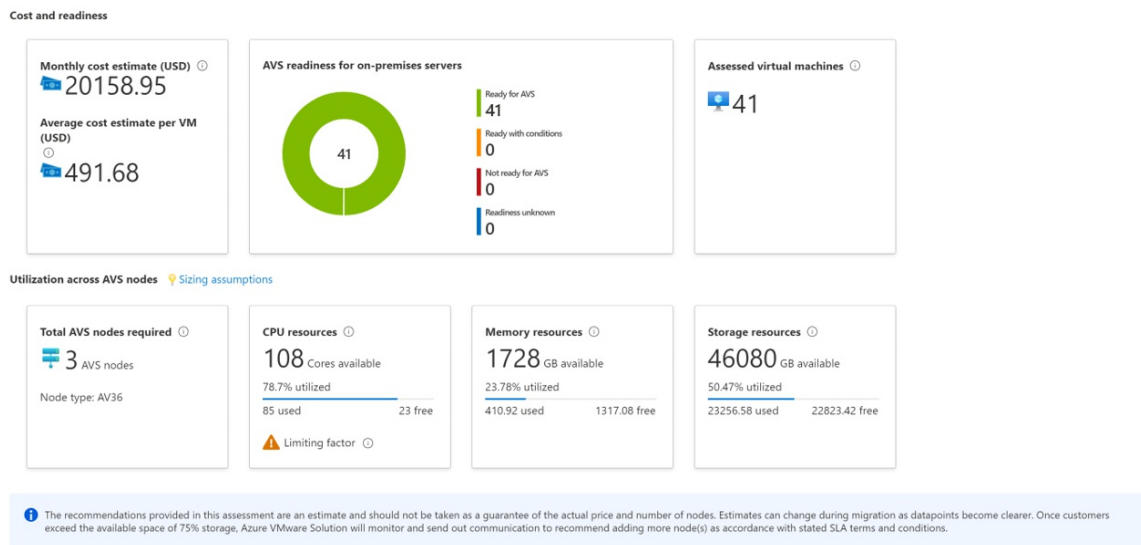
You can click on **Sizing assumptions** to understand the assumptions that went in node sizing and resource utilization calculations. You can also edit the assessment properties, or recalculate the assessment.

View an assessment

1. In **Windows, Linux and SQL Server > Azure Migrate: Discovery and assessment**, click the

number next to **** Azure VMware Solution****.

2. In **Assessments**, select an assessment to open it. As an example (estimations and costs for example only):



3. Review the assessment summary. You can click on **Sizing assumptions** to understand the assumptions that went in node sizing and resource utilization calculations. You can also edit the assessment properties, or recalculate the assessment.

Review Azure VMware Solution (AVS) readiness

1. In **Azure readiness**, verify whether servers are ready for migration to AVS.
2. Review the server status:
 - **Ready for AVS:** The server can be migrated as-is to Azure (AVS) without any changes. It will start in AVS with full AVS support.
 - **Ready with conditions:** There might be some compatibility issues example internet protocol or deprecated OS in VMware and need to be remediated before migrating to Azure VMware Solution. To fix any readiness problems, follow the remediation guidance the assessment suggests.
 - **Not ready for AVS:** The VM will not start in AVS. For example, if the on-premises VMware VM has an external device attached such as a cd-rom the VMware vMotion operation will fail (if using VMware vMotion).
 - **Readiness unknown:** Azure Migrate couldn't determine the readiness of the server because of insufficient metadata collected from the on-premises environment.
3. Review the Suggested tool:
 - **VMware HCX Advanced or Enterprise:** For VMware vSphere VMs, VMware Hybrid Cloud Extension (HCX) solution is the suggested migration tool to migrate your on-premises workload to your Azure VMware Solution (AVS) private cloud. [Learn More](#).
 - **Unknown:** For servers imported via a CSV file, the default migration tool is unknown. Though for VMware vSphere VMs, it is suggested to use the VMware Hybrid Cloud Extension (HCX) solution.
4. Click on an **AVS readiness** status. You can view VM readiness details, and drill down to see VM details, including compute, storage, and network settings.

Review cost details

This view shows the estimated cost of running servers in Azure VMware Solution.

1. Review the monthly total costs. Costs are aggregated for all servers in the assessed group.
 - Cost estimates are based on the number of AVS nodes required considering the resource

requirements of all the servers in total.

- As the pricing for Azure VMware Solution is per node, the total cost does not have compute cost and storage cost distribution.
 - The cost estimation is for running the on-premises servers in AVS. AVS assessment doesn't consider PaaS or SaaS costs.
2. You can review monthly storage cost estimates. This view shows aggregated storage costs for the assessed group, split over different types of storage disks.
 3. You can drill down to see details for specific servers.

Review confidence rating

When you run performance-based assessments, a confidence rating is assigned to the assessment.

- A rating from 1-star (lowest) to 5-star (highest) is awarded.
- The confidence rating helps you estimate the reliability of the size recommendations provided by the assessment.
- The confidence rating is based on the availability of data points needed to compute the assessment.
- For performance-based sizing, AVS assessments need the utilization data for CPU and server memory. The following performance data is collected but not used in sizing recommendations for AVS assessments:
 - The disk IOPS and throughput data for every disk attached to the server.
 - The network I/O to handle performance-based sizing for each network adapter attached to a server.

Confidence ratings for an assessment are as follows.

DATA POINT AVAILABILITY	CONFIDENCE RATING
0%-20%	1 Star
21%-40%	2 Star
41%-60%	3 Star
61%-80%	4 Star
81%-100%	5 Star

[Learn more](#) about performance data

Next steps

- Learn how to use [dependency mapping](#) to create high confidence groups.
- [Learn more](#) about how Azure VMware Solution assessments are calculated.

Create a placement policy in Azure VMware Solution

12/16/2022 • 10 minutes to read • [Edit Online](#)

In Azure VMware Solution, clusters in a private cloud are a managed resource. As a result, the CloudAdmin role can't make certain changes to the cluster from the vSphere Client, including the management of Distributed Resource Scheduler (DRS) rules.

The placement policy feature is available in all Azure VMware Solution regions. Placement policies let you control the placement of virtual machines (VMs) on hosts within a cluster through the Azure portal. When you create a placement policy, it includes a DRS rule in the specified vSphere cluster. It also includes additional logic for interoperability with Azure VMware Solution operations.

A placement policy has at least five required components:

- **Name** - Defines the name of the policy and is subject to the naming constraints of [Azure Resources](#).
- **Type** - Defines the type of control you want to apply to the resources contained in the policy.
- **Cluster** - Defines the cluster for the policy. The scope of a placement policy is a vSphere cluster, so only resources from the same cluster may be part of the same placement policy.
- **State** - Defines if the policy is enabled or disabled. In certain scenarios, a policy might be disabled automatically when a conflicting rule gets created. For more information, see [Considerations](#) below.
- **Virtual machine** - Defines the VMs and hosts for the policy. Depending on the type of rule you create, your policy may require you to specify some number of VMs and hosts. For more information, see [Placement policy types](#) below.

Prerequisite

You must have *Contributor* level access to the private cloud to manage placement policies.

Placement policy types

VM-VM policies

VM-VM policies specify if selected VMs should run on the same host or must be kept on separate hosts. In addition to choosing a name and cluster for the policy, **VM-VM** policies require that you select at least two VMs to assign. The assignment of hosts isn't required or permitted for this policy type.

- **VM-VM Affinity** policies instruct DRS to try to keeping the specified VMs together on the same host. It's useful for performance reasons, for example.
- **VM-VM Anti-Affinity** policies instruct DRS to try keeping the specified VMs apart from each other on separate hosts. It's useful in availability scenarios where a problem with one host doesn't affect multiple VMs within the same policy.

VM-Host policies

VM-Host policies specify if selected VMs can run on selected hosts. To avoid interference with platform-managed operations such as host maintenance mode and host replacement, **VM-Host** policies in Azure VMware Solution are always preferential (also known as "should" rules). Accordingly, **VM-Host** policies [may not be honored in certain scenarios](#). For more information, see [Monitor the operation of a policy](#) below.

Certain platform operations dynamically update the list of hosts defined in **VM-Host** policies. For example, when you delete a host that is a member of a placement policy, the host is removed if more than one host is part of that policy. Also, if a host is part of a policy and needs to be replaced as part of a platform-managed operation, the policy is updated dynamically with the new host.

In addition to choosing a name and cluster for the policy, a **VM-Host** policy requires that you select at least one VM and one host to assign to the policy.

- **VM-Host Affinity** policies instruct DRS to try running the specified VMs on the hosts defined.
- **VM-Host Anti-Affinity** policies instruct DRS to try running the specified VMs on hosts other than those defined.

Considerations

Cluster scale in

Azure VMware Solution attempts to prevent certain DRS rule violations from occurring when performing cluster scale-in operations.

You can't remove the last host from a VM-Host policy. However, if you need to remove the last host from the policy, you can remediate it by adding another host to the policy before removing the host from the cluster. Alternatively, you can delete the placement policy before removing the host.

You can't have a VM-VM Anti Affinity policy with more VMs than the number of hosts in a cluster. If removing a host would result in fewer hosts in the cluster than VMs, you'll receive an error preventing the operation. You can remediate it by first removing VMs from the rule and then removing the host from the cluster.

Rule conflicts

If DRS rule conflicts are detected when you create a VM-VM policy, it results in that policy being created in a disabled state following standard [VMware DRS Rule behavior](#). For more information on viewing rule conflicts, see [Monitor the operation of a policy](#) below.

Create a placement policy

There is no defined limit to the number of policies that you create. However, the more placement constraints you create, the more challenging it is for vSphere DRS to effectively move virtual machines within the cluster and provide the resources needed by the workloads.

Make sure to review the requirements for the [policy type](#).

1. In your Azure VMware Solution private cloud, under **Manage**, select **Placement policies > + Create**.

TIP

You may also select the Cluster from the Placement Policy overview pane and then select **Create**.

2. Provide a descriptive name, select the policy type, and select the cluster where the policy is created. Then select **Enabled**.

WARNING

If you disable the policy, then the policy and the underlying DRS rule are created, but the policy actions are ignored until you enable the policy.

3. If you selected **VM-Host affinity** or **VM-Host anti-affinity** as the type, select **+ Add hosts** and the

hosts to include in the policy. You can select multiple hosts.

NOTE

The select hosts pane shows how many VM-Host policies are associated with the host and the total number of VMs contained in those associated policies.

4. Select **+ Add virtual machine** and the VMs to include in the policy. You can select multiple VMs.

NOTE

The select hosts pane shows how many VM-Host policies are associated with the host and the total number of VMs contained in those associated policies.

5. Once you've finished adding the VMs you want, select **Add virtual machines**.

6. Select **Next: Review and create** to review your policy.

7. Select **Create policy**. If you want to make changes, select **Back: Basics**.

8. After the placement policy gets created, select **Refresh** to see it in the list.

The screenshot shows the Azure portal interface for a private cloud named 'Lamna_Healthcare_PC'. The 'Placement policies' overview is displayed, showing a table of policies. The table has the following data:

Name	Type	Hosts	VMs	Provisioning state	State
web-affinity	VM-VM affinity	N/A	2	Succeeded	Enabled

Edit a placement policy

You can change the state of a policy, add a new resource, or unassign an existing resource.

Change the policy state

You can change the state of a policy to **Enabled** or **Disabled**.

1. In your Azure VMware Solution private cloud, under **Manage**, select **Placement policies**.
2. For the policy you want to edit, select **More (...)** and then select **Edit**.

TIP

You can disable a policy from the Placement policy overview by selecting **Disable** from the Settings drop-down. You can't enable a policy from the Settings drop-down.

3. If the policy is enabled but you want to disable it, select **Disabled** and then select **Disabled** on the confirmation message. Otherwise, if the policy is disabled and you want to enable it, select **Enable**.

4. Select **Review + update**.

5. Review the changes and select **Update policy**. If you want to make changes, select **Back: Basics**.

Update the resources in a policy

You can add new resources, such as a VM or a host, to a policy or remove existing ones.

1. In your Azure VMware Solution private cloud, under **Manage**, select **Placement policies**.

2. For the policy you want to edit, select **More (...)** and then **Edit**.

To remove an existing resource, select one or more resources you want to remove and select **Unassign**.

To add a new resource, select **Edit virtual machine** or **Edit host**, select the resource you'd like to add, and then select **Save**.

3. Select **Next : Review and update**.

4. Review the changes and select **Update policy**. If you want to make changes, select **Back : Basics**.

Delete a policy

You can delete a placement policy and its corresponding DRS rule.

1. In your Azure VMware Solution private cloud, under **Manage**, select **Placement policies**.

2. For the policy you want to edit, select **More (...)** and then select **Delete**.

3. Select **Delete** on the confirmation message.

Monitor the operation of a policy

Use the vSphere Client to monitor the operation of a placement policy's corresponding DRS rule.

As a holder of the CloudAdmin role, you can view, but not edit, the DRS rules created by a placement policy on the cluster's Configure tab under VM/Host Rules. It lets you view additional information, such as if the DRS rules are in a conflict state.

Additionally, you can monitor various DRS rule operations, such as recommendations and faults, from the cluster's Monitor tab.

Restrict VM Movement

For certain extremely sensitive applications vMotion may cause unexpected service interruptions or disruptions. For these types of applications, it may be desirable to restrict VM movement to manually-initiated vMotion only. With the Restrict VM movement Placement Policy, DRS-initiated vMotions can be disabled. For most workloads this is not necessary and may cause unintended performance impacts due to noisy neighbors on the same host.

Enable Restricted VM movement for specific VMs

1. Navigate to Manage Placement policies and click Restrict VM movement.

2. Select the VM or VMs you want to restrict, then click Select.

3. The VM or VMS you selected appears in the VMs with restricted movement tab.

In the vSphere Client, a VM override will be created to set DRS to partially automated for that VM.

DRS will no longer migrate the VM automatically.

Manual vMotion of the VM and automatic initial placement of the VM will continue to function.

FAQs

Are placement policies the same as DRS affinity rules?

Yes, and no. While vSphere DRS implements the current set of policies, we have simplified the experience. Modifying VM groups and Host groups are a cumbersome operation, especially as hosts are ephemeral in nature and could be replaced in a cloud environment. As hosts are replaced in the vSphere inventory in an on-premises environment, the vSphere admin must modify the host group to ensure that the desired VM-Host placement constraints remain in effect. Placement policies in Azure VMware Solution update the Host groups when a host is rotated or changed. Similarly, if you scale in a cluster, the Host Group is automatically updated, as applicable. This eliminates the overhead of managing the Host Groups for the customer.

As this is an existing functionality available in vCenter, why can't I use it directly?

Azure VMware Solution provides a VMware private cloud in Azure. In this managed VMware infrastructure, Microsoft manages the clusters, hosts, datastores, and distributed virtual switches in the private cloud. At the same time, the tenant is responsible for managing the workloads deployed on the private cloud. As a result, the tenant administering the private cloud [does not have the same set of privileges](#) as available to the VMware administrator in an on-premises deployment.

Further, the lack of the desired granularity in the vSphere privileges presents some challenges when managing the placement of the workloads on the private cloud. For example, vSphere DRS rules commonly used on-premises to define affinity and anti-affinity rules can't be used as-is in an Azure VMware Solution environment, as some of those rules can block day-to-day operation the private cloud. Placement Policies provides a way to define those rules using the Azure VMware Solution portal, thereby circumventing the need to use DRS rules. Coupled with a simplified experience, they also ensure that the rules don't impact the day-to-day infrastructure maintenance and operation activities.

What is the difference between the VM-Host affinity policy and Restrict VM movement?

A VM-Host affinity policy is used to restrict the movement of VMs to a group of hosts included in the VM-Host affinity policy. Thus, a VM can be vMotioned within the set of hosts selected in the VM-Host affinity policy. Alternatively, **Restrict VM movement** ensures that the selected VM remains on the host on which it currently resides.

What caveats should I know about?

The VM-Host **MUST** rules aren't supported because they block maintenance operations.

VM-Host **SHOULD** rules are preferential rules, where vSphere DRS tries to accommodate the rules to the extent possible. Occasionally, vSphere DRS may vMotion VMs subjected to the VM-Host **SHOULD** rules to ensure that the workloads get the resources they need. It's a standard vSphere DRS behavior, and the Placement policies feature does not change the underlying vSphere DRS behavior.

If you create conflicting rules, those conflicts may show up on the vCenter Server, and the newly defined rules may not take effect. It's a standard vSphere DRS behavior, the logs for which can be observed in the vCenter Server.

Configure GitHub Enterprise Server on Azure VMware Solution

12/16/2022 • 7 minutes to read • [Edit Online](#)

In this article, you'll set up GitHub Enterprise Server, the "on-premises" version of [GitHub.com](#), on your Azure VMware Solution private cloud. The scenario covers a GitHub Enterprise Server instance that can serve up to 3,000 developers running up to 25 jobs per minute on GitHub Actions. It includes the setup of (at time of writing) *preview* features, such as GitHub Actions. To customize the setup for your particular needs, review the requirements listed in [Installing GitHub Enterprise Server on VMware](#).

Before you begin

GitHub Enterprise Server requires a valid license key. You may sign up for a [trial license](#). If you're looking to extend the capabilities of GitHub Enterprise Server via an integration, you may qualify for a free five-seat developer license. Apply for this license through [GitHub's Partner Program](#).

Install GitHub Enterprise Server on VMware

1. Download [the current release of GitHub Enterprise Server](#) for VMware ESXi/vSphere (OVA) and [deploy the OVA template](#) you downloaded.

GitHub On-premises

Choose this option if you are running GitHub on your own hardware. Download the image below, then launch a new VM with the image.

Download for VMware ESXi/vSphere (OVA)

SHA256:
eab2f77fe728a7b8a6c8b8a8e105dd462ad7f753f5e55904d7cf7cd9b1eaddf6

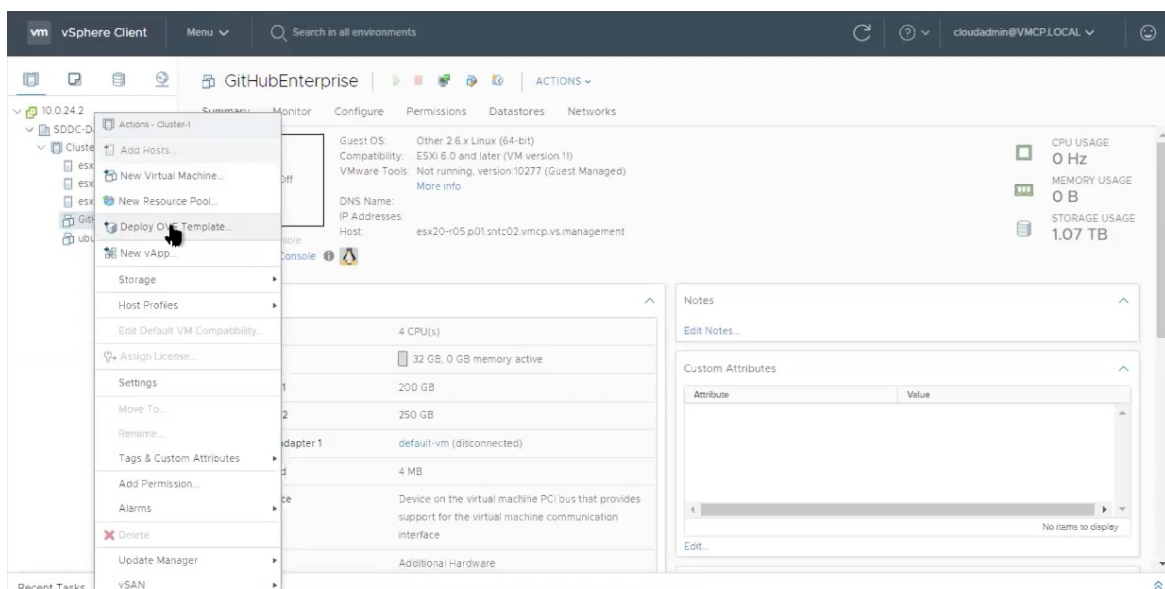
Need help? Check out our [administration guides](#).

[Change Hypervisor](#)

GitHub in the Cloud

Choose this option if you are installing or running GitHub on a cloud service such as Amazon Web Services, Microsoft Azure, or Google Cloud Platform.

Select your platform



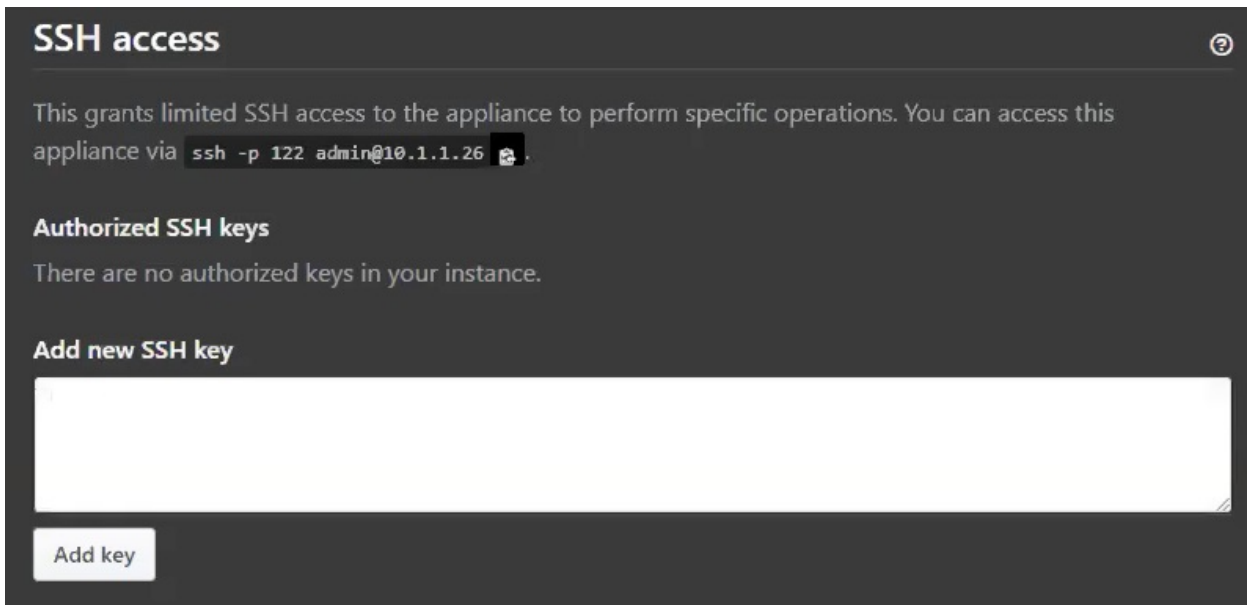
2. Provide a recognizable name for your new virtual machine, such as GitHubEnterpriseServer. You don't need to include the release details in the VM name, as these details become stale when the instance is upgraded.
3. Select all the defaults for now (we'll edit these details shortly) and wait for the OVA to be imported.
4. Once imported, [adjust the hardware configuration](#) based on your needs. In our example scenario, we'll need the following configuration.

RESOURCE	STANDARD SETUP	STANDARD SET UP + "BETA FEATURES" (ACTIONS)
vCPUs	4	8
Memory	32 GB	61 GB
Attached storage	250 GB	300 GB
Root storage	200 GB	200 GB

Your needs may vary. Refer to the guidance on hardware considerations in [Installing GitHub Enterprise Server on VMware](#). Also see [Adding CPU or memory resources for VMware](#) to customize the hardware configuration based on your situation.

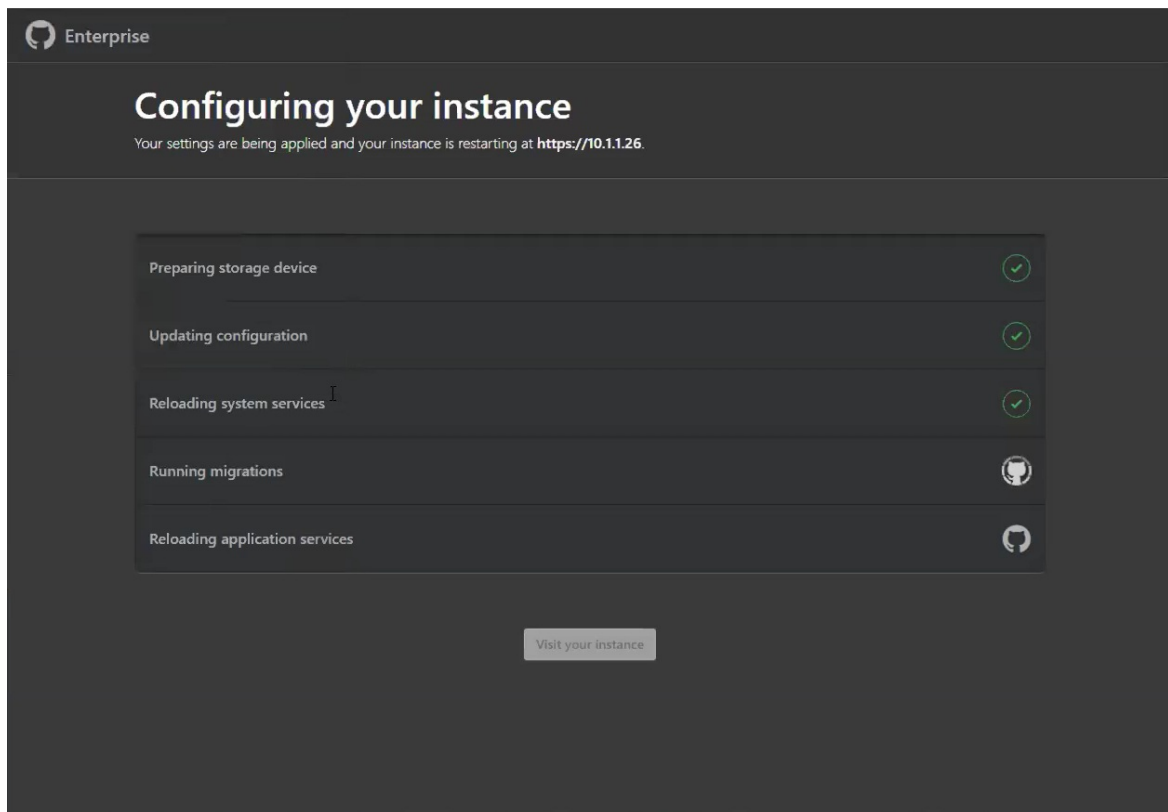
Configure the GitHub Enterprise Server instance

After the newly provisioned virtual machine (VM) has powered on, [configure it through your browser](#). You'll be required to upload your license file and set a management console password. Be sure to write down this password somewhere safe.



We recommend at least take the following steps:

1. Upload a public SSH key to the management console so that you can [access the administrative shell via SSH](#).
2. [Configure TLS on your instance](#) so that you can use a certificate signed by a trusted certificate authority. Apply your settings.



3. While the instance restarts, configure blob storage for GitHub Actions.

NOTE

GitHub Actions is [currently available as a limited beta on GitHub Enterprise Server release 2.22](#).

External blob storage is necessary to enable GitHub Actions on GitHub Enterprise Server (currently available as a "beta" feature). Actions use this external blob storage to store artifacts and logs. Actions on

GitHub Enterprise Server [supports Azure Blob Storage as a storage provider](#) (and some others). So we'll provision a new Azure storage account with a [storage account type](#) of BlobStorage.

Instance details

The default deployment model is Resource Manager, which supports the latest Azure features. You may choose to deploy using the classic deployment model instead. [Choose classic deployment model](#)

Storage account name *	<input type="text" value="ghesactionsstoragebd0903"/>
Location *	<input type="text" value="(US) East US"/>
Performance	<input checked="" type="radio"/> Standard <input type="radio"/> Premium
Account kind	<input type="text" value="BlobStorage"/>
Replication	<input type="text" value="Read-access geo-redundant storage (RA-GRS)"/>
	i Accounts with the selected kind, replication and performance type only support block and append blobs. Page blobs, file shares, tables, and queues will not be available.
Blob access tier (default)	<input type="radio"/> Cool <input checked="" type="radio"/> Hot

- Once the new BlobStorage resource deployment completes, save the connection string (available under Access keys). You'll need this string shortly.
- After the instance restarts, create a new admin account on the instance. Be sure to make a note of this user's password as well.

The screenshot shows the 'Create admin account' page in the GitHub Enterprise interface. At the top right, there is a 'Sign in' link. The main heading is 'Create admin account'. Below the heading, there are four input fields: 'Username *' with the value 'octocat', 'Email address *' with the value 'octocat@github.com', 'Password *' (masked with dots), and 'Confirm your password *' (also masked with dots). Below the password fields, there is a checkbox labeled 'Help me set up an organization next' which is checked. Underneath this checkbox, there is a short paragraph explaining that organizations are separate from personal accounts and are best suited for businesses. At the bottom of the form is a blue button labeled 'Create admin account'.

Other configuration steps

To harden your instance for production use, the following optional setup steps are recommended:

- Configure [high availability](#) for protection against:
 - Software crashes (OS or application level)
 - Hardware failures (storage, CPU, RAM, and so on)

- Virtualization host system failures
 - Logically or physically severed network
2. [Configure backup-utilities](#), providing versioned snapshots for disaster recovery, hosted in availability that is separate from the primary instance.
 3. [Setup subdomain isolation](#), using a valid TLS certificate, to mitigate cross-site scripting and other related vulnerabilities.

Set up the GitHub Actions runner

NOTE

GitHub Actions is [currently available as a limited beta on GitHub Enterprise Server release 2.22](#).

At this point, you should have an instance of GitHub Enterprise Server running, with an administrator account created. You should also have external blob storage that GitHub Actions uses for persistence.

Create somewhere for GitHub Actions to run; again, we'll use Azure VMware Solution.

1. Provision a new VM on the cluster and base it on [a recent release of Ubuntu Server](#).

New Virtual Machine

1 Select a creation type

2 Select a name and folder

3 Select a compute resource

4 Select storage

5 Select compatibility

6 Select a guest OS

7 Customize hardware

8 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine name:

Select a location for the virtual machine.

- 10.0.24.2
 - SDDC-Datacenter

CANCEL BACK NEXT

2. Continue through the set up selecting the compute resource, storage, and compatibility.
3. Select the guest OS you want installed on the VM.

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Select storage
- ✓ 5 Select compatibility
- 6 Select a guest OS**
- 7 Customize hardware
- 8 Ready to complete

Select a guest OS
Choose the guest OS that will be installed on the virtual machine

Identifying the guest operating system here allows the wizard to provide the appropriate defaults for the operating system installation.

Guest OS Family:

Guest OS Version:

Compatibility: ESXi 6.7 and later (VM version 14)

[CANCEL](#) [BACK](#) [NEXT](#)

4. Once the VM is created, power it up and connect to it via SSH.
5. Install [the Actions runner](#) application, which runs a job from a GitHub Actions workflow. Identify and download the most current Linux x64 release of the Actions runner, either from [the releases page](#) or by running the following quick script. This script requires both `curl` and `jq` to be present on your VM.

```
LATEST_RELEASE_ASSET_URL=$( curl https://api.github.com/repos/actions/runner/releases/latest | \
jq -r '.assets | .[] | select(.name | match("actions-runner-linux-arm64")) | .url' )

DOWNLOAD_URL=$( curl $LATEST_RELEASE_ASSET_URL | \
jq -r '.browser_download_url' )

curl -OL $DOWNLOAD_URL
```

You should now have a file locally on your VM, `actions-runner-linux-arm64-*.tar.gz`. Extract this tarball locally:

```
tar xzf actions-runner-linux-arm64-*.tar.gz
```

This extraction unpacks a few files locally, including a `config.sh` and `run.sh` script.

Enable GitHub Actions

NOTE

GitHub Actions is [currently available as a limited beta on GitHub Enterprise Server release 2.22](#).

Configure and enable GitHub Actions on the GitHub Enterprise Server instance.

1. [Access the GitHub Enterprise Server instance's administrative shell over SSH](#), and then run the following commands:
2. Set an environment variable containing your Blob storage connection string.

```
export CONNECTION_STRING="<your connection string from the blob storage step>"
```

3. Configure actions storage.

```
ghe-config secrets.actions.storage.blob-provider azure  
ghe-config secrets.actions.storage.azure.connection-string "$CONNECTION_STRING"
```

4. Apply the settings.

```
ghe-config-apply
```

5. Execute a precheck to install additional software required by Actions on GitHub Enterprise Server.

```
ghe-actions-precheck -p azure -cs "$CONNECTION_STRING"
```


6. Enable actions, and re-apply the configuration.

```
ghe-config app.actions.enabled true  
ghe-config-apply
```

7. Check the health of your blob storage.

```
ghe-actions-check -s blob
```

You should see output: *Blob Storage is healthy.*

8. Now that **GitHub Actions** is configured, enable it for your users. Sign in to your GitHub Enterprise Server instance as an administrator, and select the  in the upper right corner of any page.
9. In the left sidebar, select **Enterprise overview**, then **Policies, Actions**, and select the option to **enable Actions for all organizations**.
10. Configure your runner from the **Self-hosted runners** tab. Select **Add new** and then **New runner** from the drop-down. You'll be presented with a set of commands to run.
11. Copy the command to **configure** the runner, for instance:

```
./config.sh --url https://10.1.1.26/enterprises/octo-org --token AAAAAA5RHF34QLYBDCHWLJC7L73MA
```

12. Copy the `config.sh` command and paste it into a session on your Actions runner (created previously).

```

ghadmin@runner001: ~/actions-runner.3

GitHub Actions
Self-hosted runner registration

# Authentication

✓ Connected to GitHub

# Runner Registration

Enter the name of runner: [press Enter for runner001] example

This runner will have the following labels: 'self-hosted', 'Linux', 'X64'
Enter any additional labels (ex. label-1,label-2): [press Enter to skip]

✓ Runner successfully added
✓ Runner connection is good

# Runner settings

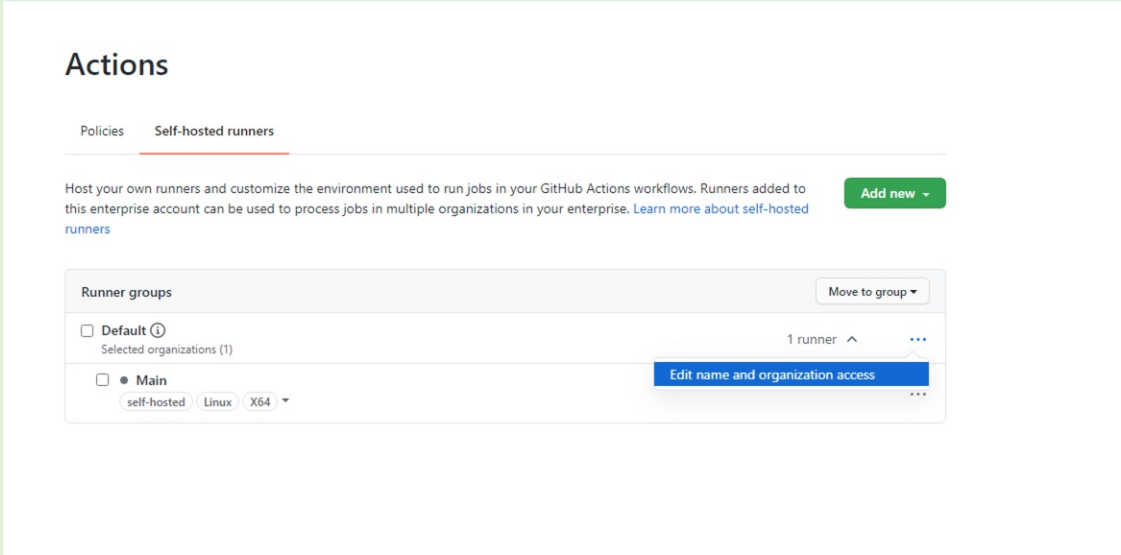
Enter name of work folder: [press Enter for _work]

```

13. Use the `./run.sh` command to *run* the runner:

TIP

To make this runner available to organizations in your enterprise, edit its organization access. You can limit access to a subset of organizations, and even to specific repositories.



The screenshot shows the GitHub Actions interface for self-hosted runners. At the top, there are tabs for 'Policies' and 'Self-hosted runners'. Below this is a description of self-hosted runners and an 'Add new' button. A section titled 'Runner groups' contains a table with the following data:

Group Name	Selected Organizations	Count	Actions
Default	(1)	1 runner	...
Main	self-hosted, Linux, X64		Edit name and organization access, ...

(Optional) Configure GitHub Connect


Although this step is optional, we recommend it if you plan to consume open-source actions available on GitHub.com. It allows you to build on the work of others by referencing these reusable actions in your workflows.

To enable GitHub Connect, follow the steps in [Enabling automatic access to GitHub.com actions using GitHub Connect](#).

Once GitHub Connect is enabled, select the **Server to use actions from GitHub.com** in workflow runs option.

Server can use actions from GitHub.com in workflow runs

Allow actions from GitHub.com to be referenced and used in workflow files.



Set up and run your first workflow

Now that Actions and GitHub Connect is set up, let's put all this work to good use. Here's an example workflow that references the excellent [octokit/request-action](#), allowing us to "script" GitHub through interactions using the GitHub API, powered by GitHub Actions.

In this basic workflow, we'll use `octokit/request-action` to open an issue on GitHub using the API.

```
name: Open Issue

on:
  push

jobs:
  my-job:|
    name: My Job
    runs-on: self-hosted
    steps:
      - name: Open issue
        uses: octokit/request-action@v2.x
        id: issue
        with:
          route: POST /repos/:repository/issues
          repository: ${{ github.repository }}
          title: Hello world
        env:
          GITHUB_TOKEN: ${{ secrets.GITHUB_TOKEN }}
      - name: Log issue
        run: |
          echo issue #${{ fromJson(steps.issue.outputs.data).id }}

opened
```

NOTE

GitHub.com hosts the action, but when it runs on GitHub Enterprise Server, it *automatically* uses the GitHub Enterprise Server API.

If you chose not to enable GitHub Connect, you could use the following alternative workflow.

```

name: Open Issue

on:
  push

jobs:
  my-job:
    name: My Job
    runs-on: self-hosted
    steps:
      - name: Open issue
        run: |
          curl -H "Authorization: bearer $GITHUB_TOKEN" -d '{"title":
"Hello world"}' "$GITHUB_API_URL/repos/${{ github.repository }}/issues"
        env:
          GITHUB_TOKEN: ${{ secrets.GITHUB_TOKEN }}

```


1. Navigate to a repo on your instance, and add the above workflow as: `.github/workflows/hello-world.yml`

```

29 lines (27 sloc) | 775 Bytes
1 name: My Workflow
2
3 on:
4   push:
5     branches:      # array of glob patterns matching against refs/heads. Optional; defaults to all
6     - master      # triggers on pushes that contain changes in master
7
8 jobs:
9   my-job:
10    name: My Job
11    runs-on: self-hosted
12    steps:
13      - name: Context
14        env:
15          GITHUB_CONTEXT: ${{ toJson(github) }}
16        run: |
17          echo "$GITHUB_CONTEXT"
18      - name: Open issue
19        uses: octokit/request-action@v2.x
20        id: issue
21        with:
22          route: POST /repos/:repository/issues
23          repository: ${{ github.repository }}
24          title: Hello world
25        env:
26          GITHUB_TOKEN: ${{ secrets.GITHUB_TOKEN }}
27      - name: Log issue
28        run: |
29          echo issue #${{ fromJson(steps.issue.outputs.data).id }} opened

```


2. In the Actions tab for your repo, wait for the workflow to execute.



You can see it being processed by the runner.

```
ghadmin@runner001: ~/actions-runner.2
ghadmin@runner001:~/actions-runner.2$ ./run.sh
✓ Connected to GitHub
2020-09-14 22:33:40Z: Listening for Jobs
2020-09-14 22:34:53Z: Running job: My Job
2020-09-14 22:35:00Z: Job My Job completed with result: Succeeded
```

If everything ran successfully, you should see a new issue in your repo, entitled "Hello world."



Congratulations! You just completed your first Actions workflow on GitHub Enterprise Server, running on your Azure VMware Solution private cloud.

This article set up a new instance of GitHub Enterprise Server, the self-hosted equivalent of GitHub.com, on top of your Azure VMware Solution private cloud. The instance includes support for GitHub Actions and uses Azure Blob Storage for persistence of logs and artifacts. But we're just scratching the surface of what you can do with GitHub Actions. Check out the list of Actions on [GitHub's Marketplace](#), or [create your own](#).

Next steps

Now that you've covered setting up GitHub Enterprise Server on your Azure VMware Solution private cloud, you may want to learn about:

- [How to get started with GitHub Actions](#)
- [How to join the beta program](#)
- [Administration of GitHub Enterprise Server](#)

Configure external identity source for vCenter Server

12/16/2022 • 11 minutes to read • [Edit Online](#)

In Azure VMware Solution, vCenter Server has a built-in local user called *cloudadmin* assigned to the CloudAdmin role. You can configure users and groups in Active Directory (AD) with the CloudAdmin role for your private cloud. In general, the CloudAdmin role creates and manages workloads in your private cloud. But in Azure VMware Solution, the CloudAdmin role has vCenter Server privileges that differ from other VMware cloud solutions and on-premises deployments.

IMPORTANT

The local *cloudadmin* user should be treated as an emergency access account for "break glass" scenarios in your private cloud. It's not for daily administrative activities or integration with other services.

- In a vCenter Server and ESXi on-premises deployment, the administrator has access to the vCenter Server administrator@vsphere.local account and the ESXi root account. They can also have more AD users and groups assigned.
- In an Azure VMware Solution deployment, the administrator doesn't have access to the administrator user account or the ESXi root account. They can, however, assign AD users and groups to the CloudAdmin role in vCenter Server. The CloudAdmin role doesn't have permissions to add an identity source like on-premises LDAP or LDAPS server to vCenter Server. However, you can use Run commands to add an identity source and assign cloudadmin role to users and groups.

The private cloud user doesn't have access to and can't configure specific management components Microsoft supports and manages. For example, clusters, hosts, datastores, and distributed virtual switches.

NOTE

In Azure VMware Solution, the *vsphere.local/SSO* domain is provided as a managed resource to support platform operations. It doesn't support the creation and management of local groups and users except for those provided by default with your private cloud.

NOTE

Run commands are executed one at a time in the order submitted.

In this article, you learn how to:

- Export the certificate for LDAPS authentication
- Upload the LDAPS certificate to blob storage and generate a SAS URL
- Configure NSX-T DNS for resolution to your Active Directory Domain
- Add Active Directory over (Secure) LDAPS (LDAP over SSL) or (unsecure) LDAP
- Add existing AD group to cloudadmin group
- List all existing external identity sources integrated with vCenter Server SSO
- Assign additional vCenter Server Roles to Active Directory Identities

- Remove AD group from the cloudadmin role
- Remove existing external identity sources

Prerequisites

- Connectivity from your Active Directory network to your Azure VMware Solution private cloud must be operational.
- For AD authentication with LDAPS:
 - You'll need access to the Active Directory Domain Controller(s) with Administrator permissions.
 - Your Active Directory Domain Controller(s) must have LDAPS enabled with a valid certificate. The certificate could be issued by an [Active Directory Certificate Services Certificate Authority \(CA\)](#) or a [Third-party/Public CA](#).
 - You need to have a valid certificate. To create a certificate, follow the steps shown in [create a certificate for secure LDAP](#). Make sure the certificate meets the requirements that are listed after the steps you used to create a certificate for secure LDAP.

NOTE

Self-sign certificates are not recommended for production environments.

- [Export the certificate for LDAPS authentication](#) and upload it to an Azure Storage account as blob storage. Then, you'll need to [grant access to Azure Storage resources using shared access signature \(SAS\)](#).
- Ensure Azure VMware Solution has DNS resolution configured to your on-premises AD. Enable DNS Forwarder from Azure portal. See [Configure DNS forwarder for Azure VMware Solution](#) for further information.


NOTE

For more information about LDAPS and certificate issuance, see with your security or identity management team.

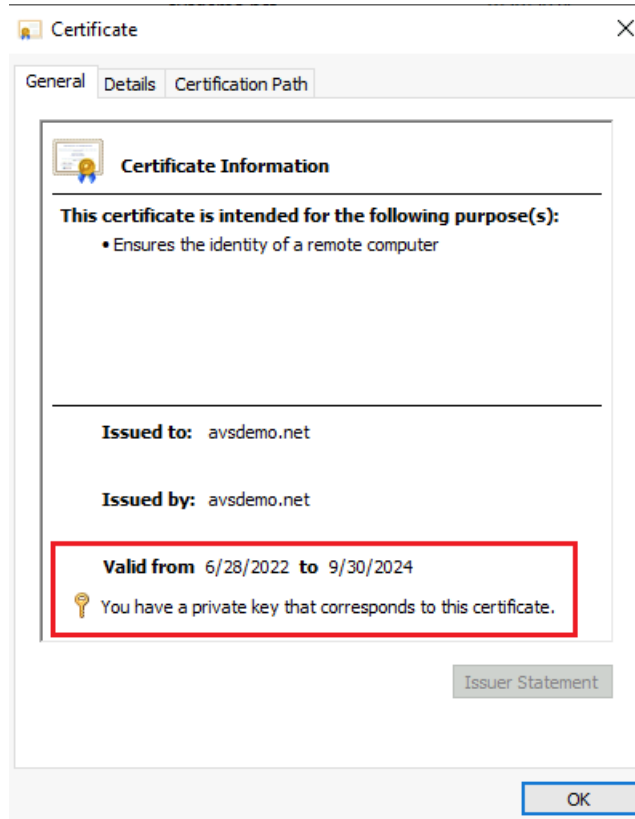
Export the certificate for LDAPS authentication

First, verify that the certificate used for LDAPS is valid. If you don't already have a certificate, follow the steps to [create a certificate for secure LDAP](#) before you continue.

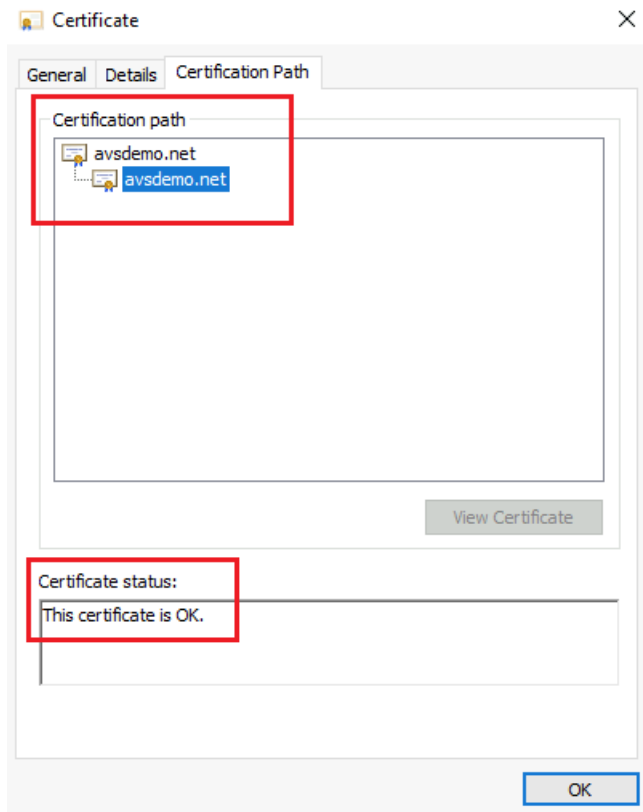
1. Sign in to a domain controller with administrator permissions where LDAPS is enabled.
2. Open the **Run command**, type **mmc** and select the **OK** button.
3. Select the **File** menu option then **Add/Remove Snap-in**.
4. Select the **Certificates** in the list of Snap-ins and select the **Add>** button.
5. In the **Certificates snap-in** window, select **Computer account** then select **Next**.
6. Keep the first option selected **Local computer...**, and select **Finish**, and then **OK**.
7. Expand the **Personal** folder under the **Certificates (Local Computer)** management console and select the **Certificates** folder to list the installed certificates.



- Double click the certificate for LDAPS purposes. The **Certificate** General properties will display. Ensure the certificate date **Valid from** and **to** is current and the certificate has a **private key** that corresponds to the certificate.



- On the same window, select the **Certification Path** tab and verify that the **Certification path** is valid, which it should include the certificate chain of root CA and optionally intermediate certificates and the **Certificate Status** is OK.



10. Close the window.

Now proceed to export the certificate

1. Still on the Certificates console, right select the LDAPS certificate and select **All Tasks > Export**. The Certificate Export Wizard prompt is displayed, select the **Next** button.
2. In the **Export Private Key** section, select the second option, **No, do not export the private key** and select the **Next** button.
3. In the **Export File Format** section, select the second option, **Base-64 encoded X.509(.CER)** and then select the **Next** button.
4. In the **File to Export** section, select the **Browse...** button and select a folder location where to export the certificate, enter a name then select the **Save** button.

NOTE

If more than one domain controller is LDAPS enabled, repeat the export procedure in the additional domain controller(s) to also export the corresponding certificate(s). Be aware that you can only reference two LDAPS server in the `New-LDAPSIIdentitySource` Run Command. If the certificate is a wildcard certificate, for example `*.avsdemo.net` you only need to export the certificate from one of the domain controllers.

Upload the LDAPS certificate to blob storage and generate a SAS URL

- Upload the certificate file (.cer format) you just exported to an Azure Storage account as blob storage. Then [grant access to Azure Storage resources using shared access signature \(SAS\)](#).
- If multiple certificates are required, upload each certificate individually and for each certificate, generate a SAS URL.

IMPORTANT

Make sure to copy each SAS URL string(s), because they will no longer be available once you leave the page.

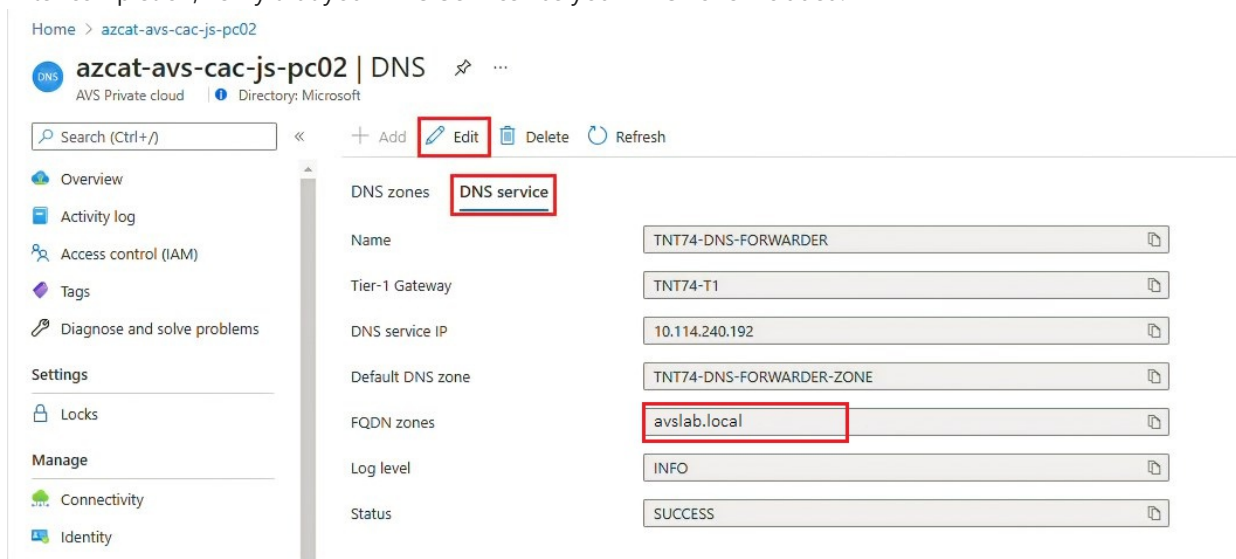
TIP

Another alternative method for consolidating certificates is saving the certificate chains in a single file as mentioned in [this VMware KB article](#), and generate a single SAS URL for the file that contains all the certificates.

Configure NSX-T DNS for resolution to your Active Directory Domain

A DNS Zone needs to be created and added to the DNS Service, follow the instructions in [Configure a DNS forwarder in the Azure portal](#) to complete these two steps.

After completion, verify that your DNS Service has your DNS zone included.



Your Azure VMware Solution Private cloud should now be able to resolve your on-premises Active Directory domain name properly.

Add Active Directory over LDAP with SSL

In your Azure VMware Solution private cloud, you'll run the `New-LDAPSIIdentitySource` cmdlet to add an AD over LDAP with SSL as an external identity source to use with SSO into vCenter Server.

1. Browse to your Azure VMware Solution private cloud and then select **Run command > Packages > New-LDAPSIIdentitySource**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
GroupName	The group in the external identity source that gives the cloudadmin access. For example, <code>avs-admins</code> .
CertificateSAS	Path to SAS strings with the certificates for authentication to the AD source. If you're using multiple certificates, separate each SAS string with a comma. For example, <code>pathtocert1,pathtocert2</code> .

FIELD	VALUE
Credential	The domain username and password used for authentication with the AD source (not cloudadmin). The user must be in the username@avslab.local format.
BaseDNGroups	Where to look for groups, for example, CN=group1, DC=avslab,DC=local . Base DN is needed to use LDAP Authentication.
BaseDNUsers	Where to look for valid users, for example, CN=users,DC=avslab,DC=local . Base DN is needed to use LDAP Authentication.
PrimaryUrl	Primary URL of the external identity source, for example, ldaps://yourserver.avslab.local:636 .
SecondaryURL	Secondary fall-back URL if there's primary failure. For example, ldaps://yourbackupldapserverserver.avslab.local:636 .
DomainAlias	For Active Directory identity sources, the domain's NetBIOS name. Add the NetBIOS name of the AD domain as an alias of the identity source. Typically the <i>*avslab*</i> format.
DomainName	The FQDN of the domain, for example avslab.local .
Name	User-friendly name of the external identity source. For example, avslab.local , is how it will be displayed in vCenter.
Retain up to	Retention period of the cmdlet output. The default value is 60 days.
Specify name for execution	Alphanumeric name, for example, addexternalIdentity .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** or the **Run Execution Status** pane to see the progress and successful completion.

Add Active Directory over LDAP

NOTE

We recommend you use the [Add Active Directory over LDAP with SSL](#) method.

You'll run the `New-LDAPIdentitySource` cmdlet to add AD over LDAP as an external identity source to use with SSO into vCenter Server.

1. Select **Run command > Packages > New-LDAPIdentitySource**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
Name	User-friendly name of the external identity source, for example, avslab.local . This is how it will be displayed in vCenter.
DomainName	The FQDN of the domain, for example avslab.local .
DomainAlias	For Active Directory identity sources, the domain's NetBIOS name. Add the NetBIOS name of the AD domain as an alias of the identity source. Typically the <i>*avsdap*</i> format.
PrimaryUrl	Primary URL of the external identity source, for example, ldap://yourserver.avslab.local:389 .
SecondaryURL	Secondary fall-back URL if there's primary failure.
BaseDNUsers	Where to look for valid users, for example, CN=users,DC=avslab,DC=local . Base DN is needed to use LDAP Authentication.
BaseDNGroups	Where to look for groups, for example, CN=group1,DC=avslab,DC=local . Base DN is needed to use LDAP Authentication.
Credential	The domain username and password used for authentication with the AD source (not cloudadmin). The user must be in the username@avslab.local format.
GroupName	The group to give cloudadmin access in your external identity source, for example, avs-admins .
Retain up to	Retention period of the cmdlet output. The default value is 60 days.
Specify name for execution	Alphanumeric name, for example, addexternalidentity .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** or the **Run Execution Status** pane to see the progress.

Add existing AD group to cloudadmin group

You'll run the `Add-GroupToCloudAdmins` cmdlet to add an existing AD group to a cloudadmin group. Users in the cloudadmin group have privileges equal to the cloudadmin (cloudadmin@vsphere.local) role defined in vCenter Server SSO.

1. Select **Run command > Packages > Add-GroupToCloudAdmins**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
-------	-------

FIELD	VALUE
GroupName	Name of the group to add, for example, VcAdminGroup .
Retain up to	Retention period of the cmdlet output. The default value is 60 days.
Specify name for execution	Alphanumeric name, for example, addADgroup .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** or the **Run Execution Status** pane to see the progress.

List external identity

You'll run the `Get-ExternalIdentitySources` cmdlet to list all external identity sources already integrated with vCenter Server SSO.

1. Sign in to the [Azure portal](#).
2. Select **Run command** > **Packages** > **Get-ExternalIdentitySources**.

The screenshot shows the Azure portal interface for a resource named 'Contoso-westus-sddc'. The 'Run command' pane is open, and the 'Packages' tab is selected. A list of cmdlets is displayed, with 'Get-ExternalIdentitySources' highlighted. The interface includes a search bar, a refresh button, and a sidebar with navigation options like Overview, Activity log, and Settings.

Name	Description
Install-JetDR	This top level Cmdlet Downloads JetDr bundle from MMS, creates a new user, assigns elevated privileges to the user, deploys JetDr Management Server Appliance(MSA), registers vCenter to the JetDr MSA, configures cluster.
Invoke-PreflightJetDRInstall	This Cmdlet checks and displays current state of the system it checks whether the minimal requirements for the script to run are met. It also checks if the cluster has minimum of 4 hosts, if the cluster details are correct, if there is already a VM with the same name provided for installing MSA, if there is any jetdr plugin present in the vCenter.
Invoke-PreflightJetDRSystemCheck	This Cmdlet checks and displays current state of the system it checks whether the minimal requirements for the script to run are met.
Invoke-PreflightJetDRUninstall	This Cmdlet checks and displays current state of the system it checks whether the minimal requirements for the script to run are met. It also checks if the cluster has minimum of 4 hosts, if the cluster details are correct and if any vCenter is registered to the MSA
Uninstall-JetDR	The top level Cmdlet creates a new user, assigns elevated privileges to the user, unconfigures cluster, unregisters vCenter from the JetDr MSA, removes the user.
Add-GroupToCloudAdmins	Add group to Cloud Admin
Get-ExternalIdentitySources	Get all current external sources connected to vCenter SSO
New-AvsLDAPIdentitySource	Allow customers to add an LDAP Secure external identity source (Active Directory over LDAP) for use with single sign on to vCenter.
New-AvsLDAPSIdentitySource	Allow customers to add an LDAPS Secure external identity source (Active Directory over LDAP) for use with single sign on to vCenter.
Remove-ExternalIdentitySources	Remove all external identity sources
Remove-GroupFromCloudAdmins	Remove previously added AD group from CloudAdmins
Set-AvsVMStoragePolicy	Set the storage policy on a VM

3. Provide the required values or change the default values, and then select **Run**.

Run command - Get-ExternalIdentitySources ×

Get all current external sources connected to vCenter SSO

Details

Retain up to

day

hour

minute

Specify name for execution *

Timeout *

hour

minute

second

Run

FIELD	VALUE
Retain up to	Retention period of the cmdlet output. The default value is 60 days.
Specify name for execution	Alphanumeric name, for example, <code>getExternalIdentity</code> .
Timeout	The period after which a cmdlet exits if taking too long to finish.

4. Check **Notifications** or the **Run Execution Status** pane to see the progress.

Packages **Run execution status**

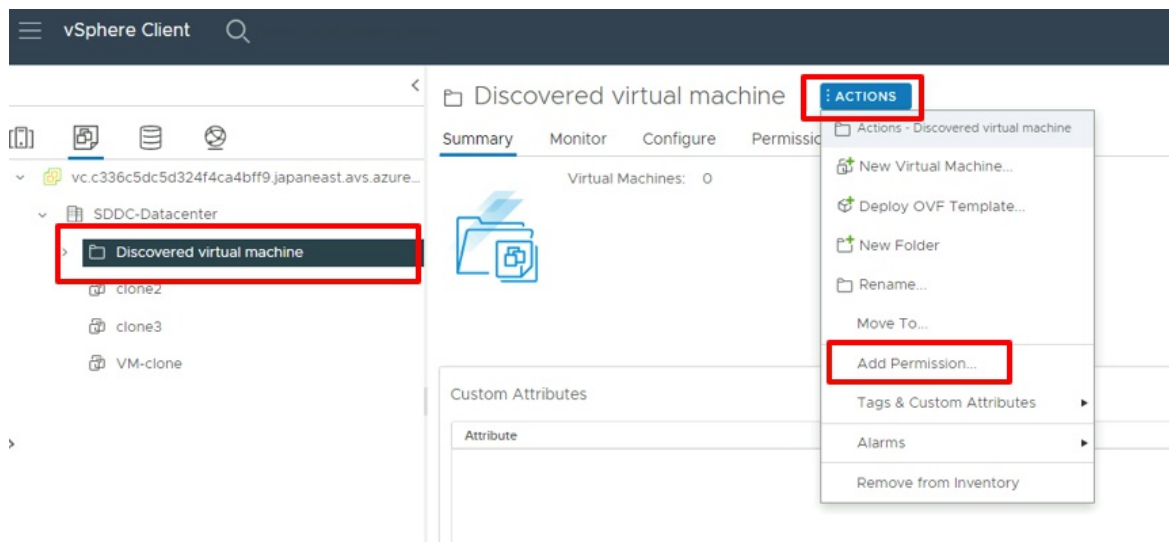
ⓘ Commands are executed one at a time in the order submitted.

Execution name	Package name	Package version	Command name	Started time... ↓	End time stamp	Status
Get-StoragePolicies-Exec1	Microsoft.AVS.Management	3.0.51	Get-StoragePolicies	2/16/2022, 10:08:51	2/16/2022, 10:09:51	✔ Succeeded
Get-CloudAdminGroups-Exec1	Microsoft.AVS.Management	3.0.51	Get-CloudAdminGroups	2/16/2022, 10:07:21	2/16/2022, 10:08:51	✔ Succeeded

Assign more vCenter Server Roles to Active Directory Identities

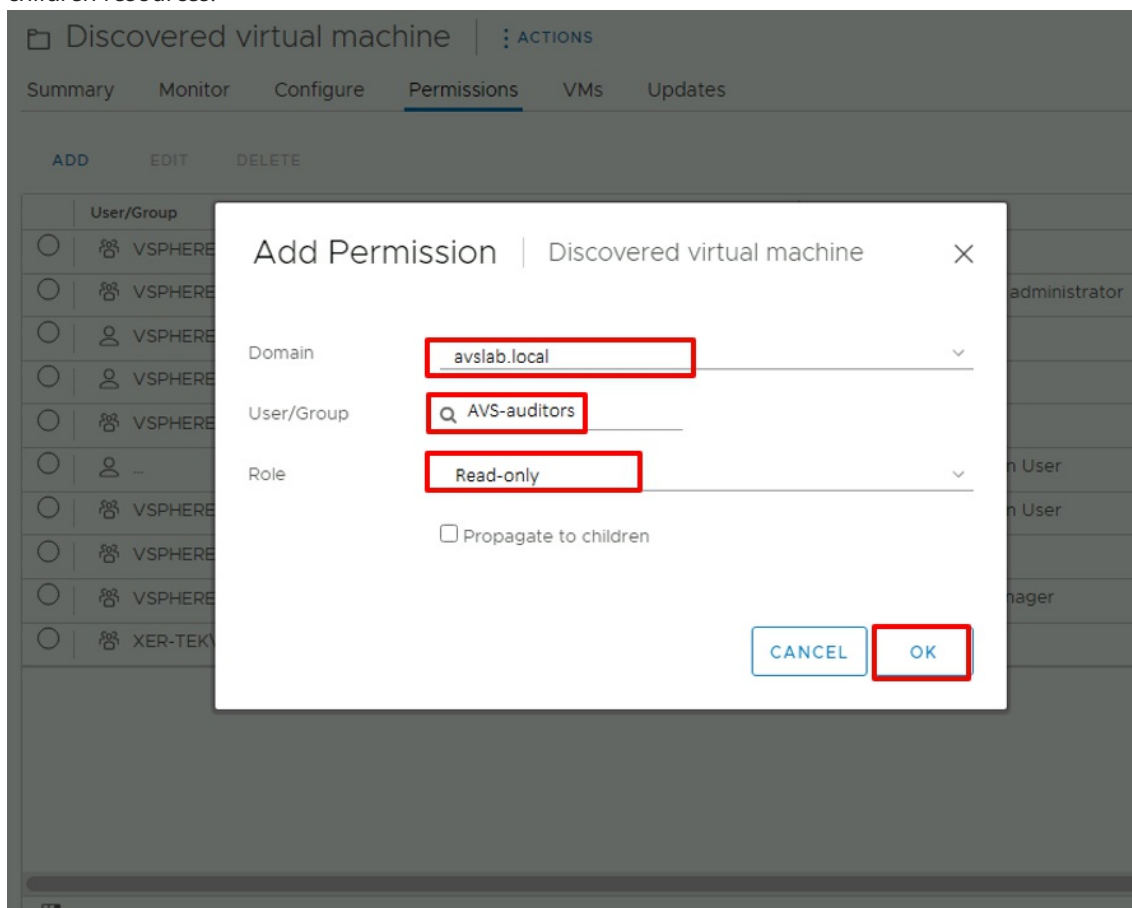
After you've added an external identity over LDAP or LDAPS, you can assign vCenter Server Roles to Active Directory security groups based on your organization's security controls.

1. After you sign in to vCenter Server with cloudadmin privileges, you can select an item from the inventory, select **ACTIONS** menu and select **Add Permission**.

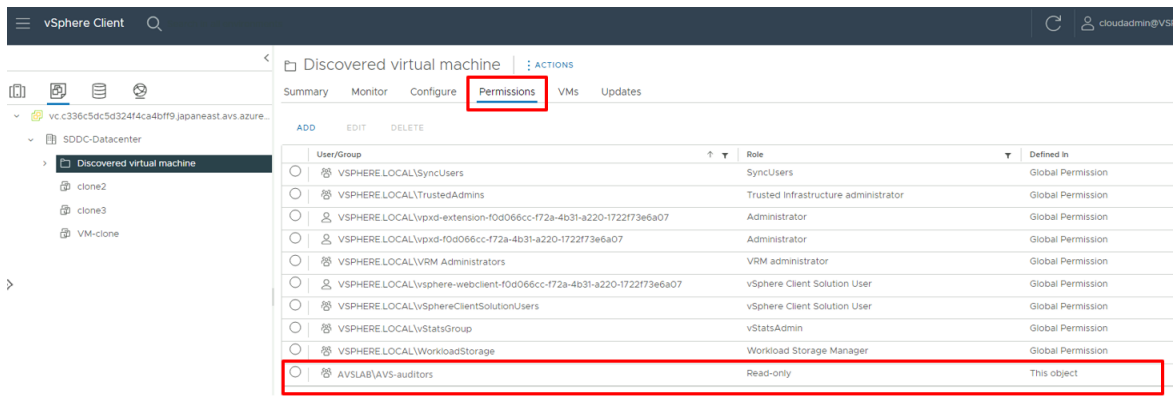


2. In the Add Permission prompt:

- a. *Domain*. Select the Active Directory that was added previously.
- b. *User/Group*. Enter the name of the desired user or group to find then select once is found.
- c. *Role*. Select the desired role to assign.
- d. *Propagate to children*. Optionally select the checkbox if permissions should be propagated down to children resources.



3. Switch to the **Permissions** tab and verify the permission assignment was added.



4. Users should now be able to sign in to vCenter Server using their Active Directory credentials.

Remove AD group from the cloudadmin role

You'll run the `Remove-GroupFromCloudAdmins` cmdlet to remove a specified AD group from the cloudadmin role.

1. Select **Run command > Packages > Remove-GroupFromCloudAdmins**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
GroupName	Name of the group to remove, for example, VcAdminGroup .
Retain up to	Retention period of the cmdlet output. The default value is 60 days.
Specify name for execution	Alphanumeric name, for example, removeADgroup .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** or the **Run Execution Status** pane to see the progress.

Remove existing external identity sources

You'll run the `Remove-ExternalIdentitySources` cmdlet to remove all existing external identity sources in bulk.

1. Select **Run command > Packages > Remove-ExternalIdentitySources**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
Retain up to	Retention period of the cmdlet output. The default value is 60 days.
Specify name for execution	Alphanumeric name, for example, remove_externalidentity .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** or the **Run Execution Status** pane to see the progress.

Next steps

Now that you've learned about how to configure LDAP and LDAPS, you can learn more about:

- [How to configure storage policy](#) - Each VM deployed to a vSAN datastore is assigned at least one VM storage policy. You can assign a VM storage policy in an initial deployment of a VM or when you do other VM operations, such as cloning or migrating.
- [Azure VMware Solution identity concepts](#) - Use vCenter Server to manage virtual machine (VM) workloads and NSX-T Manager to manage and extend the private cloud. Access and identity management use the cloudadmin role for vCenter Server and restricted administrator rights for NSX-T Manager.
- [Configure external identity source for NSX-T](#)
- [Azure VMware Solution identity concepts](#)
- [VMware product documentation](#)

Configure external identity source for NSX-T Data Center

12/16/2022 • 4 minutes to read • [Edit Online](#)

In this article, you'll learn how to configure an external identity source for NSX-T Data Center in an Azure VMware Solution. The NSX-T Data Center can be configured with external LDAP directory service to add remote directory users or groups. The users can be assigned an NSX-T Data Center Role-based access control (RBAC) role like you've on-premises.

Prerequisites

- A working connectivity from your Active Directory network to your Azure VMware Solution private cloud.
- If you require Active Directory authentication with LDAPS:
 - You'll need access to the Active Directory Domain Controller(s) with Administrator permissions.
 - Your Active Directory Domain Controller(s) must have LDAPS enabled with a valid certificate. The certificate could be issued by an [Active Directory Certificate Services Certificate Authority \(CA\)](#) or a [third-party CA](#).

NOTE

Self-sign certificates are not recommended for production environments.

- Ensure your Azure VMware Solution has DNS resolution configured to your on-premises AD. Enable DNS Forwarder from Azure portal. For more information, see [Configure NSX-T Data Center DNS for resolution to your Active Directory Domain and Configure DNS forwarder for Azure VMware Solution](#).

NOTE

For more information about LDAPS and certificate issuance, see with your security or identity management team.

Add Active Directory as LDAPS identity source

1. Sign-in to NSX-T Manager and Navigate to System > Users and Roles > LDAP.
2. Select on the Add Identity Source.
3. Enter a name for the identity source. For example, avslab.local.
4. Enter a domain name. The name must correspond to the domain name of your Active Directory server, if using Active Directory. For example, `avslab.local`.
5. Select the type as Active Directory over LDAP, if using Active Directory.
6. Enter the Base DN. Base DN is the starting point that an LDAP server uses when searching for user authentication within an Active Directory domain. For example: DC=avslab,DC=local.

NOTE

All of the user and group entries you intend to use to control access to NSX-T Data Center must be contained within the LDAP directory tree rooted at the specified Base DN. If the Base DN is set to something too specific, such as an Organizational Unit deeper in your LDAP tree, NSX may not be able to find the entries it needs to locate users and determine group membership. Selecting a broad Base DN is a best practice if you are unsure.

7. After filling in the required fields, you can select **Add** to configure LDAP servers. One LDAP server is supported for each domain.

FIELD	VALUE
Hostname/IP	The hostname or IP address of your LDAP server. For example, <code>dc.avs1ab.local</code> .
LDAP Protocol	Select LDAPS (LDAP is unsecured).
Port	The default port is populated based on the selected protocol 636 for LDAPS and 389 for LDAP. If your LDAP server is running on a non-standard port, you can edit this text box to give the port number.
Connection Status	After filling in the mandatory text boxes, including the LDAP server information, select Connection Status to test the connection.
Use StartTLS	If selected, the LDAPv3 StartTLS extension is used to upgrade the connection to use encryption. To determine if you should use this option, consult your LDAP server administrator. This option can only be used if LDAP protocol is selected.
Certificate	If you're using LDAPS or LDAP + StartTLS, this text box should contain the PEM-encoded X.509 certificate of the server. If you leave this text box blank and select the Check Status link, NSX connects to the LDAP server. NSX will then retrieve the LDAP server's certificate, and prompt you if you want to trust that certificate. If you've verified that the certificate is correct, select OK , and the certificate text box will be populated with the retrieved certificate.
Bind Identity	The format is <code>user@domainName</code> , or you can specify the distinguished name. For Active Directory, you can use either the userPrincipalName (user@domainName) or the distinguished name. For OpenLDAP, you must supply a distinguished name. This text box is required unless your LDAP server supports anonymous bind, then it's optional. Consult your LDAP server administrator if you aren't sure.
Password	Enter a password for the LDAP server. This text box is required unless your LDAP server supports anonymous bind, then it's optional. Consult your LDAP server administrator.

8. Select **Add**.

Set LDAP Server

Identity Source

#Ldap Servers 1

ADD LDAP SERVER

Maximum: 1

Hostname/IP	LDAP Protocol	Port	Connection Status
dc.avslab.local	LDAPS	636	Check Status
Use StartTLS	<input type="checkbox"/>	Certificate	Enter Certificate
Bind Identity	admin@avs.local	Password

Format: user@domainName or specify the distinguished Name

ADD CANCEL

CANCEL

APPLY

9. Select **Save** to complete the changes.

Name	Domain Name (FGDN)	Type	LDAP Servers	Connection Status
avslab.local	avslab.local	Active Directory over LDAP	1	

Base DN: dc=avslab,dc=local

Description: Example: CN=Users,DC=VMware,DC=com

Alternative Domain Names: Enter Alternative Domain Names

Example: vmware.com

SAVE CANCEL

Assign other NSX-T Data Center roles to Active Directory identities

After adding an external identity, you can assign NSX-T Data Center Roles to Active Directory security groups based on your organization's security controls.

1. Sign in to NSX-T Manager and navigate to **System > Users and Roles**.
2. Select **Add > Role Assignment for LDAP**.
 - a. Select a domain.
 - b. Enter the first few characters of the user's name, sign in ID, or a group name to search the LDAP directory, then select a user or group from the list that appears.
 - c. Select a role.
 - d. Select **Save**.

Users and Roles

User Role Assignment Local Users Roles LDAP VMware Identity Manager

ADD ▾

User/User Group Name	Roles	Type
Search Domain * avslab.local * avs-auditors * avslab.local	Select Roles *	LDAP User

Users and Roles

User Role Assignment Local Users Roles LDAP VMware Identity Manager

ADD ▾

User/User Group Name	Roles	Type
avslab.local * avs-auditors@avslab.local * SAVE CANCEL	Auditor X Select Roles	LDAP User

3. Verify the permission assignment is displayed under **Users and Roles**.

The screenshot shows the NSX-T Manager interface. The 'Users and Roles' page is active, displaying a table of users and their assigned roles. The user 'avs-auditors@avslab.local' is highlighted with a red box, showing they are assigned the 'Auditor' role and are of type 'LDAP User'. Other users listed include 'admin', 'adminsvc', 'audit', 'avs-admins@xer-tek.com', 'cloudadmin', and 'nsx-t-superuser'.

User/User Group Name	Roles	Type
avs-auditors@avslab.local	Auditor	LDAP User
admin	Enterprise Admin	Local User
adminsvc	Enterprise Admin	Local User
audit	Auditor	Local User
avs-admins@xer-tek.com	CloudAdmin	LDAP User
cloudadmin	CloudAdmin View More	Local User
nsx-t-superuser	Enterprise Admin	Principal Identity User

4. Users should now be able to sign in to NSX-T Manager using their Active Directory credentials.

Next steps

Now that you've configured the external source, you can also learn about:

- [Configure external identity source for vCenter Server](#)
- [Azure VMware Solution identity concepts](#)
- [VMware product documentation](#)

Enable Managed SNAT for Azure VMware Solution workloads

12/16/2022 • 2 minutes to read • [Edit Online](#)

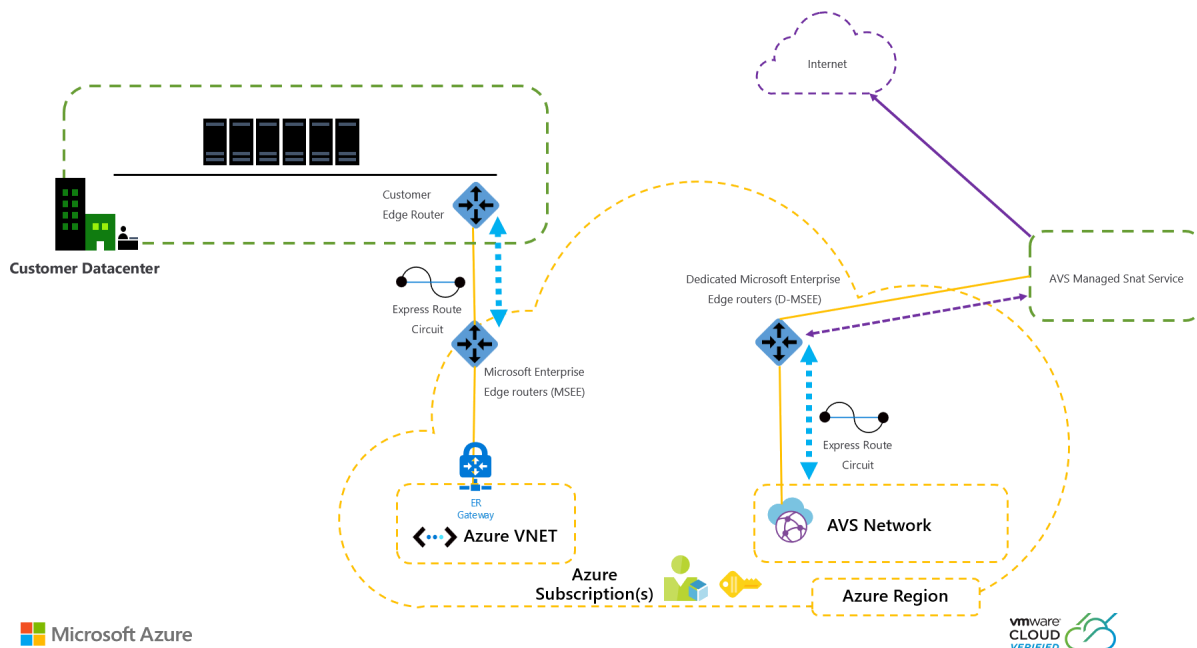
In this article, you'll learn how to enable Azure VMware Solution's Managed Source NAT (SNAT) to connect to the Internet outbound. A SNAT service translates from RFC1918 space to the public Internet for simple outbound Internet access. The SNAT service won't work when you have a default route from Azure.

With this capability, you:

- Have a basic SNAT service with outbound Internet connectivity from your Azure VMware Solution private cloud.
- Have no control of outbound SNAT rules.
- Are unable to view connection logs.
- Have a limit of 128 000 concurrent connections.

Reference architecture

The architecture shows Internet access outbound from your Azure VMware Solution private cloud using an Azure VMware Solution Managed SNAT Service.



Configure Outbound Internet access using Managed SNAT in the Azure port

1. Log in to the Azure portal and then search for and select **Azure VMware Solution**.
2. Select the Azure VMware Solution private cloud.
3. In the left navigation, under **Workload Networking**, select **Internet Connectivity**.
4. Select **Connect using SNAT** button and select **Save**. You have successfully enabled outbound Internet access for your Azure VMware Solution private cloud using our Managed SNAT service.

Next steps

[Internet connectivity design considerations \(Preview\)](#)

[Enable Public IP to the NSX Edge for Azure VMware Solution \(Preview\)](#)

[Disable Internet access or enable a default route](#)

Enable Public IP to the NSX-T Data Center Edge for Azure VMware Solution

12/16/2022 • 6 minutes to read • [Edit Online](#)

In this article, you'll learn how to enable Public IP to the NSX-T Data Center Edge for your Azure VMware Solution.

TIP

Before you enable Internet access to your Azure VMware Solution, review the [Internet connectivity design considerations](#).

Public IP to the NSX-T Data Center Edge is a feature in Azure VMware Solution that enables inbound and outbound internet access for your Azure VMware Solution environment.

IMPORTANT

The use of Public IPv4 addresses can be consumed directly in Azure VMware Solution and charged based on the Public IPv4 prefix shown on [Pricing - Virtual Machine IP Address Options](#).

The Public IP is configured in Azure VMware Solution through the Azure portal and the NSX-T Data Center interface within your Azure VMware Solution private cloud.

With this capability, you have the following features:

- A cohesive and simplified experience for reserving and using a Public IP down to the NSX Edge.
- The ability to receive up to 1000 or more Public IPs, enabling Internet access at scale.
- Inbound and outbound internet access for your workload VMs.
- DDoS Security protection against network traffic in and out of the Internet.
- HCX Migration support over the Public Internet.

IMPORTANT

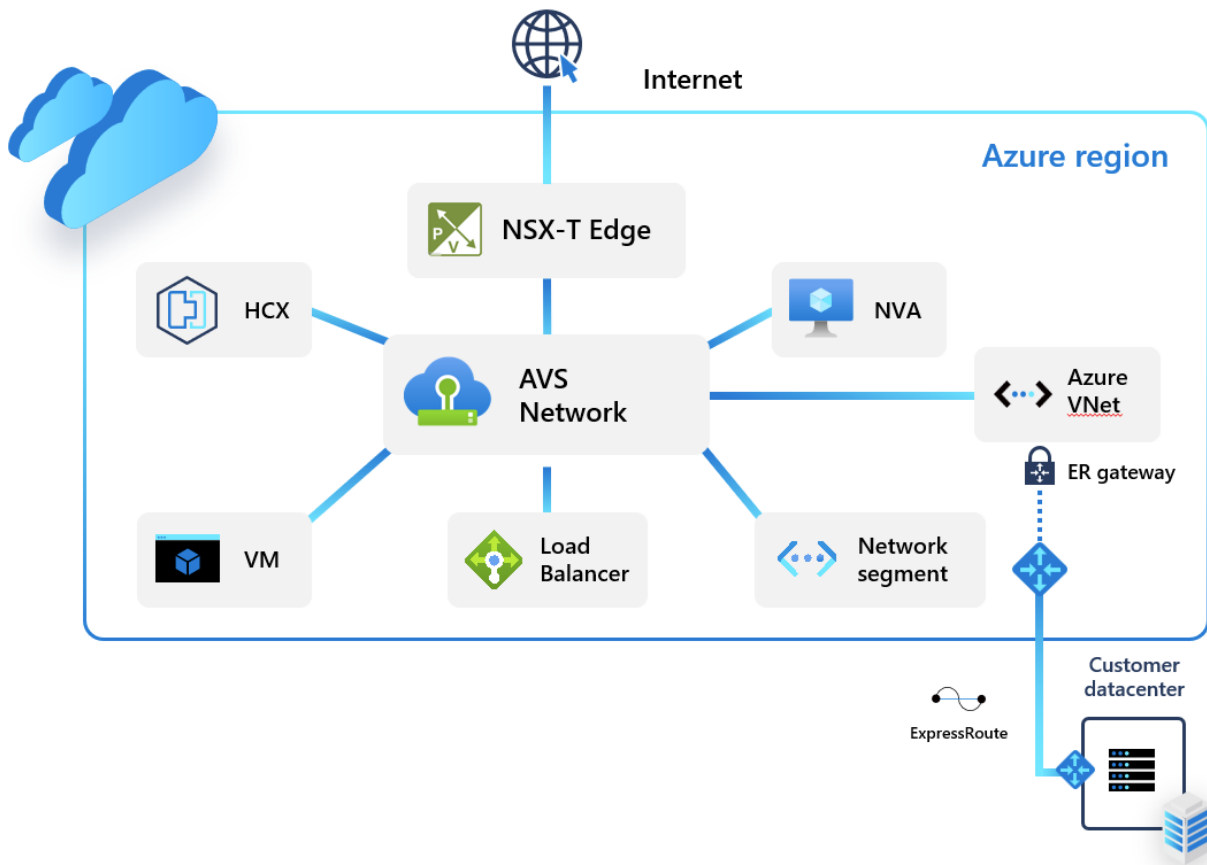
You can configure up to 64 total Public IP addresses across these network blocks. If you want to configure more than 64 Public IP addresses, please submit a support ticket stating how many.

Prerequisites

- Azure VMware Solution private cloud
- DNS Server configured on the NSX-T Data Center

Reference architecture

The architecture shows Internet access to and from your Azure VMware Solution private cloud using a Public IP directly to the NSX-T Data Center Edge.



IMPORTANT

The use of Public IP down to the NSX-T Data Center Edge is not compatible with reverse DNS Lookup.

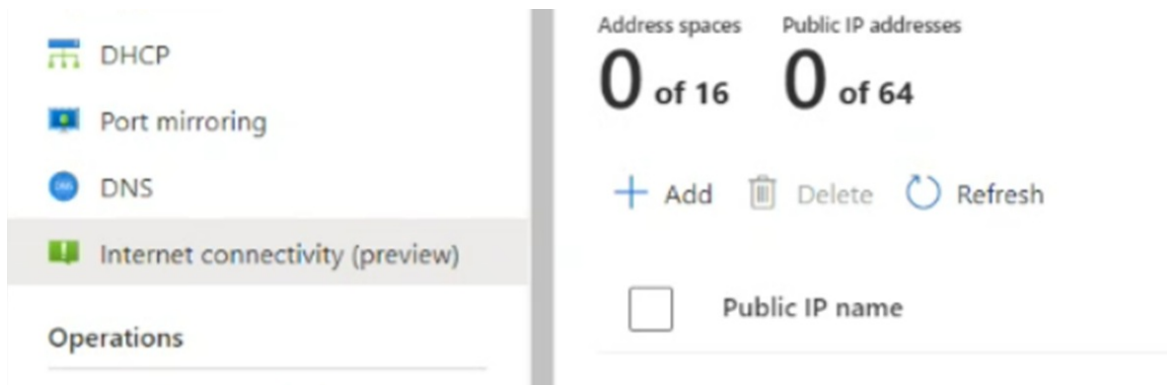
Configure a Public IP in the Azure portal

1. Log in to the Azure portal.
2. Search for and select Azure VMware Solution.
3. Select the Azure VMware Solution private cloud.
4. In the left navigation, under **Workload Networking**, select **Internet connectivity**.
5. Select the **Connect using Public IP down to the NSX-T Edge** button.

IMPORTANT

Before selecting a Public IP, ensure you understand the implications to your existing environment. For more information, see [Internet connectivity design considerations](#). This should include a risk mitigation review with your relevant networking and security governance and compliance teams.

6. Select **Public IP**.



7. Enter the **Public IP name** and select a subnet size from the **Address space** dropdown and select **Configure**.

8. This Public IP should be configured within 20 minutes and will show the subnet.



9. If you don't see the subnet, refresh the list. If the refresh fails, try the configuration again.

10. After configuring the Public IP, select the **Connect using the Public IP down to the NSX-T Edge** checkbox to disable all other Internet options.

11. Select **Save**.

You have successfully enabled Internet connectivity for your Azure VMware Solution private cloud and reserved a Microsoft allocated Public IP. You can now configure this Public IP down to the NSX-T Data Center Edge for your workloads. The NSX-T Data Center is used for all VM communication. There are several options for configuring your reserved Public IP down to the NSX-T Data Center Edge.

There are three options for configuring your reserved Public IP down to the NSX-T Data Center Edge: Outbound Internet Access for VMs, Inbound Internet Access for VMs, and Gateway Firewall used to Filter Traffic to VMs at T1 Gateways.

Outbound Internet access for VMs

A Sourced Network Translation Service (SNAT) with Port Address Translation (PAT) is used to allow many VMs to one SNAT service. This connection means you can provide Internet connectivity for many VMs.

IMPORTANT

To enable SNAT for your specified address ranges, you must [configure a gateway firewall rule](#) and SNAT for the specific address ranges you desire. If you don't want SNAT enabled for specific address ranges, you must create a [No-NAT rule](#) for the address ranges to exclude. For your SNAT service to work as expected, the No-NAT rule should be a lower priority than the SNAT rule.

Add rule

1. From your Azure VMware Solution private cloud, select **vCenter Server Credentials**
2. Locate your NSX-T Manager URL and credentials.
3. Log in to **VMware NSX-T Manager**.
4. Navigate to **NAT Rules**.
5. Select the T1 Router.
6. Select **ADD NAT RULE**.

Configure rule

1. Enter a name.
2. Select **SNAT**.
3. Optionally, enter a source such as a subnet to SNAT or destination.
4. Enter the translated IP. This IP is from the range of Public IPs you reserved from the Azure VMware Solution Portal.
5. Optionally, give the rule a higher priority number. This prioritization will move the rule further down the rule list to ensure more specific rules are matched first.
6. Click **SAVE**.

Logging can be enabled by way of the logging slider. For more information on NSX-T Data Center NAT configuration and options, see the [NSX-T Data Center NAT Administration Guide](#)

No Network Address Translation rule for specific address ranges

A No SNAT rule in NSX-T Manager can be used to exclude certain matches from performing Network Address Translation. This policy can be used to allow private IP traffic to bypass existing network translation rules.

1. From your Azure VMware Solution private cloud, select **vCenter Server Credentials**.
2. Locate your NSX-T Manager URL and credentials.
3. Log in to **VMware NSX-T Manager** and then select **NAT Rules**.
4. Select the T1 Router and then select **ADD NAT RULE**.
5. Select **NO SNAT** rule as the type of NAT rule.
6. Select the **Source IP** as the range of addresses you do not want to be translated. The **Destination IP** should be any internal addresses you are reaching from the range of Source IP ranges.
7. Select **SAVE**.

Inbound Internet Access for VMs

A Destination Network Translation Service (DNAT) is used to expose a VM on a specific Public IP address and/or a specific port. This service provides inbound internet access to your workload VMs.

Log in to VMware NSX-T Manager

1. From your Azure VMware Solution private cloud, select **VMware credentials**.
2. Locate your NSX-T Manager URL and credentials.
3. Log in to **VMware NSX-T Manager**.

Configure the DNAT rule

1. Name the rule.
2. Select **DNAT** as the action.
3. Enter the reserved Public IP in the destination match. This IP is from the range of Public IPs reserved from the Azure VMware Solution Portal.
4. Enter the VM Private IP in the translated IP.
5. Select **SAVE**.
6. Optionally, configure the Translated Port or source IP for more specific matches.

The VM is now exposed to the internet on the specific Public IP and/or specific ports.


Gateway Firewall used to filter traffic to VMs at T1 Gateways

You can provide security protection for your network traffic in and out of the public internet through your Gateway Firewall.

1. From your Azure VMware Solution Private Cloud, select **VMware credentials**.

2. Locate your NSX-T Manager URL and credentials.
3. Log in to **VMware NSX-T Manager**.
4. From the NSX-T home screen, select **Gateway Policies**.
5. Select **Gateway Specific Rules**, choose the T1 Gateway and select **ADD POLICY**.
6. Select **New Policy** and enter a policy name.
7. Select the Policy and select **ADD RULE**.
8. Configure the rule.
 - a. Select **New Rule**.
 - b. Enter a descriptive name.
 - c. Configure the source, destination, services, and action.
9. Select **Match External Address** to apply firewall rules to the external address of a NAT rule.

For example, the following rule is set to Match External Address, and this setting will allow SSH traffic inbound to the Public IP.



If **Match Internal Address** was specified, the destination would be the internal or private IP address of the VM.

For more information on the NSX-T Data Center Gateway Firewall see the [NSX-T Data Center Gateway Firewall Administration Guide](#). The Distributed Firewall could be used to filter traffic to VMs. This feature is outside the scope of this document. For more information, see [NSX-T Data Center Distributed Firewall Administration Guide](#).

Next steps

[Internet connectivity design considerations \(Preview\)](#)

[Enable Managed SNAT for Azure VMware Solution Workloads \(Preview\)](#)

[Disable Internet access or enable a default route](#)

[Enable HCX access over the internet](#)

Disable internet access or enable a default route

12/16/2022 • 2 minutes to read • [Edit Online](#)

In this article, you'll learn how to disable Internet access or enable a default route for your Azure VMware Solution private cloud. There are multiple ways to set up a default route. You can use a Virtual WAN hub, Network Virtual Appliance in a Virtual Network, or use a default route from on-premises. If you don't set up a default route, there will be no Internet access to your Azure VMware Solution private cloud.

With a default route setup, you can achieve the following tasks:

- Disable Internet access to your Azure VMware Solution private cloud.

NOTE

Ensure that a default route is not advertised from on-premises or Azure as that will override this setup.

- Enable Internet access by generating a default route from Azure Firewall or third-party Network Virtual Appliance.

Prerequisites

- If Internet access is required, a default route must be advertised from an Azure Firewall, Network Virtual Appliance or Virtual WAN Hub.
- Azure VMware Solution private cloud.

Disable Internet access or enable a default route in the Azure portal

1. Log in to the Azure portal.
2. Search for **Azure VMware Solution** and select it.
3. Locate and select your Azure VMware Solution private cloud.
4. On the left navigation, under **Workload networking**, select **Internet connectivity**.
5. Select the **Don't connect or connect using default route from Azure** button and select **Save**.
If you don't have a default route from on-premises or from Azure, you have successfully disabled Internet connectivity to your Azure VMware Solution private cloud.

Next steps

[Internet connectivity design considerations \(Preview\)](#)

[Enable Managed SNAT for Azure VMware Solution Workloads](#)

[Enable Public IP to the NSX Edge for Azure VMware Solution](#)

Configure DHCP for Azure VMware Solution

12/16/2022 • 5 minutes to read • [Edit Online](#)

Applications and workloads running in a private cloud environment require name resolution and DHCP services for lookup and IP address assignments. A proper DHCP and DNS infrastructure are required to provide these services. You can configure a virtual machine to provide these services in your private cloud environment.

Use the DHCP service built-in to NSX or use a local DHCP server in the private cloud instead of routing broadcast DHCP traffic over the WAN back to on-premises.

In this how-to article, you'll use NSX-T Manager to configure DHCP for Azure VMware Solution in one of the following ways:

- [Use the Azure portal to create a DHCP server or relay](#)
- [Use NSX-T Data Center to host your DHCP server](#)
- [Use a third-party external DHCP server](#)

TIP

If you want to configure DHCP using a simplified view of NSX-T Data Center operations, see [Configure DHCP for Azure VMware Solution](#).

IMPORTANT

For clouds created on or after July 1, 2021, the simplified view of NSX-T Data Center operations must be used to configure DHCP on the default Tier-1 Gateway in your environment.

DHCP does not work for virtual machines (VMs) on the VMware HCX L2 stretch network when the DHCP server is in the on-premises datacenter. NSX-T Data Center, by default, blocks all DHCP requests from traversing the L2 stretch. For the solution, see the [Configure DHCP on L2 stretched VMware HCX networks](#) procedure.

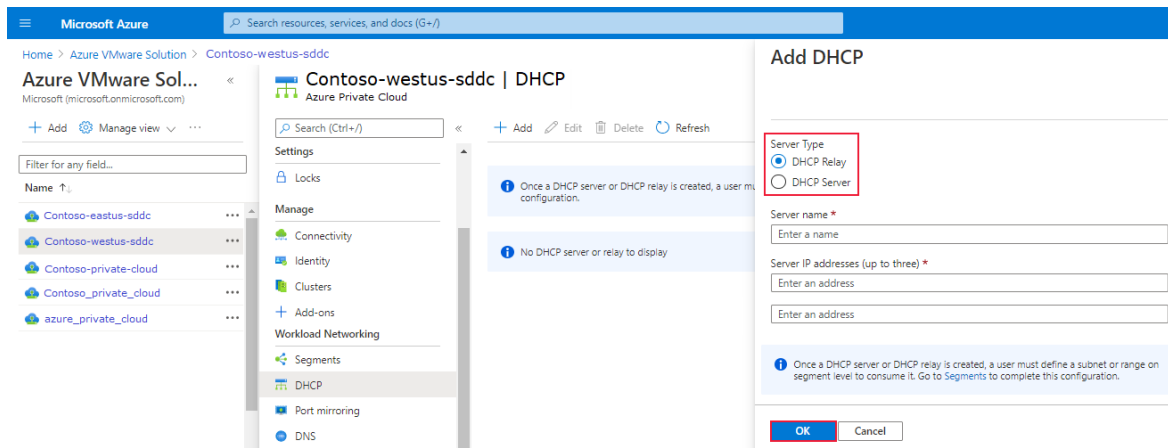
Use the Azure portal to create a DHCP server or relay

You can create a DHCP server or relay directly from Azure VMware Solution in the Azure portal. The DHCP server or relay connects to the Tier-1 gateway created when you deployed Azure VMware Solution. All the segments where you gave DHCP ranges will be part of this DHCP. After you've created a DHCP server or DHCP relay, you must define a subnet or range on segment level to consume it.

1. In your Azure VMware Solution private cloud, under **Workload Networking**, select **DHCP > Add**.
2. Select either **DHCP Server** or **DHCP Relay** and then provide a name for the server or relay and three IP addresses.

NOTE

For DHCP relay, you only require one IP address for a successful configuration.



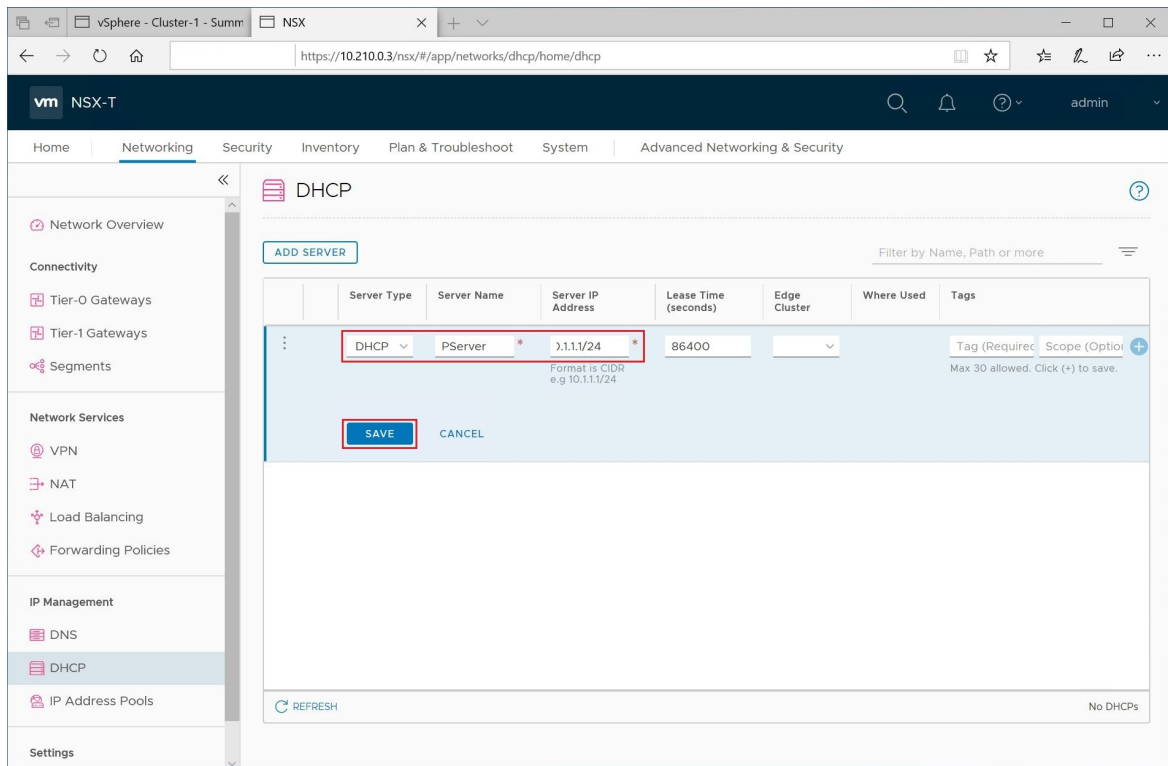
3. Complete the DHCP configuration by [providing DHCP ranges on the logical segments](#) and then select **OK**.

Use NSX-T Data Center to host your DHCP server

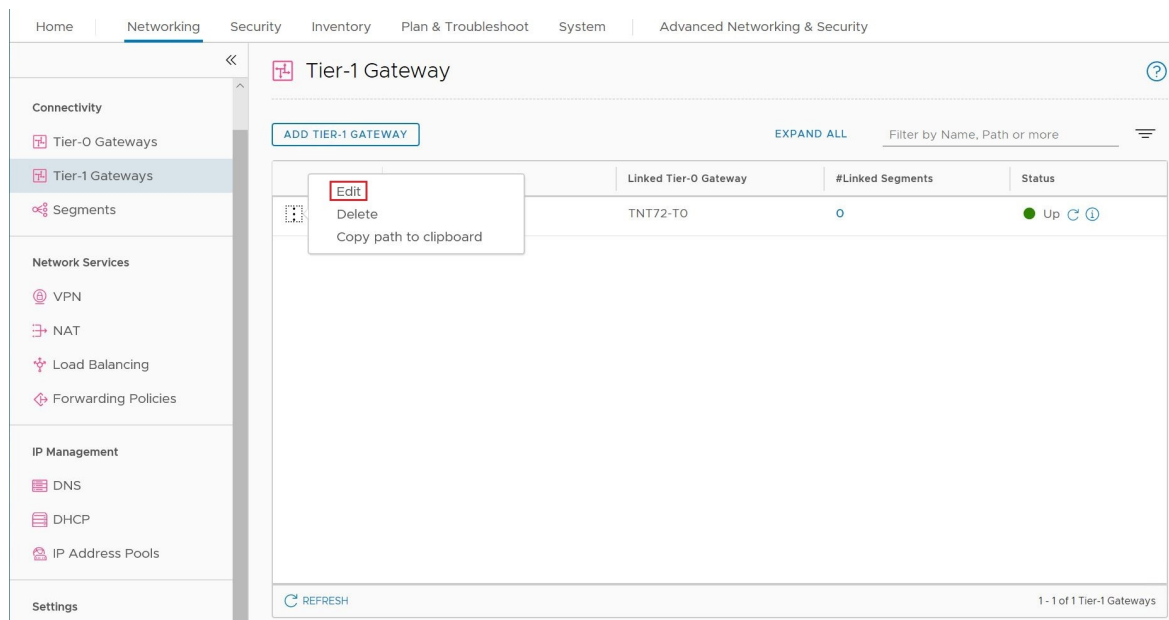
If you want to use NSX-T Data Center to host your DHCP server, you'll create a DHCP server and a relay service. Then you'll add a network segment and specify the DHCP IP address range.

Create a DHCP server


1. In NSX-T Manager, select **Networking** > **DHCP**, and then select **Add Server**.
2. Select **DHCP** for the **Server Type**, provide the server name and IP address, and select **Save**.



3. Select **Tier 1 Gateways**, select the vertical ellipsis on the Tier-1 gateway, and then select **Edit**.



4. Select **No IP Allocation Set** to add a subnet.



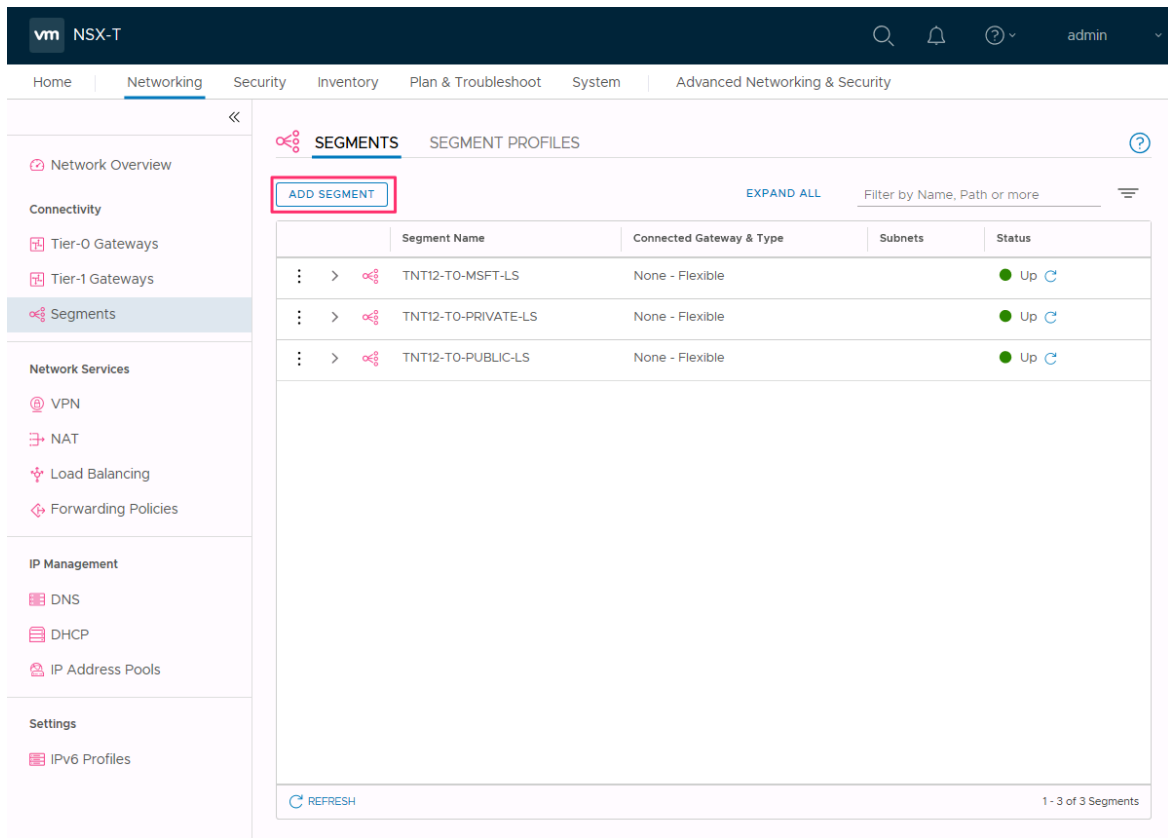
5. For **Type**, select **DHCP Local Server**.

6. For the **DHCP Server**, select **Default DHCP**, and then select **Save**.

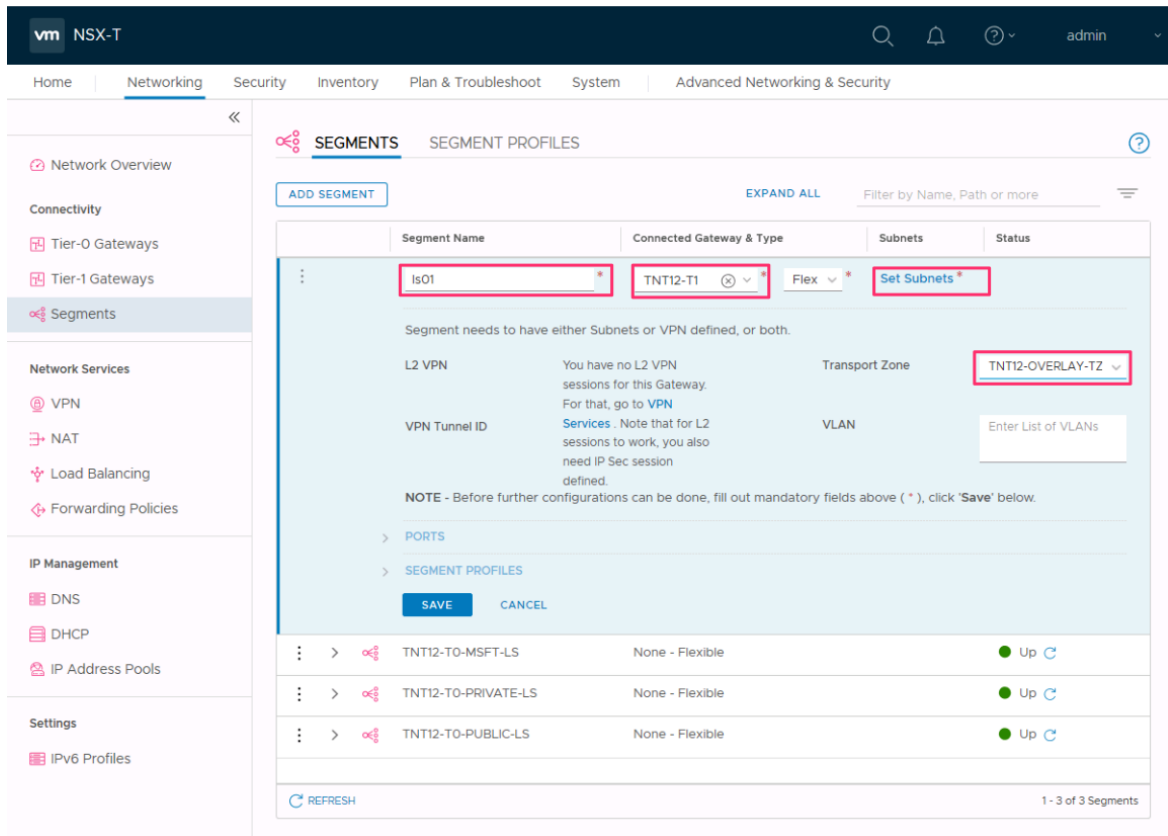
7. Select **Save** again and then select **Close Editing**.

Add a network segment

1. In NSX-T Manager, select **Networking > Segments**, and then select **Add Segment**.



2. Enter a name for the segment.
3. Select the Tier-1 Gateway (TNT \bar{x} -T1) as the **Connected Gateway** and leave the **Type** as Flexible.
4. Select the pre-configured overlay **Transport Zone** (TNT \bar{x} -OVERLAY-TZ) and then select **Set Subnets**.



5. Enter the gateway IP address and then select **Add**.

IMPORTANT

The IP address needs to be on a non-overlapping RFC1918 address block, which ensures connection to the VMs on the new segment.

Set Subnets

Segment #Subnets 1

ADD SUBNET Search

Gateway IP/Prefix Length	DHCP Ranges
10.10.230.1/24 Format CIDR e.g. 10.12.2.1/24	Enter DHCP Ranges Formats, e.g. 10.12.2.64/26, 10.12.2.2-10.12.2.50

ADD CANCEL

CANCEL APPLY

6. Select **Apply** and then **Save**.

7. Select **No** to decline the option to continue configuring the segment.

Specify the DHCP IP address range

When you create a relay to a DHCP server, you'll also specify the DHCP IP address range.

NOTE

The IP address range shouldn't overlap with the IP range used in other virtual networks in your subscription and on-premises networks.

1. In NSX-T Manager, select **Networking** > **Segments**.
2. Select the vertical ellipsis on the segment name and select **Edit**.
3. Select **Set Subnets** to specify the DHCP IP address for the subnet.

SEGMENTS SEGMENT PROFILES

ADD SEGMENT EXPAND ALL Filter by Name, Path or more

Segment Name	Connected Gateway & Type	Subnets	Status
Default VM Segment *	TNT72-T1 X ⊗ * Flexit *	Set Subnets *	
Segment needs to have either Subnets or VPN defined, or both.			
L2 VPN	You have no L2 VPN sessions for this Gateway. For that, go to VPN Services . Note that for L2 sessions to work, you also need IP Sec session defined.	Transport Zone	TNT72-OVERLAY-TZ Ov ▾
VPN Tunnel ID		VLAN	Enter List of VLANs
NOTE - Before further configurations can be done, fill out mandatory fields above (*), click 'Save' below.			
PORTS			
SEGMENT PROFILES			
<input type="button" value="SAVE"/> <input type="button" value="CANCEL"/>			
TNT72-TO-MSFT-LS	None - Flexible		Up ↻

REFRESH 1 - 3 of 3 Segments

4. Modify the gateway IP address if needed, and enter the DHCP range IP.

Set Subnets

Segment #Subnets 1

ADD SUBNET Search

Gateway	DHCP Ranges
<input type="text" value="10.12.2.1/24"/> * Format CIDR e.g. 10.12.2.1/24	<input type="text" value="10.12.2.64/26"/> X Enter DHCP Ranges Formats, e.g. 10.12.2.64/26, 10.12.2.2-10.12.2.50
<input type="button" value="ADD"/> <input type="button" value="CANCEL"/>	

5. Select **Apply**, and then **Save**. The segment is assigned a DHCP server pool.

Segment Name	Connected Gateway & Type	Subnets	Status
Default VM Segment	TNT72-T1 Tier1 - Flexible	1	Up
L2 VPN	Transport Zone	TNT72-OVERLAY-TZ Overlay	VIEW STATISTICS VIEW RELATED GROUPS
VPN Tunnel ID	VLAN		
Domain Name	IP Address Pool		
Tags	0		
PORTS SEGMENT PROFILES ADVANCED CONFIGURATION			
TNT72-T0-MSFT-LS	None - Flexible		Up
TNT72-T0-PRIVATE-LS	None - Flexible		Up
TNT72-T0-PUBLIC-LS	None - Flexible		Up

REFRESH 1 - 4 of 4 Segments

Use a third-party external DHCP server

If you want to use a third-party external DHCP server, you'll create a DHCP relay service in NSX-T Manager. You'll also specify the DHCP IP address range.

IMPORTANT

For clouds created on or after July 1, 2021, the simplified view of NSX-T Data Center operations must be used to configure DHCP on the default Tier-1 Gateway in your environment.

Create DHCP relay service

Use a DHCP relay for any non-NSX-based DHCP service. For example, a VM running DHCP in Azure VMware Solution, Azure IaaS, or on-premises.

1. In NSX-T Manager, select **Networking > DHCP**, and then select **Add Server**.
2. Select **DHCP Relay** for the **Server Type**, provide the server name and IP address, and select **Save**.

Home | Networking | Security | Inventory | Plan & Troubleshoot | System | Advanced Networking & Security

Connectivity

- Tier-0 Gateways
- Tier-1 Gateways
- Segments

Network Services

- VPN
- NAT
- Load Balancing
- Forwarding Policies

IP Management

- DNS
- DHCP

DHCP


ADD SERVER

Filter by Name, Path or more

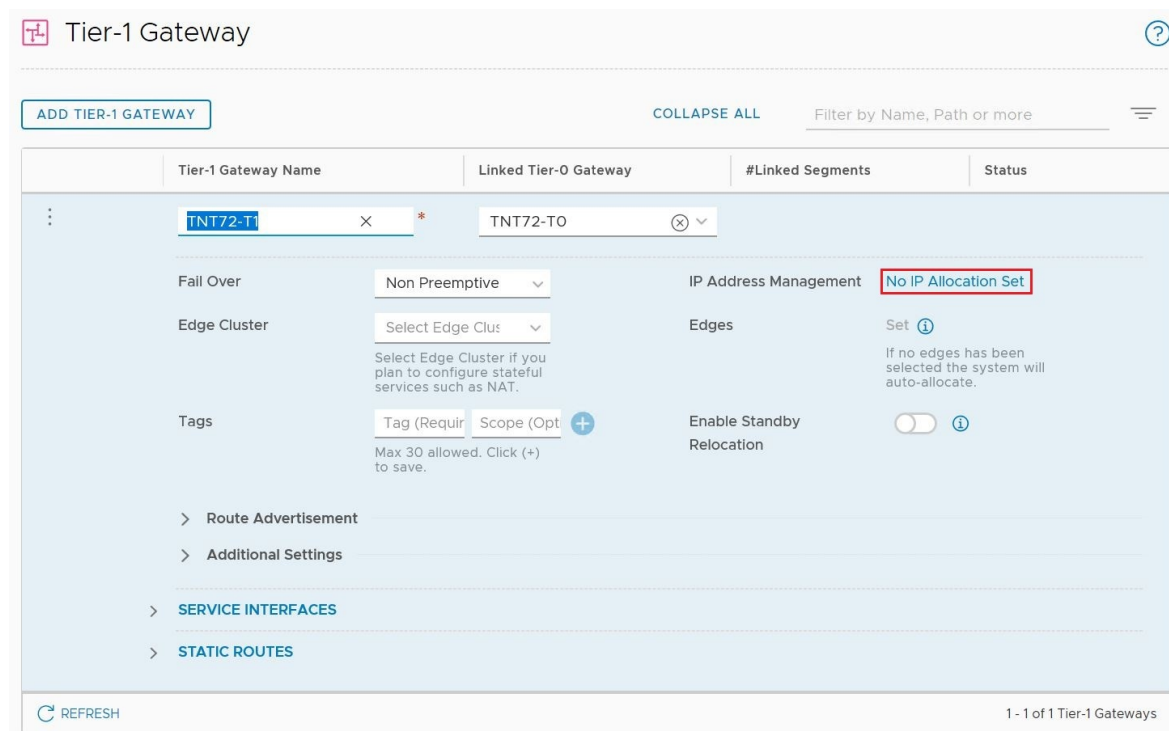
Server Type	Server Name	Server IP Address	Lease Time (seconds)	Edge Cluster	Where Used	Tags
DHCP Rel.	HCP-Relay *	10.10.10.10 1 or More IP Adc e.g. 10.10.10.10				Tag Scope Max 30 allowed. Click (+) to save.

SAVE CANCEL

3. Select **Tier 1 Gateways**, select the vertical ellipsis on the Tier-1 gateway, and then select **Edit**.



4. Select **No IP Allocation Set** to define the IP address allocation.



5. For **Type**, select **DHCP Server**.

6. For the **DHCP Server**, select **DHCP Relay**, and then select **Save**.

7. Select **Save** again and then select **Close Editing**.

Specify the DHCP IP address range

When you create a relay to a DHCP server, you'll also specify the DHCP IP address range.

NOTE

The IP address range shouldn't overlap with the IP range used in other virtual networks in your subscription and on-premises networks.

1. In NSX-T Manager, select **Networking** > **Segments**.

2. Select the vertical ellipsis on the segment name and select **Edit**.

3. Select **Set Subnets** to specify the DHCP IP address for the subnet.

The screenshot shows the 'SEGMENTS' configuration page. At the top, there are tabs for 'SEGMENTS' and 'SEGMENT PROFILES'. Below the tabs, there is an 'ADD SEGMENT' button and an 'EXPAND ALL' button. A search filter is present: 'Filter by Name, Path or more'. The main content area is a table with columns: 'Segment Name', 'Connected Gateway & Type', 'Subnets', and 'Status'. The first row shows 'Default VM Segment' with a red asterisk, 'TNT72-T|' with a red asterisk, 'Flexit' with a red asterisk, and 'Set Subnets' with a red asterisk and a red box around it. Below the table, there is a message: 'Segment needs to have either Subnets or VPN defined, or both.' followed by 'L2 VPN' and 'VPN Tunnel ID' sections. A 'NOTE' is present: 'NOTE - Before further configurations can be done, fill out mandatory fields above (*), click 'Save' below.' There are expandable sections for 'PORTS' and 'SEGMENT PROFILES'. At the bottom, there are 'SAVE' and 'CANCEL' buttons. Below the main content, there is a summary row for 'TNT72-T0-MSFT-LS' with 'None - Flexible' and a status indicator 'Up'. A 'REFRESH' button and '1 - 3 of 3 Segments' are at the bottom.

4. Modify the gateway IP address if needed, and enter the DHCP range IP.

The screenshot shows the 'Set Subnets' configuration page. At the top, there is a title 'Set Subnets' and a 'Segment' dropdown menu with '#Subnets 1'. Below the title, there is an 'ADD SUBNET' button and a search box. The main content area is a form with two sections: 'Gateway' and 'DHCP Ranges'. The 'Gateway' field is set to '10.12.2.1/24' with a red asterisk and a hint 'Format CIDR e.g. 10.12.2.1/24'. The 'DHCP Ranges' field is set to '10.12.2.64/26' with a red asterisk and a hint 'Enter DHCP Ranges'. Below the 'DHCP Ranges' field, there is a hint: 'Formats, e.g. 10.12.2.64/26, 10.12.2.2-10.12.2.50'. At the bottom, there are 'ADD' and 'CANCEL' buttons. Below the main content, there are 'CANCEL' and 'APPLY' buttons.

5. Select **Apply**, and then **Save**. The segment is assigned a DHCP server pool.

SEGMENTS		SEGMENT PROFILES			
ADD SEGMENT		EXPAND ALL		Filter by Name, Path or more	
Segment Name	Connected Gateway & Type	Subnets	Status		
⋮ <input checked="" type="checkbox"/>	Default VM Segment	TNT72-T1 Tier1 - Flexible	1	● Up ↻	
L2 VPN		Transport Zone	TNT72-OVERLAY-TZ Overlay	VIEW STATISTICS	
VPN Tunnel ID		VLAN		VIEW RELATED GROUPS	
Domain Name		IP Address Pool			
Tags		0			
> PORTS > SEGMENT PROFILES ADVANCED CONFIGURATION					
⋮ >	TNT72-T0-MSFT-LS	None - Flexible		● Up ↻	
⋮ >	TNT72-T0-PRIVATE-LS	None - Flexible		● Up ↻	
⋮ >	TNT72-T0-PUBLIC-LS	None - Flexible		● Up ↻	
REFRESH		1 - 4 of 4 Segments			

Next steps

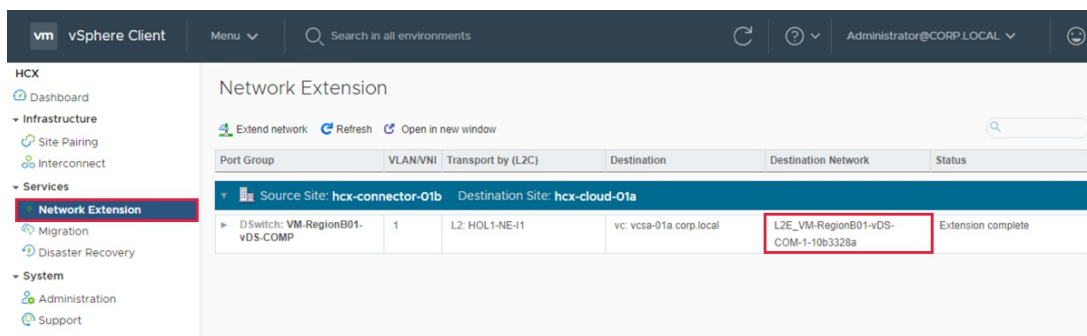
If you want to send DHCP requests from your Azure VMware Solution VMs to a non-NSX-T DHCP server, see the [Configure DHCP on L2 stretched VMware HCX networks](#) procedure.

Configure DHCP on L2 stretched VMware HCX networks

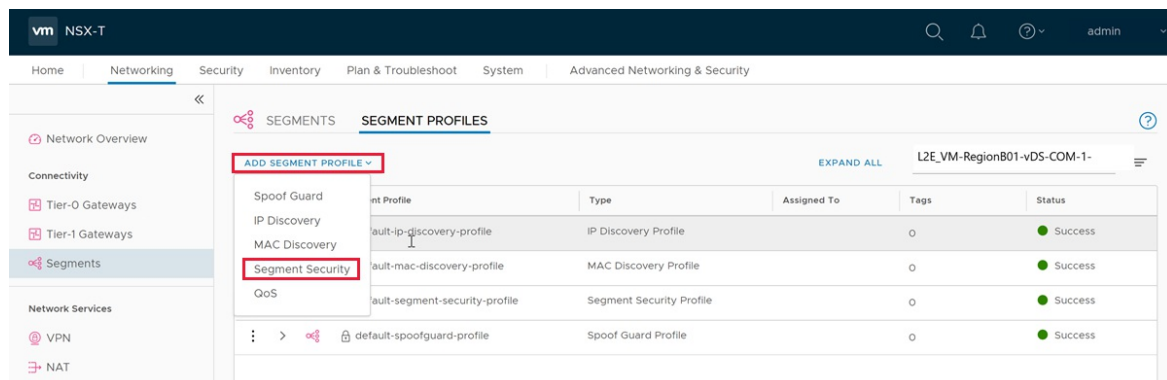
12/16/2022 • 2 minutes to read • [Edit Online](#)

DHCP does not work for virtual machines (VMs) on the VMware HCX L2 stretch network when the DHCP server is in the on-premises data center. This is because NSX-T Data Center, by default, blocks all DHCP requests from traversing the L2 stretch. Therefore, to send DHCP requests from your Azure VMware Solution VMs to a non-NSX-T Data Center DHCP server, you'll need to configure DHCP on L2 stretched VMware HCX networks.

1. (Optional) If you need to locate the segment name of the L2 extension:
 - a. Sign in to your on-premises vCenter Server, and under **Home**, select **HCX**.
 - b. Select **Network Extension** under **Services**.
 - c. Select the network extension you want to support DHCP requests from Azure VMware Solution to on-premises.
 - d. Take note of the destination network name.



2. In NSX-T Manager, select **Networking** > **Segments** > **Segment Profiles**.
3. Select **Add Segment Profile** and then **Segment Security**.



4. Provide a name and a tag, and then set the **BPDU Filter** toggle to ON and all the DHCP toggles to OFF.

SEGMENTS **SEGMENT PROFILES** ?

ADD SEGMENT PROFILE ▾ EXPAND ALL Filter by Name, Path or more

Segment Profile	Type	Assigned To	Tags	Status
SecProfile-DHCP	Segment Security Profile		dhcp	Scope (Optional) +

Max 30 allowed. Click (+) to save.

BPDU

BPDU Filter

BPDU Filter Allow List

- 01:80:c2: >
- 01:80:c2: >
- 01:80:c2: >
- 01:80:c2: >
- 01:80:c2: >
- 01:80:c2: >

DHCP

Server Block	<input type="checkbox"/>	Client Block	<input type="checkbox"/>
Server Block - IPv6	<input type="checkbox"/>	Client Block - IPv6	<input type="checkbox"/>
Non-IP Traffic Block	<input type="checkbox"/>	RA Guard	<input checked="" type="checkbox"/>
Rate Limits	<input type="checkbox"/>		

Receive Broadcast	<input type="text" value="0"/>	Transmit Broadcast	<input type="text" value="0"/>
Receive Multicast	<input type="text" value="0"/>	Transmit Multicast	<input type="text" value="0"/>

SAVE CANCEL

SEGMENTS **SEGMENT PROFILES** ?

ADD SEGMENT EXPAND ALL Filter by Name, Path or more

Segment Name	Connected Gateway & Type	Subnets	Status
Connectivity	<input checked="" type="checkbox"/> ⓘ		

Tags +
Max 30 allowed. Click (+) to save.

> PORTS

< SEGMENT PROFILES

IP Discovery	default-ip-discovery-profile	Spoof Guard	default-spoofguard-profile
MAC Discovery	default-mac-discovery-profile	Segment Security	default-segment-security-profile
QoS	None		SecProfile-DHCP

ADVANCED CONFIGURATION

CLOSE EDITING

Configure a DNS forwarder in the Azure portal

12/16/2022 • 4 minutes to read • [Edit Online](#)

IMPORTANT

For Azure VMware Solution private clouds created on or after July 1, 2021, you now have the ability to configure private DNS resolution. For private clouds created before July 1, 2021, that need private DNS resolution, open a [support request](#) and request Private DNS configuration.

By default, Azure VMware Solution management components such as vCenter Server can only resolve name records available through Public DNS. However, certain hybrid use cases require Azure VMware Solution management components to resolve name records from privately hosted DNS to properly function, including customer-managed systems such as vCenter Server and Active Directory.

Private DNS for Azure VMware Solution management components lets you define conditional forwarding rules for the desired domain name to a selected set of private DNS servers through the NSX-T Data Center DNS Service.

This capability uses the DNS Forwarder Service in NSX-T Data Center. A DNS service and default DNS zone are provided as part of your private cloud. To enable Azure VMware Solution management components to resolve records from your private DNS systems, you must define an FQDN zone and apply it to the NSX-T Data Center DNS Service. The DNS Service conditionally forwards DNS queries for each zone based on the external DNS servers defined in that zone.

NOTE

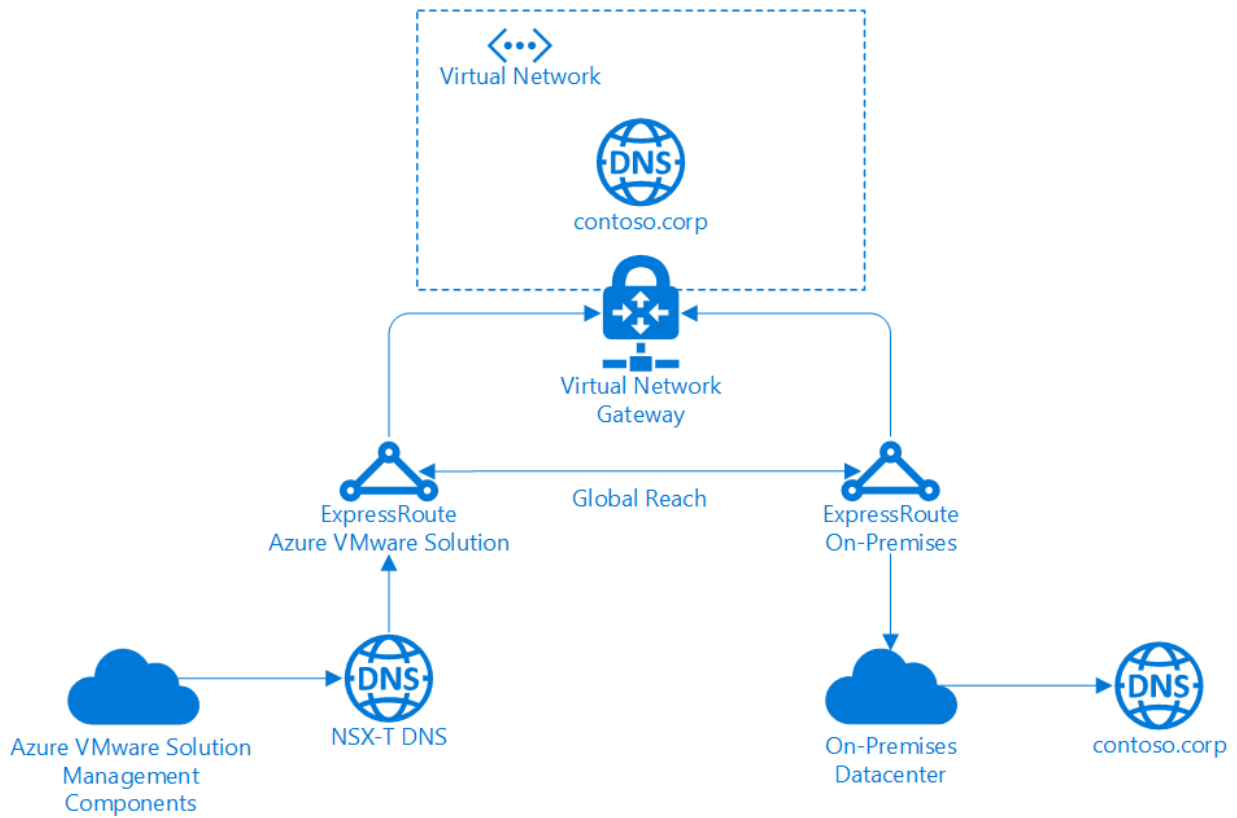
The DNS Service is associated with up to five FQDN zones. Each FQDN zone is associated with up to three DNS servers.

TIP

If desired, you can also use the conditional forwarding rules for workload segments by configuring virtual machines on those segments to use the NSX-T Data Center DNS Service IP address as their DNS server.

Architecture

The diagram shows that the NSX-T Data Center DNS Service can forward DNS queries to DNS systems hosted in Azure and on-premises environments.



Configure DNS forwarder

1. In your Azure VMware Solution private cloud, under **Workload Networking**, select **DNS > DNS zones**. Then select **Add**.

NOTE

For private clouds created on or after July 1, 2021, the default DNS zone is created for you during the private cloud creation.

Home > azcat-avs-cac-js-pc02

azcat-avs-cac-js-pc02 | DNS

Search (Ctrl+/) Add Edit Delete Refresh

DNS zones DNS service

Filter by name Name: All

Name	Domain	DNS servers	Source IP	DNS service
TNT74-DNS-FORWARDER-ZONE	any	1.1.1.1,1.0.0.1	-	1

Operations: Run command, Monitoring: Alerts, Metrics, Advisor recommendations, Automation

2. Select FQDN zone, provide a name and up to three DNS server IP addresses in the format of 10.0.0.53. Then select OK.

Add DNS zone

Type

Default DNS zone

FQDN zone

DNS zone name *
contoso ✓

Domain *
contoso.corp ✓

DNS server IP (up to 3) *

172.21.88.20 ✓

172.21.88.34 ✓

Enter an address ✓

Source IP
Enter an address ✓

OK Cancel

IMPORTANT

While NSX-T Data Center allows spaces and other non-alphanumeric characters in a DNS zone name, certain NSX-T Data Center resources such as a DNS Zone are mapped to an Azure resource whose names don't permit certain characters.

As a result, DNS zone names that would otherwise be valid in NSX-T Data Center may need adjustment to adhere to the [Azure resource naming conventions](#).

It takes several minutes to complete, and you can follow the progress from **Notifications**. You'll see a message in the Notifications when the DNS zone has been created.

3. Ignore the message about a default DNS zone because one gets created for you as part of your private cloud.
4. Select the **DNS service** tab and then select **Edit**.

TIP

For private clouds created on or after July 1, 2021, you can ignore the message about a default DNS zone as one is created for you during private cloud creation.

IMPORTANT

While certain operations in your private cloud may be performed from NSX-T Manager, for private clouds created on or after July 1, 2021, you *must* edit the DNS service from the Simplified Networking experience in the Azure portal for any configuration changes made to the default Tier-1 Gateway.

The screenshot shows the Azure portal interface for a private cloud named 'azcat-avs-cac-js-pc02'. The 'DNS' service is selected, and the 'DNS service' tab is active. The 'Edit' button is highlighted with a red box. The configuration details are as follows:

Property	Value
Name	TNT74-DNS-FORWARDER
Tier-1 Gateway	TNT74-T1
DNS service IP	10.114.240.192
Default DNS zone	TNT74-DNS-FORWARDER-ZONE
FQDN zones	-
Log level	INFO
Status	SUCCESS

5. From the **FQDN zones** drop-down, select the newly created FQDN, and then select **OK**.

Edit DNS service



Name *

TNT74-DNS-FORWARDER

Tier-1 Gateway

TNT74-T1

DNS service IP *

10.114.240.192

Default DNS zone *

Select DNS zone

FQDN zones (up to 5)

contoso

contoso

Info

OK

Cancel

It takes several minutes to complete, and once finished, you'll see the *Completed* message from **Notifications**. At this point, management components in your private cloud should be able to resolve DNS entries from the FQDN zone provided to the NSX-T Data Center DNS Service.

6. Repeat the above steps for other FQDN zones, including any applicable reverse lookup zones.

Verify name resolution operations

After you've configured the DNS forwarder, you'll have a few options available to verify name resolution operations.

NSX-T Manager

NSX-T Manager provides the DNS Forwarder Service statistics at the global service level and on a per-zone basis.

1. In NSX-T Manager, select **Networking** > **DNS**, and then expand your DNS Forwarder Service.

The screenshot shows the NSX-T Manager interface. The top navigation bar includes 'vm NSX-T' and user 'admin'. The main navigation menu on the left is expanded to 'DNS'. The main content area shows 'DNS SERVICES' with a table of services. The table has columns for Name, Tier0/Tier1 Gateway, DNS Service IP, Default DNS Zone, and Status. One service is listed: 'TNT86-DNS-FORWARDER' with gateway 'TNT86-T1', IP '10.103.64.192', and status 'Up'. Below the table, details for the selected service are shown, including 'Description: Not Set', 'Admin Status: Enabled', 'FQDN Zones: contoso-onprem', and 'Log Level: Info'. A 'VIEW STATISTICS' link is visible.

Name	Tier0/Tier1 Gateway	DNS Service IP	Default DNS Zone	Status
TNT86-DNS-FORWARDER	TNT86-T1	10.103.64.192	TNT86-DNS-FORWARDER-ZONE	Up

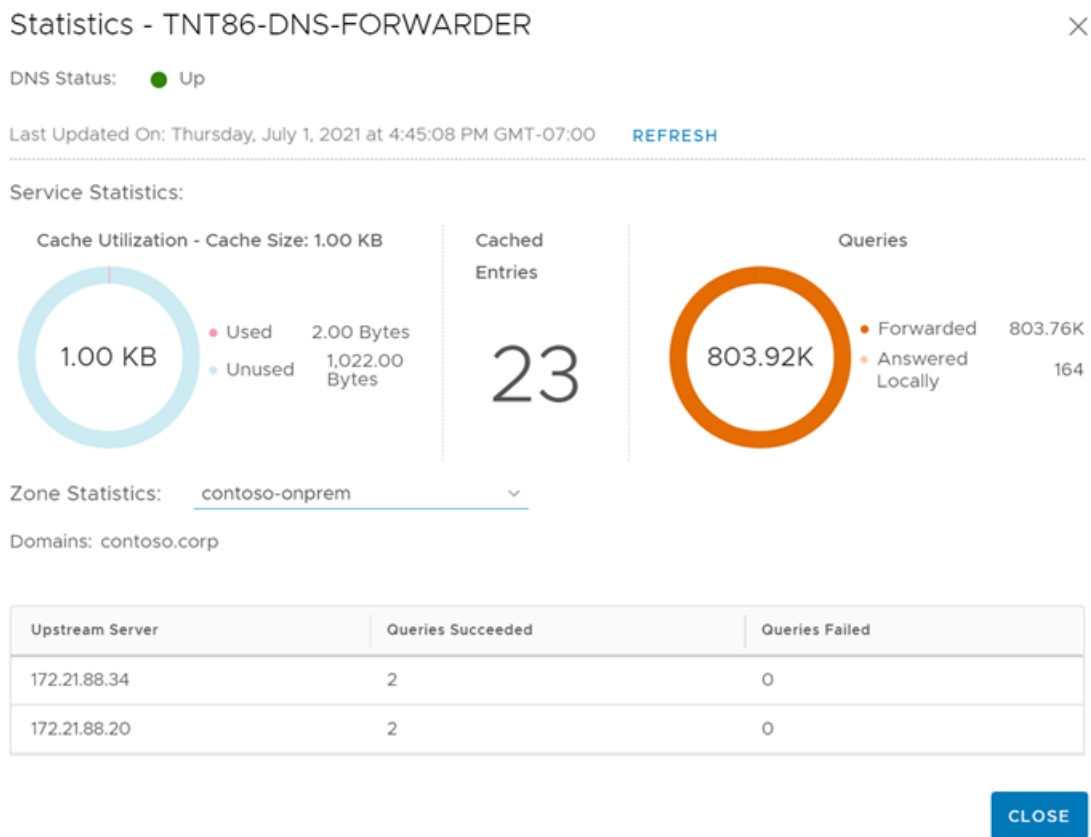
Details for TNT86-DNS-FORWARDER:

- Description: Not Set
- Admin Status: Enabled
- FQDN Zones: contoso-onprem
- Log Level: Info

2. Select **View Statistics**, and then from the **Zone Statistics** drop-down, select your FQDN Zone.

The top half shows the statistics for the entire service, and the bottom half shows the statistics for your

specified zone. In this example, you can see the forwarded queries to the DNS services specified during the configuration of the FQDN zone.



PowerCLI

The NSX-T Policy API lets you run nslookup commands from the NSX-T Data Center DNS Forwarder Service. The required cmdlets are part of the `VMware.VimAutomation.Nsxt` module in PowerCLI. The following example demonstrates output from version 12.3.0 of that module.

1. Connect to your NSX-T Manager cluster.

TIP

You can obtain the IP address of your NSX-T Manager cluster from the Azure portal under **Manage > Identity**.

```
Connect-NsxtServer -Server 10.103.64.3
```

2. Obtain a proxy to the DNS Forwarder's nslookup service.

```
$nslookup = Get-NsxtPolicyService -Name com.vmware.nsx_policy.infra.tier_1s.dns_forwarder.nslookup
```

3. Perform lookups from the DNS Forwarder Service.

```
$response = $nslookup.get('TNT86-T1', 'vc01.contoso.corp')
```

The first parameter in the command is the ID for your private cloud's T1 gateway, which you can obtain from the DNS service tab in the Azure portal.

1. Obtain a raw answer from the lookup using the following properties of the response.

```
$response.dns_answer_per_enforcement_point.raw_answer; ((() DiG 9.10.3-P4-Ubuntu ((() @10.103.64.192
-b 10.103.64.192 vc01.contoso.corp +timeout=5 +tries=3 +nosearch ; (1 server found) ;; global
options: +cmd ;; Got answer: ;; -))HEADER((- opcode: QUERY, status: NOERROR, id: 10684 ;; flags: qr
rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 1 ;; OPT PSEUDOSECTION: ; EDNS: version: 0,
flags;; udp: 4096 ;; QUESTION SECTION: ;vc01.contoso.corp. IN A ;; ANSWER SECTION:
vc01.contoso.corp. 3046 IN A 172.21.90.2 ;; Query time: 0 msec ;; SERVER:
10.103.64.192:53(10.103.64.192) ;; WHEN: Thu Jul 01 23:44:36 UTC 2021 ;; MSG SIZE rcvd: 62
```

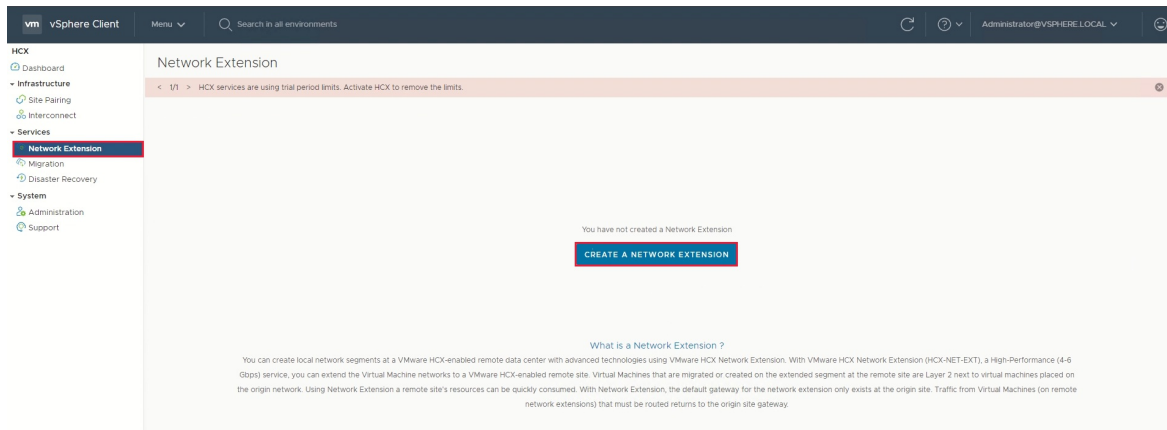
In this example, you can see an answer for the query of vc01.contoso.corp showing an A record with the address 172.21.90.2. Also, this example shows a cached response from the DNS Forwarder Service, so your output may vary slightly.

Create a HCX network extension

12/16/2022 • 2 minutes to read • [Edit Online](#)

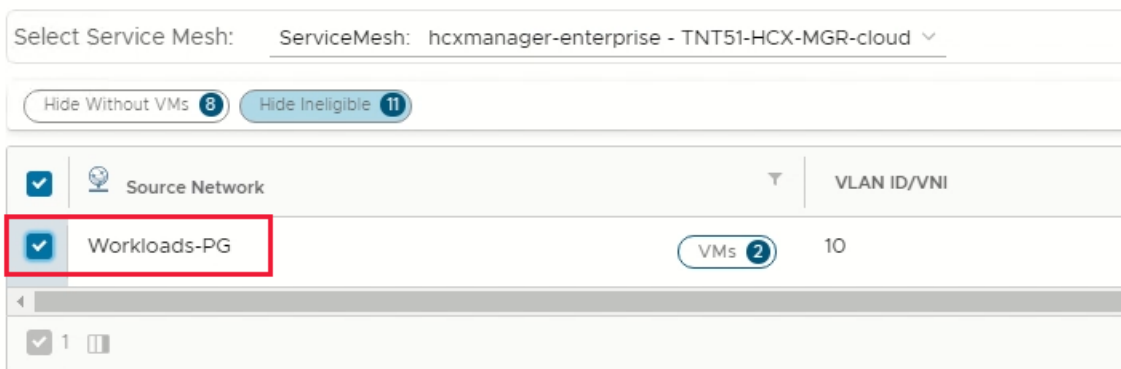
This is an optional step to extend any networks from your on-premises environment to Azure VMware Solution.

1. Under **Services**, select **Network Extension** > **Create a Network Extension**.



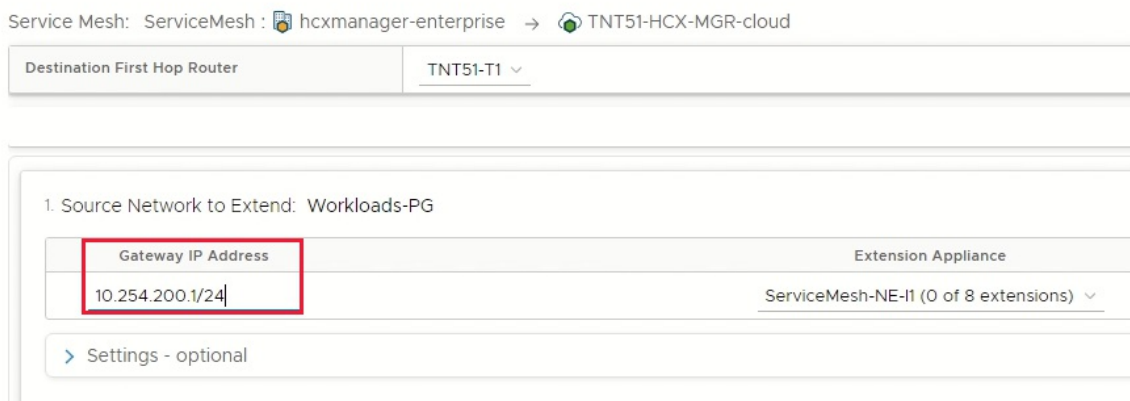
2. Select each of the networks you want to extend to Azure VMware Solution, and then select **Next**.

Extend Networks Select source networks for extension to remote site



3. Enter the on-premises gateway IP for each of the networks you're extending, and then select **Submit**.

Extend Networks Select source networks for extension to remote site



It takes a few minutes for the network extension to finish. When it does, you see the status change to **Extension complete**.

Extensions: 1

+EXTEND NETWORKS

1

Transport Zones / DVS

Extension Appliance	Status
ServiceMesh-NE-11	✓ Extension complete

1 extension

Next steps

Now that you've configured the HCX Network Extension, you can also learn about:

- [VMware HCX Mobility Optimized Networking \(MON\) guidance](#)

VMware HCX Mobility Optimized Networking (MON) guidance

12/16/2022 • 3 minutes to read • [Edit Online](#)

NOTE

HCX Mobility Optimized Networking is officially supported by VMware and Azure VMware Solutions from HCX version 4.1.0.

IMPORTANT

Before you enable HCX MON, please read the below limitations and unsupported configurations:

[Unsupported source configurations for HCX NE](#)

[Limitations for any HCX deployment including MON](#)

[HCX Mobility Optimized Networking \(MON\)](#) is an optional feature to enable when using [HCX Network Extensions \(NE\)](#). MON provides optimal traffic routing under certain scenarios to prevent network tromboning between the on-premises and cloud-based resources on extended networks.

As MON is an enterprise capability of the NE feature, make sure you've enabled the [VMware HCX Enterprise](#) add-on through a [support request](#).

Throughout the migration cycle, MON optimizes application mobility for:

- Optimizing for virtual machine (VM) to VM L2 communication when using stretched networks
- Optimizing and avoiding asymmetric traffic flows between on-premises, Azure VMware Solution, and Azure

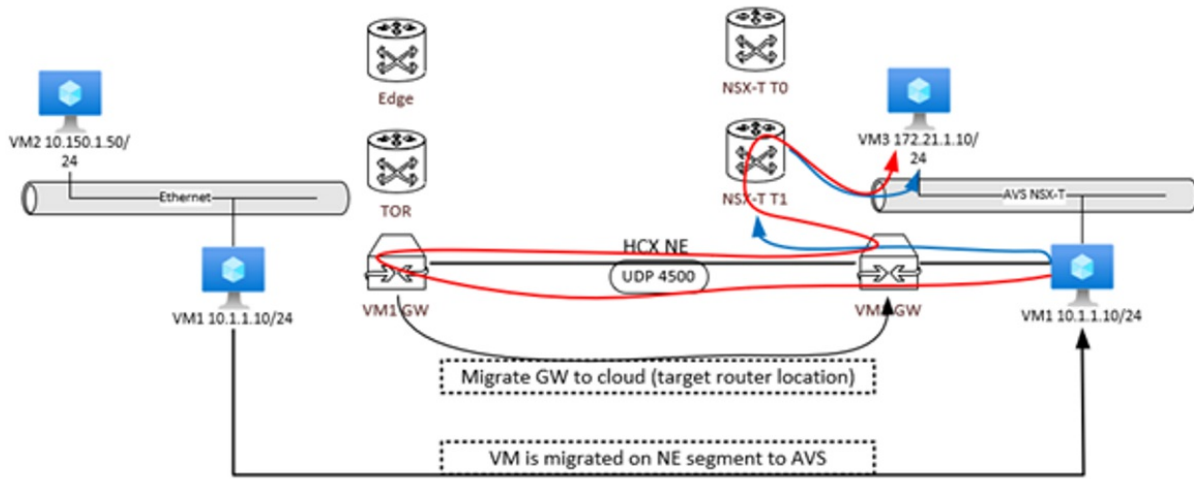
In this article, you'll learn about the Azure VMware Solution-specific use cases for MON.

Optimize traffic flows across standard and stretched segments on the private cloud side

In this scenario, VM1 is migrated to the cloud using the NE, which provides optimal VM to VM latency. As a result, VM1 needs low latency to VM3 on the local Azure VMware Solution segment. We migrate the VM1 gateway from on-premises to Azure VMware Solution (cloud) to ensure an optimal path for traffic (blue line). If the gateway remains on-premises (red line), a tromboning effect and higher latency are observed.

NOTE

When you enable MON without migrating the VM gateway to the cloud side, it doesn't ensure an optimal path for traffic flow. It also doesn't allow the evaluation of policy routes.



Optimize and avoid asymmetric traffic flows

In this scenario, we assume a VM from on-premises has been migrated to Azure VMware Solution and participates in L2, and L3 traffic flows back to on-premises to access services. We also assume some VM communication from Azure (in the Azure VMware Solution connected vNET) could reach down into the Azure VMware Solution private cloud.

IMPORTANT

The main point here is to plan and avoid asymmetric traffic flows carefully.

By default and without using MON, a VM in Azure VMware Solution on a stretched network without MON can communicate back to on-premises using the ExpressRoute preferred path. Ideally, and based on customers use case one should evaluate how a VM on an Azure VMware Solution stretched segment enabled with MON should be traversing back to on-premises either over the NE or the T0 gateway via the ExpressRoute, but keeping traffic flows symmetric.

If choosing the NE path for example, the MON policy routes have to specifically address the subnet on the on-premises side; otherwise, the 0.0.0.0/0 default route is used. Policy routes can be found under the NE segment, selecting advanced. By default, all RFC1918 IP addresses are included in the MON policy routes definition.

Policy Routes



Configure IP subnets assigned to the source environment. [i](#)

Mobility Optimized Networking Site: TNT37-HCX-MGR-cloud

[+ ADD](#) [REMOVE](#) [REFRESH](#)

<input type="checkbox"/>	Network	Send to Source with HCX
<input type="checkbox"/>	10.0.0.0/8	
<input type="checkbox"/>	172.16.0.0/12	
<input type="checkbox"/>	192.168.0.0/16	

SUBMIT

CANCEL

Policy routes are evaluated only if the VM gateway is migrated to the cloud. The effect of this configuration is that any matching subnets for the destination get tunneled over the NE appliance. If not matched, they get routed through the T0 gateway.

NOTE

Special consideration for using MON in Azure VMware Solution is to give the /32 routes advertised over BGP to its peers; this includes on-premises and Azure over the ExpressRoute connection. For example, a VM in Azure learns the path to an Azure VMware Solution VM on an Azure VMware Solution MON enabled segment. Once the return traffic is sent back to the T0 as expected, if the return subnet is an RFC1918 match, traffic is forced over the NE instead of the T0. Then egresses over the ExpressRoute back to Azure on the on-premises side. This can cause confusion for stateful firewalls in the middle and asymmetric routing behavior. It's also a good idea to determine how VMs on NE MON segments will need to access the internet, either via the T0 in Azure VMware Solution or only through the NE back to on-premises.




Fig 1- No Policy Based Route defined

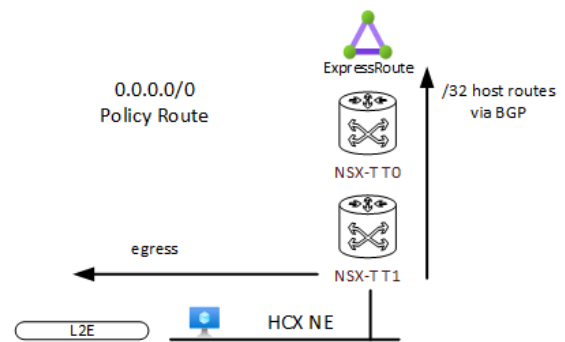


Fig 2- With Policy Route defined to direct all traffic to NE

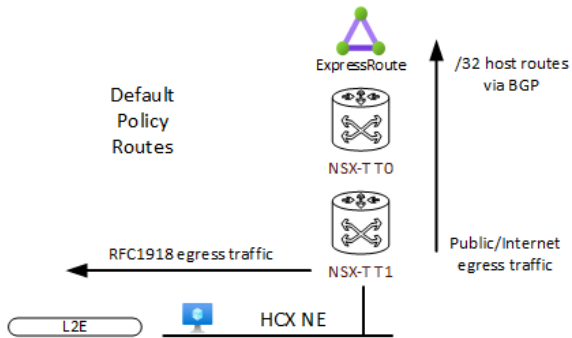


Fig 3- With Policy Route defined using defaults (RFC1918)

As outlined in the above diagram, the importance is to match a policy route to each required subnet. Otherwise, the traffic gets routed over the T0 and not the NE.

To learn more about policy routes, see [Mobility Optimized Networking Policy Routes](#).

Configure NSX-T Data Center network components using Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

An Azure VMware Solution private cloud comes with NSX-T Data Center by default. The private cloud comes pre-provisioned with an NSX-T Data Center Tier-0 gateway in **Active/Active** mode and a default NSX-T Data Center Tier-1 gateway in Active/Standby mode. These gateways let you connect the segments (logical switches) and provide East-West and North-South connectivity.

After deploying Azure VMware Solution, you can configure the necessary NSX-T Data Center objects from the Azure portal. It presents a simplified view of NSX-T Data Center operations a VMware administrator needs daily and is targeted at users not familiar with NSX-T Manager.

You'll have four options to configure NSX-T Data Center components in the Azure VMware Solution console:

- **Segments** - Create segments that display in NSX-T Manager and vCenter Server. For more information, see [Add an NSX-T Data Center segment using the Azure portal](#).
- **DHCP** - Create a DHCP server or DHCP relay if you plan to use DHCP. For more information, see [Use the Azure portal to create a DHCP server or relay](#).
- **Port mirroring** – Create port mirroring to help troubleshoot network issues. For more information, see [Configure port mirroring in the Azure portal](#).
- **DNS** – Create a DNS forwarder to send DNS requests to a designated DNS server for resolution. For more information, see [Configure a DNS forwarder in the Azure portal](#).

IMPORTANT

You'll still have access to the NSX-T Manager console, where you can use the advanced settings mentioned and other NSX-T Data Center features.

Configure port mirroring in the Azure portal

12/16/2022 • 2 minutes to read • [Edit Online](#)

After deploying Azure VMware Solution, you can configure port mirroring from the Azure portal. Port mirroring places a protocol analyzer on the port that receives the mirrored data. It analyzes traffic from a source, a virtual machine (VM), or a group of VMs, and then sent to a defined destination.

In this how-to, you'll configure port mirroring to monitor network traffic, which involves forwarding a copy of each packet from one network switch port to another.

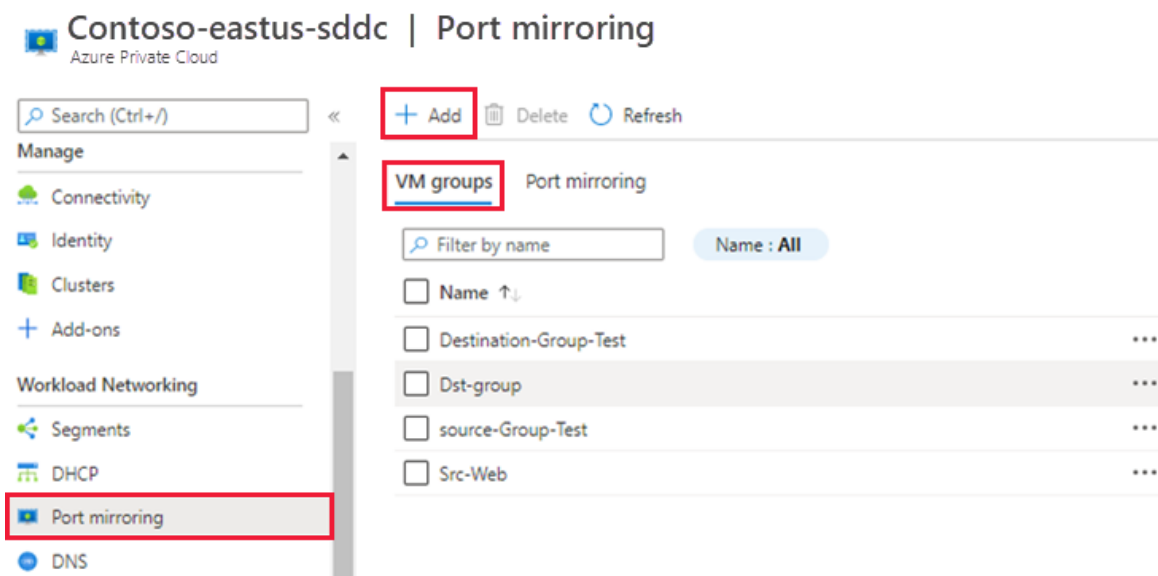
Prerequisites

An Azure VMware Solution private cloud with access to the vCenter Server and NSX-T Manager interfaces. For more information, see the [Configure networking](#) tutorial.

Create the VMs or VM groups

You'll create the source and destination VMs or VM groups. The source group has a single VM or multiple VMs where the traffic is mirrored.

1. In your Azure VMware Solution private cloud, under **Workload Networking**, select **Port mirroring** > **VM groups** > **Add**.



2. Provide a name for the new VM group, select VMs from the list, and then **OK**.
3. Repeat these steps to create the destination VM group.

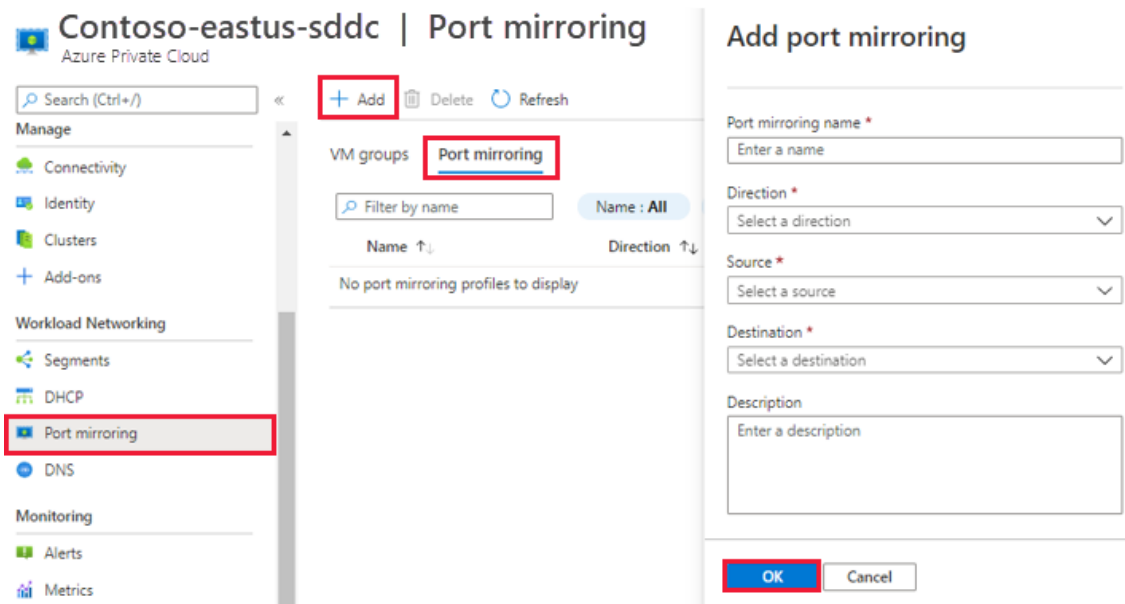
NOTE

Before creating a port mirroring profile, make sure that you've created both the source and destination VM groups.

Create a port mirroring profile

You'll create a port mirroring profile that defines the traffic direction for the source and destination VM groups.

1. Select **Port mirroring** > **Port mirroring** > **Add** and then provide:



- **Port mirroring name** - Descriptive name for the profile.
- **Direction** - Select from Ingress, Egress, or Bi-directional.
- **Source** - Select the source VM group.
- **Destination** - Select the destination VM group.
- **Description** - Enter a description for the port mirroring.


2. Select **OK** to complete the profile.

The profile and VM groups are visible in the Azure VMware Solution console.

Configure a site-to-site VPN in vWAN for Azure VMware Solution

12/16/2022 • 9 minutes to read • [Edit Online](#)

In this article, you'll establish a VPN (IPsec IKEv1 and IKEv2) site-to-site tunnel terminating in the Microsoft Azure Virtual WAN hub. The hub contains the Azure VMware Solution ExpressRoute gateway and the site-to-site VPN gateway. It connects an on-premises VPN device with an Azure VMware Solution endpoint.



Prerequisites

You must have a public-facing IP address terminating on an on-premises VPN device.

Create an Azure Virtual WAN

1. In the portal, in the **Search resources** bar, type **Virtual WAN** in the search box and select **Enter**.
2. Select **Virtual WANs** from the results. On the Virtual WANs page, select **+ Create** to open the **Create WAN** page.
3. On the **Create WAN** page, on the **Basics** tab, fill in the fields. Modify the example values to apply to your environment.

Create WAN ...

Basics Review + create

The virtual WAN resource represents a virtual overlay of your Azure network and is a collection of multiple resources. [Learn more](#)

Project details

Subscription *	<input type="text" value="Content Development"/>
Resource group *	<input type="text" value="TestRG"/>

[Create new](#)

Virtual WAN details

Resource group location *	<input type="text" value="East US"/>
Name *	<input type="text" value="TestVWAN1"/>
Type ⓘ	<input type="text" value="Standard"/>

- **Subscription:** Select the subscription that you want to use.
- **Resource group:** Create new or use existing.
- **Resource group location:** Choose a resource location from the dropdown. A WAN is a global resource and doesn't live in a particular region. However, you must select a region in order to manage and locate the WAN resource that you create.
- **Name:** Type the Name that you want to call your virtual WAN.
- **Type:** Basic or Standard. Select **Standard**. If you select Basic, understand that Basic virtual WANs can only contain Basic hubs. Basic hubs can only be used for site-to-site connections.

4. After you finish filling out the fields, at the bottom of the page, select **Review + Create**.

5. Once validation passes, click **Create** to create the virtual WAN.

Create a virtual hub

A virtual hub is a virtual network that is created and used by Virtual WAN. It's the core of your Virtual WAN network in a region. It can contain gateways for site-to-site and ExpressRoute.

TIP

You can also [create a gateway in an existing hub](#).

1. Go to the virtual WAN that you created. On the virtual WAN page left pane, under the **Connectivity**, select **Hubs**.
2. On the **Hubs** page, select **+New Hub** to open the **Create virtual hub** page.

Create virtual hub ...

Basics Site to site Point to site ExpressRoute Tags Review + create

A virtual hub is a Microsoft-managed virtual network. The hub contains various service endpoints to enable connectivity from your on-premises network (vpnsite). [Learn more](#)

Project details

The hub will be created under the same subscription and resource group as the vWAN.

Subscription	Content Development
Resource group	TestRG

Virtual Hub Details

Region *	West US
Name *	Hub1
Hub private address space * ⓘ	10.1.0.0/16
Virtual hub capacity * ⓘ	2 Routing Infrastructure Units, 3 Gbps Router, Supports 2000 VMs
Hub routing preference * ⓘ	ExpressRoute
Router ASN * ⓘ	65515

3. On the **Create virtual hub** page **Basics** tab, complete the following fields:

- **Region:** Select the region in which you want to deploy the virtual hub.
- **Name:** The name by which you want the virtual hub to be known.
- **Hub private address space:** The hub's address range in CIDR notation. The minimum address space is /24 to create a hub.
- **Virtual hub capacity:** Select from the dropdown. For more information, see [Virtual hub settings](#).
- **Hub routing preference:** This field is only available as part of the virtual hub routing preference preview and can only be viewed in the [preview portal](#). See [Virtual hub routing preference](#) for more information.
- **Router ASN:** Unless necessary, leave the default.

Create a VPN gateway


1. On the **Create virtual hub** page, click **Site to site** to open the **Site to site** tab.

Create virtual hub ...

Basics Site to site Point to site ExpressRoute Tags Review + create

You will need to enable Site to site (VPN gateway) before connecting to VPN sites. You can do this after hub creation, but doing it now will save time and reduce the risk of service interruptions later. [Learn more](#)

Do you want to create a Site to site (VPN gateway)? Yes No

AS Number ⓘ 

*Gateway scale units ⓘ 

Routing preference ⓘ Microsoft network Internet

2. On the **Site to site** tab, complete the following fields:

- Select **Yes** to create a Site-to-site VPN.
- **AS Number**: The AS Number field can't be edited.
- **Gateway scale units**: Select the **Gateway scale units** value from the dropdown. The scale unit lets you pick the aggregate throughput of the VPN gateway being created in the virtual hub to connect sites to.

If you pick 1 scale unit = 500 Mbps, it implies that two instances for redundancy will be created, each having a maximum throughput of 500 Mbps. For example, if you had five branches, each doing 10 Mbps at the branch, you'll need an aggregate of 50 Mbps at the head end. Planning for aggregate capacity of the Azure VPN gateway should be done after assessing the capacity needed to support the number of branches to the hub.

- **Routing preference**: Azure routing preference lets you choose how your traffic routes between Azure and the internet. You can choose to route traffic either via the Microsoft network, or via the ISP network (public internet). These options are also referred to as cold potato routing and hot potato routing, respectively.

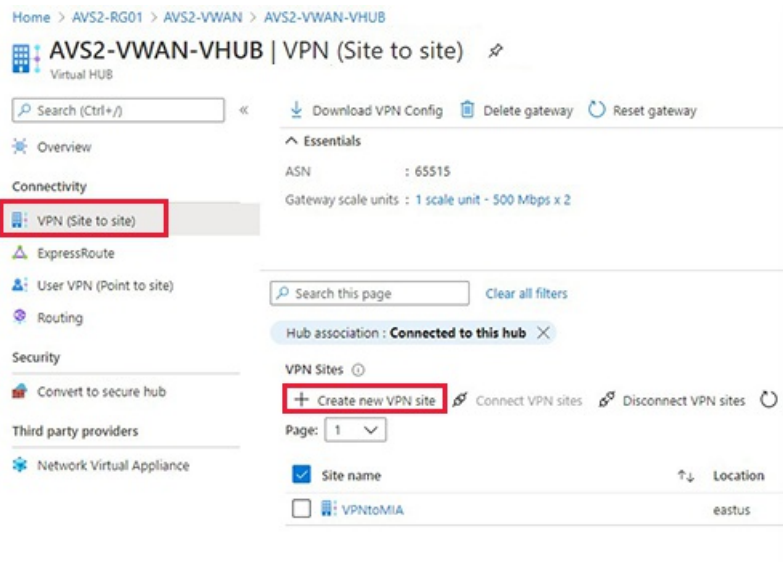
The public IP address in Virtual WAN is assigned by the service, based on the routing option selected. For more information about routing preference via Microsoft network or ISP, see the [Routing preference](#) article.

3. Select **Review + Create** to validate.

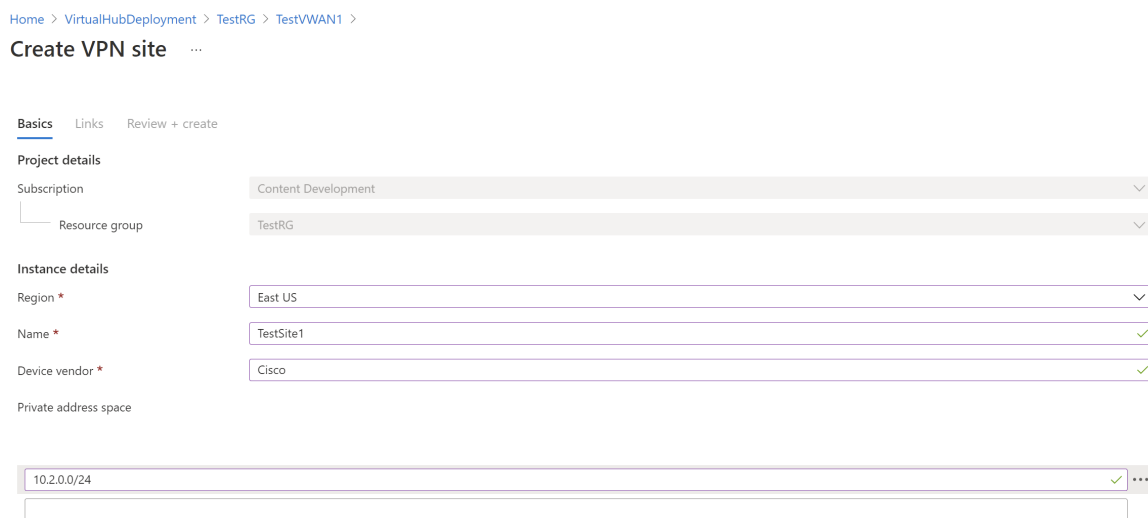
4. Select **Create** to create the hub and gateway. This can take up to 30 minutes. After 30 minutes, **Refresh** to view the hub on the **Hubs** page. Select **Go to resource** to navigate to the resource.

Create a site-to-site VPN

1. In the Azure portal, select the virtual WAN you created earlier.
2. In the **Overview** of the virtual hub, select **Connectivity > VPN (Site-to-site) > Create new VPN site**.



3. On the **Basics** tab, enter the required fields.



- **Region** - Previously referred to as location. It's the location you want to create this site resource in.
- **Name** - The name by which you want to refer to your on-premises site.
- **Device vendor** - The name of the VPN device vendor, for example, Citrix, Cisco, or Barracuda. It helps the Azure Team better understand your environment to add more optimization possibilities in the future or help you troubleshoot.
- **Private address space** - The CIDR IP address space located on your on-premises site. Traffic destined for this address space is routed to your local site. The CIDR block is only required if you **BGP** isn't enabled for the site.

NOTE

If you edit the address space after creating the site (for example, add an additional address space) it can take 8-10 minutes to update the effective routes while the components are recreated.

4. Select **Links** to add information about the physical links at the branch. If you have a Virtual WAN partner CPE device, check with them to see if this information gets exchanged with Azure as a part of the branch information upload set up from their systems.

Specifying link and provider names allow you to distinguish between any number of gateways that may eventually be created as part of the hub. **BGP** and autonomous system number (ASN) must be unique

inside your organization. BGP ensures that both Azure VMware Solution and the on-premises servers advertise their routes across the tunnel. If disabled, the subnets that need to be advertised must be manually maintained. If subnets are missed, HCX fails to form the service mesh.

IMPORTANT

By default, Azure assigns a private IP address from the GatewaySubnet prefix range automatically as the Azure BGP IP address on the Azure VPN gateway. The custom Azure APIPA BGP address is needed when your on-premises VPN devices use an APIPA address (169.254.0.1 to 169.254.255.254) as the BGP IP. Azure VPN Gateway will choose the custom APIPA address if the corresponding local network gateway resource (on-premises network) has an APIPA address as the BGP peer IP. If the local network gateway uses a regular IP address (not APIPA), Azure VPN Gateway will revert to the private IP address from the GatewaySubnet range.

Microsoft Azure

Search resources, services, and docs (G+/)

Create VPN site

Basics **Links** Review + create

At least one link is mandatory

Link name	Link speed	Link provider name	Link IP address / FQDN	Link BGP address	Link ASN
<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>	<input type="text"/>

You can also work with a Virtual WAN partner to create multiple sites simultaneously. [Learn more.](#)

Previous Next: Review + create >

5. Select **Review + create**.

6. Navigate to the virtual hub you want, and deselect **Hub association** to connect your VPN site to the hub.

Home > TestVWAN1 > Hub1

Hub1 | VPN (Site to site)

Virtual Hub

Search (Ctrl+/) Download VPN Config Delete gateway Reset gateway

Overview

Connectivity

ExpressRoute

User VPN (Point to site)

Routing

Security

Convert to secure hub

Third party providers

Network Virtual Appliance

Essentials

ASN : 65515

Gateway scale units : 2 scale units - 1 Gbps x 2

NAT Rules : 0 NAT Rule(s) (Edit)

Bytes in/out : --- MB / --- GB

VPN Gateway : 0489ea6fd-eastus-gw

Gateway configuration : View/Configure

Metrics : View in Azure Monitor

Search this page Clear all filters

Hub association : **Connected to this hub** X

VPN Sites

Check active filters when searching for a VPN site

+ Create new VPN site Connect VPN sites Disconnect VPN sites Refresh

Page: 1

(Optional) Create policy-based VPN site-to-site tunnels

IMPORTANT

This is an optional step and applies only to policy-based VPNs.

[Policy-based VPN setups](#) require on-premises and Azure VMware Solution networks to be specified, including the hub ranges. These ranges specify the encryption domain of the policy-based VPN tunnel on-premises endpoint. The Azure VMware Solution side only requires the policy-based traffic selector indicator to be enabled.

1. In the Azure portal, go to your Virtual WAN hub site and, under **Connectivity**, select **VPN (Site to site)**.
2. Select the VPN Site for which you want to set up a custom IPsec policy.

Home > Resource groups > SEA-Cust13 > wwan-SEA-Cust13 > westushub-SEA-Cust13 - VPN (Site to site)

westushub-SEA-Cust13 - VPN (Site to site)

Virtual HUB

Search (Ctrl+/) «

Download VPN Config Delete gateway Reset gateway

ASN : 65515
Gateway scale units : 1 scale unit - 500 Mbps x 2

Search by site name x Clear all filters

Hub association : **Connected to this hub**

VPN Sites ⓘ

+ Create new VPN site Connect VPN sites Disconnect VPN sites Refresh

Site name	Location
<input checked="" type="checkbox"/> nfgwonprem1	westus
<input type="checkbox"/> nfgwonprem2	westus

3. Select your VPN site name, select **More (...)** at the far right, and then select **Edit VPN Connection**.

Search by site name x Clear all filters

Hub association : **Connected to this hub**

VPN Sites ⓘ

+ Create new VPN site Connect VPN sites Disconnect VPN sites Refresh

Site name	Location	Hub connection status	Site Connection Provisioning Status	Context menu
<input checked="" type="checkbox"/> nfgwonprem1	westus	✔ Succeeded	✔ Connected	...
<input type="checkbox"/> nfgwonprem2	westus	✔ Succeeded	✔ Connected	

- Internet Protocol Security (IPSec), select **Custom**.
- Use policy-based traffic selector, select **Enable**
- Specify the details for **IKE Phase 1** and **IKE Phase 2(ipsec)**.

4. Change the IPsec setting from default to custom and customize the IPsec policy. Then select **Save**.

Edit VPN connection



Virtual HUB

You are editing the connection between the [nfgwonprem1] VPN site and the [westushub-SEA-Cust13] hub.

Connection name

Border gateway protocol Disable Enable

i To edit the site BGP settings, navigate to the site nfgwonprem1.

Links

Link name

Use Azure Private IP Address Yes No

Security settings

Pre-shared key (PSK)

Protocol IKEv2 IKEv1

To change the protocol, please delete the connection first and then create a new connection.

IPSec Default Custom

IKE Phase 1 <input type="radio"/>	Encryption * <input type="text" value="AES128"/>	Integrity/PRF * <input type="text" value="SHA256"/>	DH Group * <input type="text" value="DHGroup14"/>
IKE Phase 2(ipsec) <input type="radio"/>	IPSec Encryption * <input type="text" value="AES256"/>	IPSec Integrity * <input type="text" value="SHA256"/>	PFS Group * <input type="text" value="PFS14"/>

Propagate Default Route Enable Disable

Use policy based traffic selector Enable Disable

Save

Your traffic selectors or subnets that are part of the policy-based encryption domain should be:

- Virtual WAN hub
- Azure VMware Solution private cloud
- Connected Azure virtual network (if present)

Connect your VPN site to the hub

1. Select your VPN site name and then select **Connect VPN sites**.
2. In the **Pre-shared key** field, enter the key previously defined for the on-premises endpoint.

TIP

If you don't have a previously defined key, you can leave this field blank. A key is generated for you automatically.

Connect sites

Virtual HUB



Security settings

Pre-shared key (PSK) ⓘ

Protocol

IKEv2 IKEv1

IPsec ⓘ

Default Custom

Propagate Default Route ⓘ

Enable Disable

Use policy based traffic selector ⓘ

Enable Disable

Configure traffic selector?

Yes No

Connection Mode ⓘ

Default Initiator Only Responder Only

These sites will be connected to the [Hub1] hub.

Site name	↑↓ Region	↑↓
TestSite1	eastus	

- If you're deploying a firewall in the hub and it's the next hop, set the **Propagate Default Route** option to **Enable**.

When enabled, the Virtual WAN hub propagates to a connection only if the hub already learned the default route when deploying a firewall in the hub or if another connected site has forced tunneling enabled. The default route does not originate in the Virtual WAN hub.

- Select **Connect**. After a few minutes, the site shows the connection and connectivity status.

Site name	Location	Connection Status	Connectivity Status
westcentralusBRANCH	westcentralus	● Succeeded	● Connected

Connection Status: Status of the Azure resource for the connection that connects the VPN site to the Azure hub's VPN gateway. Once this control plane operation is successful, the Azure VPN gateway and the on-premises VPN device establish connectivity.

Connectivity Status: Actual connectivity (data path) status between Azure's VPN gateway in the hub and VPN site. It can show any of the following states:

- **Unknown:** Typically seen if the backend systems are working to transition to another status.
- **Connecting:** Azure VPN gateway is trying to reach out to the actual on-premises VPN site.
- **Connected:** Connectivity established between Azure VPN gateway and on-premises VPN site.
- **Disconnected:** Typically seen if disconnected sites for any reason (on-premises or in Azure)

- Download the VPN configuration file and apply it to the on-premises endpoint.

- On the VPN (Site to site) page, near the top, select **Download VPN Config**. Azure creates a storage account in the resource group 'microsoft-network-[location]', where location is the location of the WAN. After you have applied the configuration to your VPN devices, you can delete this storage account.

- b. Once created, select the link to download it.
- c. Apply the configuration to your on-premises VPN device.

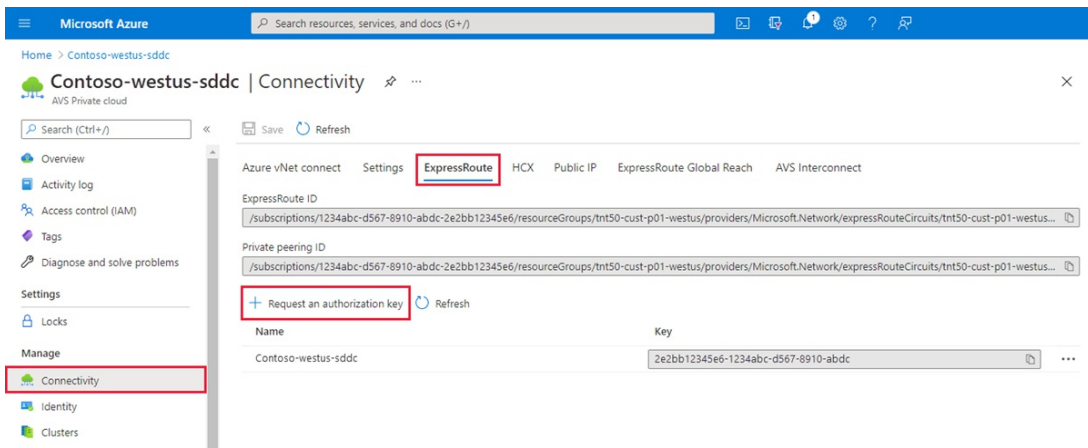
For more information about the configuration file, see [About the VPN device configuration file](#).

6. Patch the Azure VMware Solution ExpressRoute in the Virtual WAN hub.

IMPORTANT

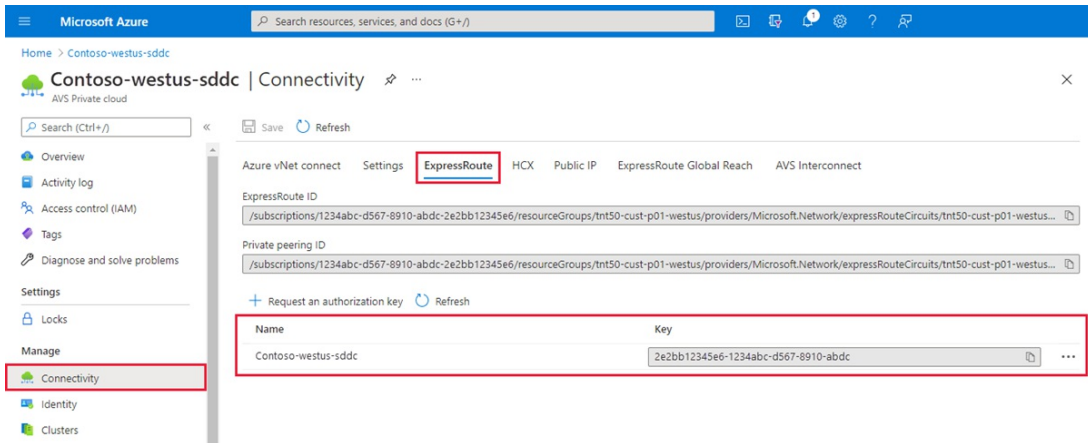
You must first have a private cloud created before you can patch the platform.

- a. In the Azure portal, navigate to the Azure VMware Solution private cloud. Select **Manage > Connectivity > ExpressRoute** and then select **+ Request an authorization key**.



- b. Provide a name for it and select **Create**.

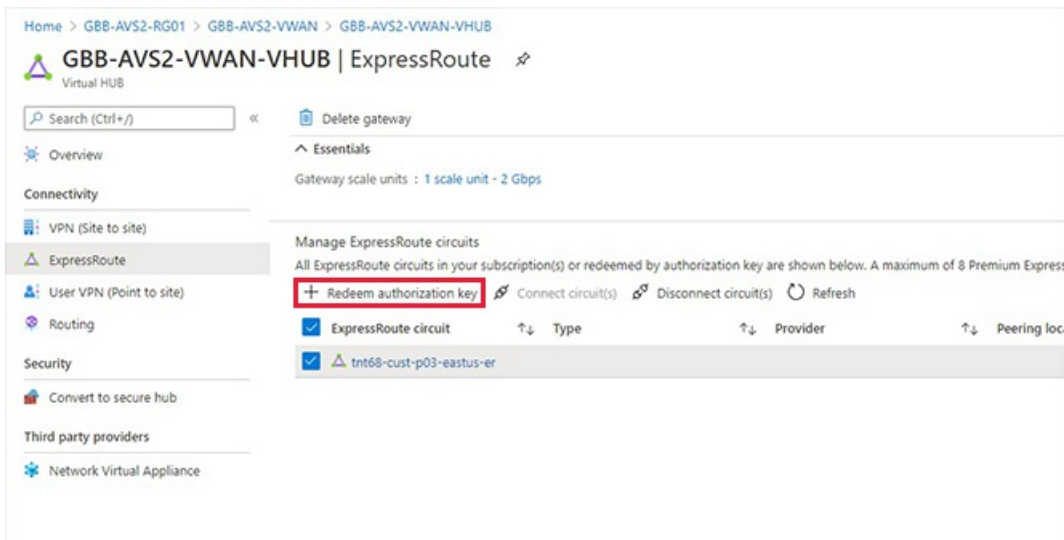
It may take about 30 seconds to create the key. Once created, the new key appears in the list of authorization keys for the private cloud.



- c. Copy the authorization key and ExpressRoute ID. You'll need them to complete the peering. The authorization key disappears after some time, so copy it as soon as it appears.

7. Link Azure VMware Solution and the VPN gateway together in the Virtual WAN hub. You'll use the authorization key and ExpressRoute ID (peer circuit URI) from the previous step.

- a. Select your ExpressRoute gateway and then select **Redeem authorization key**.



- b. Paste the authorization key in the **Authorization Key** field.
 - c. Paste the ExpressRoute ID into the **Peer circuit URI** field.
 - d. Select **Automatically associate this ExpressRoute circuit with the hub** check box.
 - e. Select **Add** to establish the link.
8. Test your connection by [creating an NSX-T Data Center segment](#) and provisioning a VM on the network. Ping both the on-premises and Azure VMware Solution endpoints.

NOTE

Wait approximately 5 minutes before you test connectivity from a client behind your ExpressRoute circuit, for example, a VM in the VNet that you created earlier.

Configure storage policy

12/16/2022 • 4 minutes to read • [Edit Online](#)

VMware vSAN storage policies define storage requirements for your virtual machines (VMs). These policies guarantee the required level of service for your VMs because they determine how storage is allocated to the VM. Each VM deployed to a vSAN datastore is assigned at least one VM storage policy.

You can assign a VM storage policy in an initial deployment of a VM or when you do other VM operations, such as cloning or migrating. Post-deployment cloudadmin users or equivalent roles can't change the default storage policy for a VM. However, **VM storage policy** per disk changes is permitted.

The Run command lets authorized users change the default or existing VM storage policy to an available policy for a VM post-deployment. There are no changes made on the disk-level VM storage policy. You can always change the disk level VM storage policy as per your requirements.

NOTE

Run commands are executed one at a time in the order submitted.

In this how-to, you learn how to:

- List all storage policies
- Set the storage policy for a VM
- Specify default storage policy for a cluster

Prerequisites

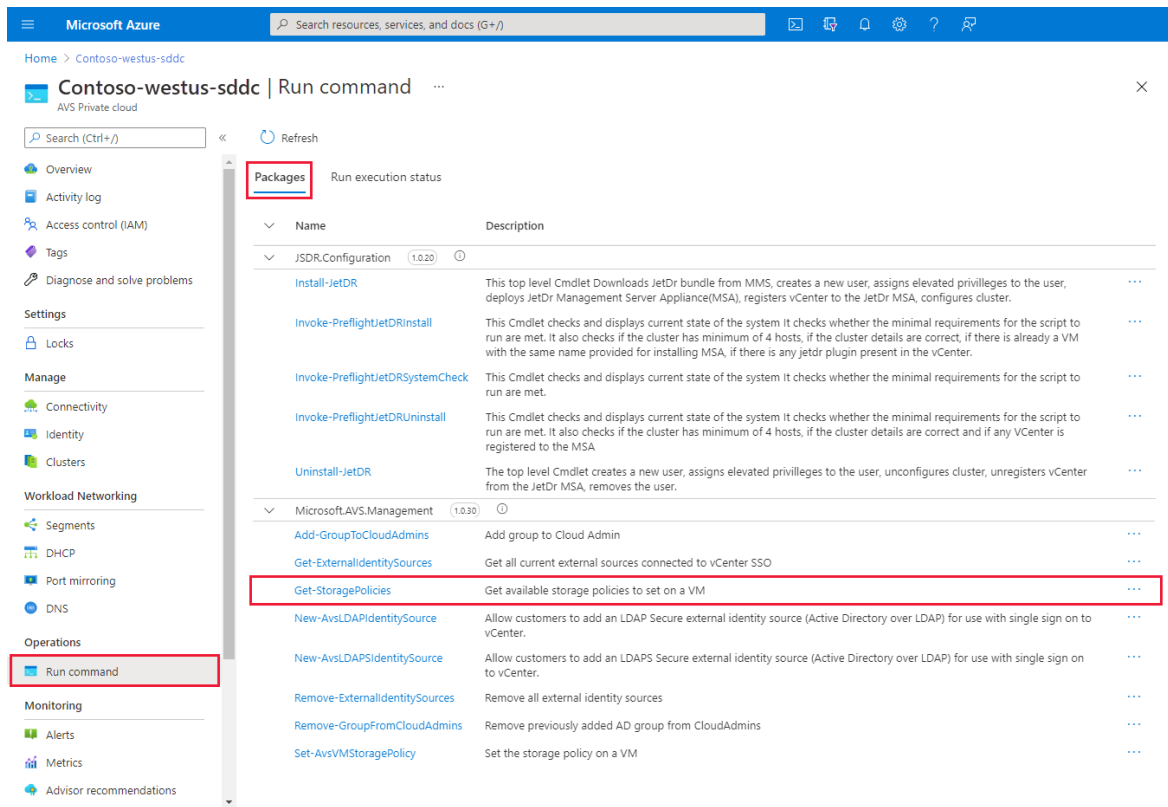
Make sure that the [minimum level of hosts are met](#).

RAID CONFIGURATION	FAILURES TO TOLERATE (FTT)	MINIMUM HOSTS REQUIRED
RAID-1 (Mirroring) Default setting.	1	3
RAID-5 (Erasure Coding)	1	4
RAID-1 (Mirroring)	2	5
RAID-6 (Erasure Coding)	2	6
RAID-1 (Mirroring)	3	7

List storage policies

You'll run the `Get-StoragePolicy` cmdlet to list the vSAN based storage policies available to set on a VM.

1. Sign in to the [Azure portal](#).
2. Select **Run command** > **Packages** > **Get-StoragePolicies**.



3. Provide the required values or change the default values, and then select Run.

Run command - Get-StoragePolicies

Get available storage policies to set on a VM

Details

Retain up to

day hour minute

Specify name for execution *

Timeout *

hour minute second

FIELD	VALUE
Retain up to	Retention period of the cmdlet output. The default value is 60.
Specify name for execution	Alphanumeric name, for example, Get-StoragePolicies-Exec1 .
Timeout	The period after which a cmdlet exits if taking too long to finish.

4. Check Notifications to see the progress.

Set storage policy on VM

You'll run the `Set-VMStoragePolicy` cmdlet to modify vSAN-based storage policies on a default cluster, individual VM, or group of VMs sharing a similar VM name. For example, if you have three VMs named "MyVM1", "MyVM2", and "MyVM3", supplying "MyVM*" to the VMName parameter would change the StoragePolicy on all three VMs.

NOTE

You cannot use the vSphere Client to change the default storage policy or any existing storage policies for a VM.

1. Select **Run command > Packages > Set-VMStoragePolicy**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
VMName	Name of the target VM.
StoragePolicyName	Name of the storage policy to set. For example, RAID-FTT-1.
Retain up to	Retention period of the cmdlet output. The default value is 60.
Specify name for execution	Alphanumeric name, for example, <code>changeVMStoragePolicy</code> .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** to see the progress.

Set storage policy on all VMs in a location

You'll run the `Set-LocationStoragePolicy` cmdlet to Modify vSAN based storage policies on all VMs in a location where a location is the name of a cluster, resource pool, or folder. For example, if you have 3 VMs in Cluster-3, supplying "Cluster-3" would change the storage policy on all 3 VMs.

NOTE

You cannot use the vSphere Client to change the default storage policy or any existing storage policies for a VM.

1. Select **Run command > Packages > Set-LocationStoragePolicy**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
Location	Name of the target VM.
StoragePolicyName	Name of the storage policy to set. For example, RAID-FTT-1.
Retain up to	Retention period of the cmdlet output. The default value is 60.

FIELD	VALUE
Specify name for execution	Alphanumeric name, for example, changeVMStoragePolicy .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** to see the progress.

Specify storage policy for a cluster

You'll run the `Set-ClusterDefaultStoragePolicy` cmdlet to specify default storage policy for a cluster,

1. Select **Run command > Packages > Set-ClusterDefaultStoragePolicy**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
ClusterName	Name of the cluster.
StoragePolicyName	Name of the storage policy to set. For example, RAID-FTT-1 .
Retain up to	Retention period of the cmdlet output. The default value is 60.
Specify name for execution	Alphanumeric name, for example, Set-ClusterDefaultStoragePolicy-Exec1 .
Timeout	The period after which a cmdlet exits if taking too long to finish.

3. Check **Notifications** to see the progress.

Next steps

Now that you've learned how to configure VMware vSAN storage policies, you can learn more about:

- [How to attach disk pools to Azure VMware Solution hosts \(Preview\)](#) - You can use disks as the persistent storage for Azure VMware Solution for optimal cost and performance.
- [How to configure external identity for vCenter](#) - vCenter Server has a built-in local user called cloudadmin and assigned to the CloudAdmin role. The local cloudadmin user is used to set up users in Active Directory (AD). With the Run command feature, you can configure Active Directory over LDAP or LDAPS for vCenter as an external identity source.

Configure VMware syslogs for Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

Diagnostic settings are used to configure streaming export of platform logs and metrics for a resource to the destination of your choice. You can create up to five different diagnostic settings to send different logs and metrics to independent destinations.

In this article, you'll configure a diagnostic setting to collect VMware syslogs for your Azure VMware Solution private cloud. You'll store the syslog to a storage account to view the vCenter Server logs and analyze for diagnostic purposes.

IMPORTANT

The **VMware syslogs** contains the following logs:

- NSX-T Data Center Distributed Firewall logs
- NSX-T Manager logs
- NSX-T Data Center Gateway Firewall logs
- ESXi logs
- vCenter Server logs
- NSX-T Data Center Edge Appliance logs

Prerequisites

Make sure you have an Azure VMware Solution private cloud with access to the vCenter Server and NSX-T Manager interfaces.

Configure diagnostic settings

1. From your Azure VMware Solution private cloud, select **Diagnostic settings**, then **Add diagnostic settings**.

Microsoft Azure | Search resources, services, and docs (G+)

Home > Contoso-westus-sddc

Contoso-westus-sddc | Diagnostic settings

AVS Private cloud

Search (Ctrl+/) Refresh Feedback

Diagnostic settings are used to configure streaming export of platform logs and metrics for a resource to the destination of your choice. You may create up to five different diagnostic settings to send different logs and metrics to independent destinations. [Learn more about diagnostic settings](#)

Diagnostic settings

Name	Storage account	Event hub	Log Analytics worksp...	Partner solution	Edit setting
amanejawestus04	testwestus05	-	-	-	Edit setting
service	testwestus03	-	-	-	Edit setting
testamanejawestus	-	-	vcenterlogs-hcl	-	Edit setting

[+ Add diagnostic setting](#)

Click 'Add Diagnostic setting' above to configure the collection of the following data:

- vmwaresyslog
- AllMetrics

1. Select the **vmwaresyslog**, **All metrics**, and select one of the following options presented.

Send to Log Analytics workspace

How to set up Log Analytics

A Log Analytics workspace:

- Contains your AVS private cloud logs.
- Is the workspace from which you can take desired actions, such as querying for logs.

In this section, you'll:

- Configure a Log Analytics workspace
- Create a diagnostic setting in your private cloud to send your logs to this workspace

Create a resource

1. In the Azure portal, go to **Create a resource**.
2. Search for "Log Analytics Workspace" and click **Create** -> **Log Analytics Workspace**.

[Home](#) > [Create a resource](#) >

Marketplace

Get Started

Service Providers

Management

Private Marketplace

Private Offer Management

My Marketplace

Favorites

Recently created

Private products

 Azure benefit eligible only ⓘ

Showing 1 to 20 of 174 results for 'Log



Log Analytics Workspace

Microsoft

Azure Service

Collect, search and visualize machine data from on-premises and cloud

1. Enter the Subscription you intend to use, the Resource Group that'll house this workspace. Give it a name and select a region.
2. Click **Review + Create**.

[Home](#) > [Create a resource](#) > [Marketplace](#) >

Create Log Analytics workspace ...

Basics Tags Review + Create

i A Log Analytics workspace is the basic management unit of Azure Monitor Logs. There are specific considerations you should take when creating a new Log Analytics workspace. [Learn more](#) ✕

With Azure Monitor Logs you can easily store, retain, and query data collected from your monitored resources in Azure and other environments for valuable insights. A Log Analytics workspace is the logical storage unit where your log data is collected and stored.

Project details

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ

Resource group * ⓘ [Create new](#)

Instance details

Name * ⓘ ✓

Region * ⓘ

Add a diagnostic setting


Next, we add a diagnostic setting in your AVS private cloud, so it knows where to send your logs to.

[Home](#) > [vh-private-cloud](#)


vh-private-cloud | Diagnostic settings 🔗 ...

AVS Private cloud


[Refresh](#) [Feedback](#)

 Datastores

Workload Networking

 Segments

 DHCP

 Port mirroring

 DNS

 Internet connectivity

Operations

 Azure Arc (preview)

Diagnostic settings are used to configure streaming export of platform logs and metrics [about diagnostic settings](#)

Diagnostic settings

Name	Storage account
vh-diagnostic-settings	-
vh-diagnostic-settings-1	-

[+ Add diagnostic setting](#)

Click 'Add Diagnostic setting' above to configure the collection of the following data:

- VMware Syslog
- AllMetrics

1. Click your AVS private cloud. Go to Diagnostic settings on the left-hand menu under Monitoring. Select **Add diagnostic setting**.
2. Give your diagnostic setting a name. Select the log categories you are interested in sending to your Log Analytics workspace.
3. Make sure to select the checkbox next to **Send to Log Analytics workspace**. Select the Subscription your Log Analytics workspace lives in and the Log Analytics workspace. Click **Save** on the top left.

[Home](#) > [vh-private-cloud](#) | [Diagnostic settings](#) >

Diagnostic setting ...

 Save  Discard  Delete  Feedback

A diagnostic setting specifies a list of categories of platform logs and/or metrics that you want to collect from a resource, and one or more destinations that you would stream them to. Normal usage charges for the destination will occur. [Learn more about the different log categories and contents of those logs](#)

Diagnostic setting name * 

Logs

Category groups ⓘ

audit allLogs

Categories

VMware Syslog

Metrics

AllMetrics

Destination details

Send to Log Analytics workspace

Subscription

Log Analytics workspace

Archive to a storage account

Stream to an event hub

Send to partner solution

At this point, your Log Analytics workspace has been successfully configured to receive logs from your AVS private cloud.

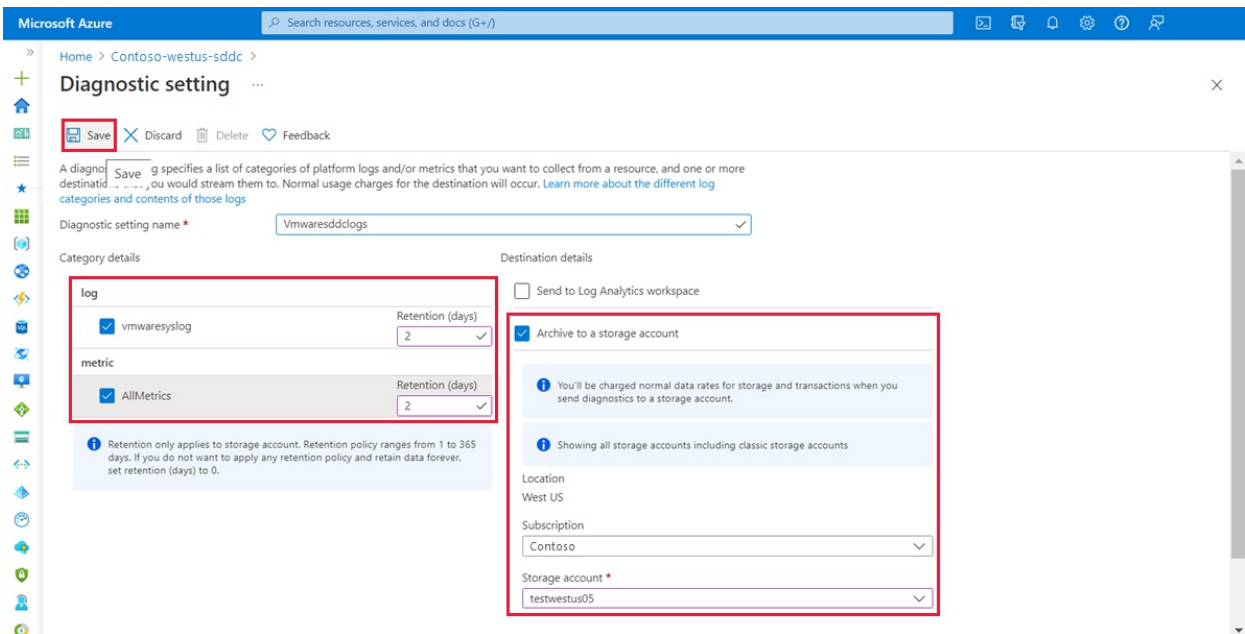
Search and analyze logs using Kusto

Now that you've successfully configured your logs to go to your Log Analytics workspace, you can use that data to gain meaningful insights with Log Analytics' search feature. Log Analytics uses a language called the Kusto Query Language (or Kusto) to search through your logs.

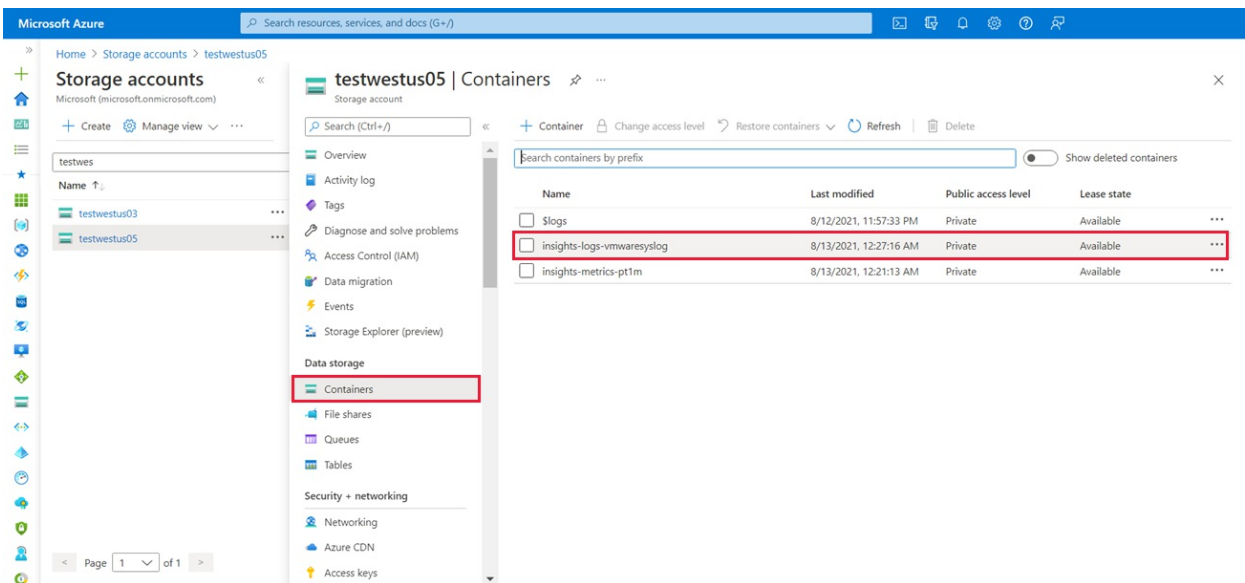
For more information, see [Data analysis in Azure Data Explorer with Kusto Query Language](#).

Archive to storage account

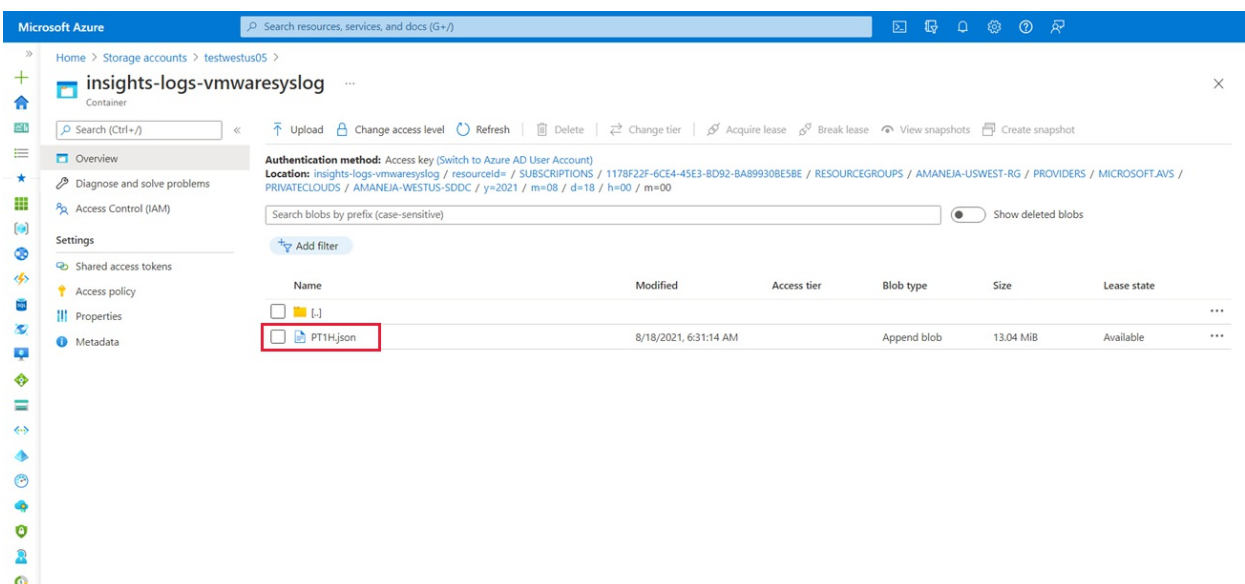
1. In **Diagnostic setting**, select the storage account where you want to store the logs and select **Save**.



2. Go to your Storage accounts, verify Insight logs vmwarelog has been created, select it.



3. Browse Insight logs vmwarelog to locate and download the json file to view the logs.



1. In **Diagnostic setting**, under **Destination details**, select **Stream to an Event Hub**.
2. From the **Event Hub namespace** drop-down menu, choose where you want to send the logs, select, and **Save**.

The screenshot shows the 'Diagnostic setting' configuration page in the Microsoft Azure portal. The page title is 'Diagnostic setting' and the breadcrumb is 'Home > SDDC-WestUS > Diagnostic setting'. At the top, there are buttons for 'Save', 'Discard', 'Delete', and 'Feedback'. Below this is a description: 'A diagnostic setting specifies a list of categories of platform logs and/or metrics that you want to collect from a resource, and one or more destinations that you would stream them to. Normal usage charges for the destination will occur. Learn more about the different log categories and contents of those logs'. The 'Diagnostic setting name' is 'EventHubSysLogs'. Under 'Category details', the 'log' category 'vmwaresyslog' is selected. Under 'Destination details', the 'Stream to an event hub' option is selected and highlighted with a red box. The configuration for this option includes: 'Subscription' (AVS Dogfood), 'Event hub namespace' (amaneja-eventhub-ns), 'Event hub name (optional)' (amaneja-westus-eventhub), and 'Event hub policy name' (RootManageSharedAccessKey). There is also an unchecked option 'Send to partner solution'.

Configure Windows Server Failover Cluster on Azure VMware Solution vSAN

12/16/2022 • 7 minutes to read • [Edit Online](#)

In this article, you'll learn how to configure [Failover Clustering in Windows Server](#) on Azure VMware Solution vSAN with native shared disks.

Windows Server Failover Cluster, previously known as Microsoft Service Cluster Service (MSCS), is a Windows Server Operating System (OS) feature. WSFC is a business-critical feature, and for many applications is required. For example, WSFC is required for the following configurations:

- SQL server configured as:
 - Always On Failover Cluster Instance (FCI), for instance-level high availability.
 - Always On Availability Group (AG), for database-level high availability.
- Windows File Services:
 - Generic File share running on active cluster node.
 - Scale-Out File Server (SOFS), which stores files in cluster shared volumes (CSV).
 - Storage Spaces Direct (S2D); local disks used to create storage pools across different cluster nodes.

You can host the WSFC cluster on different Azure VMware Solution instances, known as Cluster-Across-Box (CAB). You can also place the WSFC cluster on a single Azure VMware Solution node. This configuration is known as Cluster-in-a-Box (CIB). We don't recommend using a CIB solution for a production implementation, use CAB instead with placement policies. Were the single Azure VMware Solution node to fail, all WSFC cluster nodes would be powered off, and the application would experience downtime. Azure VMware Solution requires a minimum of three nodes in a private cloud cluster.

It's important to deploy a supported WSFC configuration. You'll want your solution to be supported on VMware vSphere and with Azure VMware Solution. VMware provides a detailed document about WSFC on vSphere 7.0, [Setup for Failover Clustering and Microsoft Cluster Service](#).

This article focuses on WSFC on Windows Server 2016 and Windows Server 2019. Unfortunately, older Windows Server versions are out of [mainstream support](#), so we don't consider them here.

You'll need first to [create a WSFC](#). Then, use the information we provide in this article to specify a WSFC deployment on Azure VMware Solution.

Prerequisites

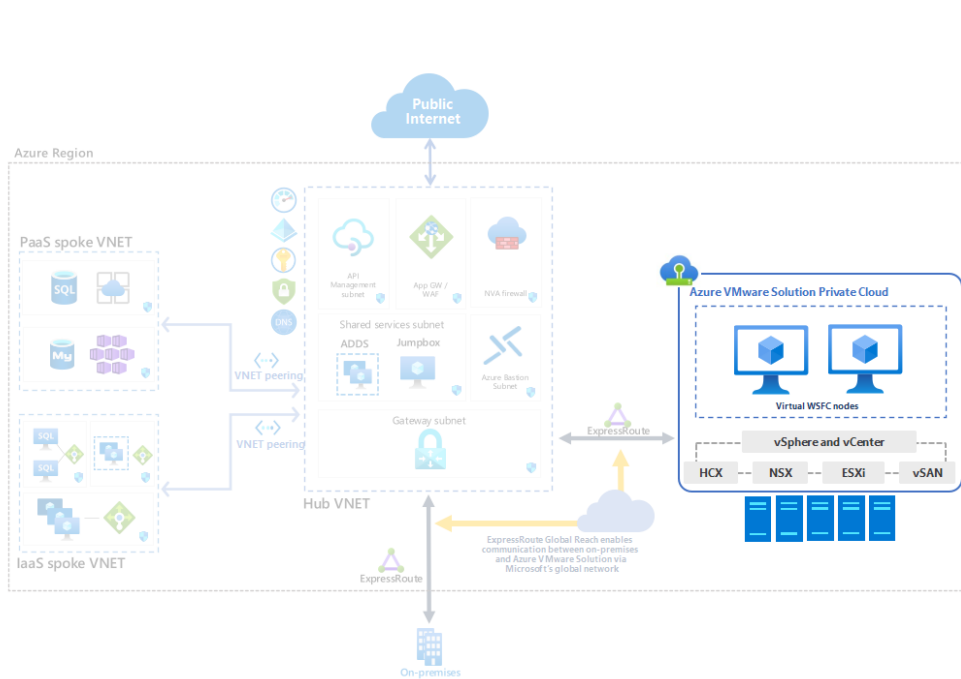
- Azure VMware Solution environment
- Microsoft Windows Server OS installation media

Reference architecture

Azure VMware Solution provides native support for virtualized WSFC. It supports SCSI-3 Persistent Reservations (SCSI3PR) on a virtual disk level. WSFC requires this support to arbitrate access to a shared disk between nodes. Support of SCSI3PRs enables configuration of WSFC with a disk resource shared between VMs natively on vSAN datastores.

The following diagram illustrates the architecture of WSFC virtual nodes on an Azure VMware Solution private cloud. It shows where Azure VMware Solution resides, including the WSFC virtual servers (blue box), in relation

to the broader Azure platform. This diagram illustrates a typical hub-spoke architecture, but a similar setup is possible using Azure Virtual WAN. Both offer all the value other Azure services can bring you.



Supported configurations

Currently, the configurations supported are:

- Microsoft Windows Server 2012 or later
- Up to five failover clustering nodes per cluster
- Up to four PVSCSI adapters per VM
- Up to 64 disks per PVSCSI adapter

Virtual machine configuration requirements

WSFC node configuration parameters

- Install the latest VMware Tools on each WSFC node.
- Mixing non-shared and shared disks on a single virtual SCSI adapter isn't supported. For example, if the system disk (drive C:) is attached to SCSI0:0, the first shared disk would be attached to SCSI1:0. A VM node of a WSFC has the same virtual SCSI controller maximum as an ordinary VM - up to four (4) virtual SCSI Controllers.
- Virtual discs SCSI IDs should be consistent between all VMs hosting nodes of the same WSFC.

COMPONENT	REQUIREMENTS
VM hardware version	11 or above to support Live vMotion.
Virtual NIC	VMXNET3 paravirtualized network interface card (NIC); enable the in-guest Windows Receive Side Scaling (RSS) on the virtual NIC.
Memory	Use full VM reservation memory for nodes in the WSFC cluster.

COMPONENT	REQUIREMENTS
Increase the I/O timeout of each WSFC node.	Modify HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\ \Disk\TimeOutValueSet to 60 seconds or more. (If you recreate the cluster, this value might be reset to its default, so you must change it again.)
Windows cluster health monitoring	The value of the SameSubnetThreshold Parameter of Windows cluster health monitoring must be modified to allow 10 missed heartbeats at minimum. It's the default in Windows Server 2016 . This recommendation applies to all applications using WSFC, including shared and non-shared disks.

WSFC node - Boot disks configuration parameters

COMPONENT	REQUIREMENTS
SCSI Controller Type	LSI Logic SAS
Disk mode	Virtual
SCSI bus sharing	None
Modify advanced settings for a virtual SCSI controller hosting the boot device.	Add the following advanced settings to each WSFC node: scsiX.returnNoConnectDuringAPD = "TRUE" scsiX.returnBusyOnNoConnectStatus = "FALSE" Where X is the boot device SCSI bus controller ID number. By default, X is set to 0.

WSFC node - Shared disks configuration parameters

COMPONENT	REQUIREMENTS
SCSI Controller Type	VMware Paravirtualized (PVSCSI)
Disk mode	Independent - Persistent (step 2 in illustration below). By using this setting, you ensure that all disks are excluded from snapshots. Snapshots aren't supported for WSFC- based VMs.
SCSI bus sharing	Physical (step 1 in illustration below)
Multi-writer flag	Not used
Disk format	Thick provisioned. (Eager Zeroed Thick (EZT) isn't required with vSAN.)

[ADD NEW DEVICE](#)

▼ New Hard disk *	750	GB
Maximum Size	33.92 TB	
VM storage policy	vSAN Default Storage Policy ▼	
Location	Store with the virtual machine ▼	
Disk Provisioning	As defined in the VM storage policy ▼	
Sharing	Unspecified ▼	
Shares	Normal ▼	1000
Limit - IOPs	Unlimited ▼	
Virtual flash read cache	0	MB ▼
Disk Mode	2 Independent - Persistent ▼	
Virtual Device Node	SCSI controller 0 ▼	SCSI(0:1) New Hard disk ▼
> SCSI controller 0	LSI Logic SAS	
▼ New SCSI controller *	VMware Paravirtual	
Change Type	1 VMware Paravirtual ▼	
SCSI Bus Sharing	Physical ▼	
> Network adapter 1	workload-segment-01 ▼	<input checked="" type="checkbox"/> Connected
> CD/DVD drive 1	Datastore ISO File ▼	<input type="checkbox"/> Connected

Non-supported scenarios

The following functionalities aren't supported for WSFC on Azure VMware Solution:

- NFS data stores
- Storage Spaces
- vSAN using iSCSI Service
- vSAN Stretched Cluster
- Enhanced vMotion Compatibility (EVC)
- vSphere Fault Tolerance (FT)
- Snapshots
- Live (online) storage vMotion
- N-Port ID Virtualization (NPIV)

Hot changes to virtual machine hardware might disrupt the heartbeat between the WSFC nodes.

The following activities aren't supported and might cause WSFC node failover:

- Hot adding memory
- Hot adding CPU
- Using snapshots
- Increasing the size of a shared disk
- Pausing and resuming the virtual machine state
- Memory over-commitment leading to ESXi swapping or VM memory ballooning
- Hot Extend Local VMDK file, even if it isn't associated with SCSI bus sharing controller

Configure WSFC with shared disks on Azure VMware Solution vSAN

1. Ensure that an Active Directory environment is available.
2. Create virtual machines (VMs) on the vSAN datastore.
3. Power on all VMs, configure the hostname and IP addresses, join all VMs to an Active Directory domain, and install the latest available OS updates.
4. Install the latest VMware Tools.
5. Enable and configure the Windows Server Failover Cluster feature on each VM.
6. Configure a Cluster Witness for quorum (this can be a file share witness).
7. Power off all nodes of the WSFC cluster.
8. Add one or more Paravirtual SCSI controllers (up to four) to each VM part of the WSFC. Use the settings per the previous paragraphs.
9. On the first cluster node, add all needed shared disks using **Add New Device > Hard Disk**. Leave Disk sharing as **Unspecified** (default) and Disk mode as **Independent - Persistent**. Then attach it to the controller(s) created in the previous steps.
10. Continue with the remaining WSFC nodes. Add the disks created in the previous step by selecting **Add New Device > Existing Hard Disk**. Be sure to maintain the same disk SCSI IDs on all WSFC nodes.
11. Power on the first WSFC node; sign in and open the disk management console (mmc). Make sure the added shared disks can be managed by the OS and are initialized. Format the disks and assign a drive letter.
12. Power on the other WSFC nodes.
13. Add the disk to the WSFC cluster using the **Add Disk wizard** and add them to a Cluster Shared Volume.
14. Test a failover using the **Move disk wizard** and make sure the WSFC cluster with shared disks works properly.
15. Run the **Validation Cluster wizard** to confirm whether the cluster and its nodes are working properly.

It's important to keep the following specific items from the Cluster Validation test in mind:

- **Validate Storage Spaces Persistent Reservation.** If you aren't using Storage Spaces with your cluster (such as on Azure VMware Solution vSAN), this test isn't applicable. You can ignore any results of the Validate Storage Spaces Persistent Reservation test including this warning. To avoid warnings, you can exclude this test.
 - **Validate Network Communication.** The Cluster Validation test displays a warning indicating that only one network interface per cluster node is available. You can ignore this warning. Azure VMware Solution provides the required availability and performance needed, since the nodes are connected to one of the NSX-T Data Center segments. However, keep this item as part of the Cluster Validation test, as it validates other aspects of network communication.
16. Create the relevant Placement Policies to situate the WSFC VMs on the correct Azure VMware Solution nodes depending upon the WSFC CIB or CAB configuration. To do so, you need a host-to-VM affinity rule. This way, cluster nodes will run on the same or separate Azure VMware Solution host(s) respectively.

Related information

- [Failover Clustering in Windows Server](#)
- [Guidelines for Microsoft Clustering on vSphere \(1037959\) \(vmware.com\)](#)
- [About Setup for Failover Clustering and Microsoft Cluster Service \(vmware.com\)](#)

- [vSAN 6.7 U3 - WSFC with Shared Disks & SCSI-3 Persistent Reservations \(vmware.com\)](#)
- [Azure VMware Solution limits](#)

Next steps

Now that you've covered setting up a WSFC in Azure VMware Solution, you may want to learn about:

- Setting up your new WSFC by adding more applications that require the WSFC capability. For instance, SQL Server and SAP ASCS.
- Setting up a backup solution.
 - [Setting up Azure Backup Server for Azure VMware Solution](#)
 - [Backup solutions for Azure VMware Solution virtual machines](#)

Connect multiple Azure VMware Solution private clouds in the same region

12/16/2022 • 2 minutes to read • [Edit Online](#)

The **AVS Interconnect** feature lets you create a network connection between two or more Azure VMware Solution private clouds located in the same region. It creates a routing link between the management and workload networks of the private clouds to enable network communication between the clouds.

You can connect a private cloud to multiple private clouds, and the connections are non-transitive. For example, if *private cloud 1* is connected to *private cloud 2*, and *private cloud 2* is connected to *private cloud 3*, private clouds 1 and 3 would not communicate until they were directly connected.

You can only connect private clouds in the same region. To connect private clouds that are in different regions, [use ExpressRoute Global Reach](#) to connect your private clouds in the same way you connect your private cloud to your on-premises circuit.

Supported regions

The Azure VMware Solution Interconnect feature is available in all regions.

Prerequisites

- Write access to each private cloud you're connecting
- Routed IP address space in each cloud is unique and doesn't overlap

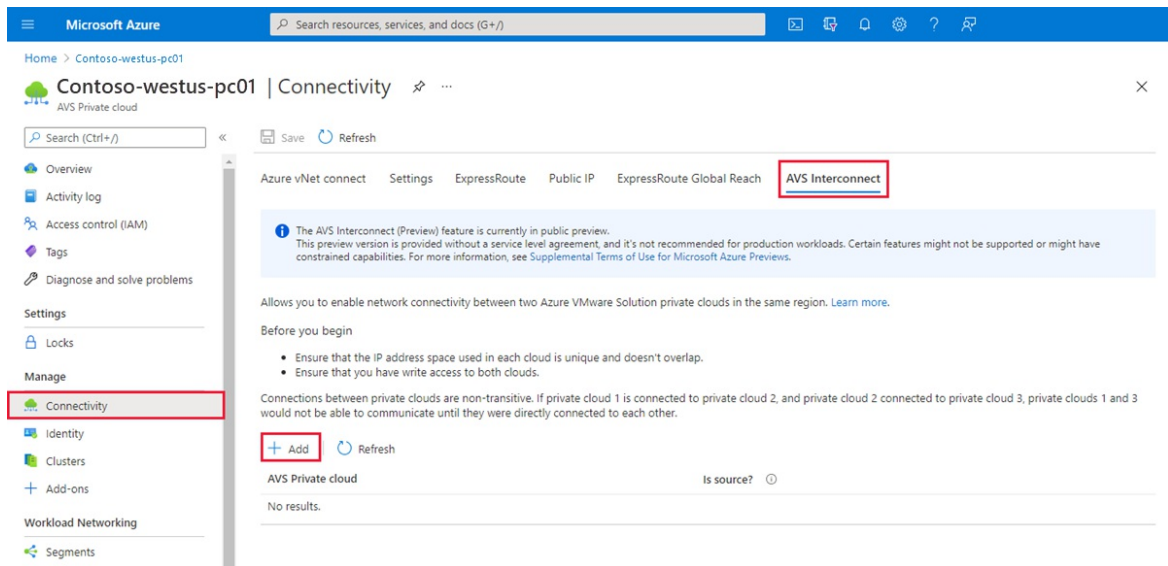
NOTE

The **AVS interconnect** feature doesn't check for overlapping IP space the way native Azure vNet peering does before creating the peering. Therefore, it's your responsibility to ensure that there isn't overlap between the private clouds.

In Azure VMware Solution environments, it's possible to configure non-routed, overlapping IP deployments on NSX segments that aren't routed to Azure. These don't cause issues with the AVS Interconnect feature, as it only routes between the NSX-T Data Center T0 gateway on each private cloud.

Add connection between private clouds

1. In your Azure VMware Solution private cloud, under **Manage**, select **Connectivity**.
2. Select the **AVS Interconnect** tab and then **Add**.



3. Select the information and Azure VMware Solution private cloud for the new connection.

NOTE

You can only connect to private clouds in the same region. To connect to private clouds that are in different regions, use [ExpressRoute Global Reach](#) to connect your private clouds in the same way you connect your private cloud to your on-premises circuit.

Add connection to other private cloud ✕

Subscription
Contoso ▼

Location ⓘ
(US) West US ▼
You can only connect to other private clouds in the same region.

Resource group
contoso-westus-rg ▼

AVS Private cloud * ⓘ
Contoso-westus-pc02 ▼

I confirm that the two private clouds to be connected don't contain overlapping network address space.

Create Cancel

4. Select the I confirm checkbox acknowledging that there are no overlapping routed IP spaces in the two private clouds.

5. Select **Create**. You can check the status of the connection creation.

Notifications



More events in the activity log →

Dismiss all ▾

- *** Save cloud connection Running ×
Building connection between Contoso-westus-pc01 and Contoso-westus-pc03.
21 minutes ago
- ✔ Save cloud connection ×
Successfully created a connection between Contoso-westus-pc01 and Contoso-westus-02.
2 minutes ago

You'll see all of your connections under **AVS Private Cloud**.

Home > Contoso-westus-pc01

Contoso-westus-pc01 | Connectivity ×

Search (Ctrl+/) Save Refresh

Azure vNet connect Settings ExpressRoute Public IP ExpressRoute Global Reach **AVS Interconnect**

i The AVS Interconnect (Preview) feature is currently in public preview. This preview version is provided without a service level agreement, and it's not recommended for production workloads. Certain features might not be supported or might have constrained capabilities. For more information, see [Supplemental Terms of Use for Microsoft Azure Previews](#).

Allows you to enable network connectivity between two Azure VMware Solution private clouds in the same region. [Learn more.](#)

Before you begin

- Ensure that the IP address space used in each cloud is unique and doesn't overlap.
- Ensure that you have write access to both clouds.

Connections between private clouds are non-transitive. If private cloud 1 is connected to private cloud 2, and private cloud 2 connected to private cloud 3, private clouds 1 and 3 would not be able to communicate until they were directly connected to each other.

+ Add | Refresh

AVS Private cloud	Is source?	
Contoso-westus-pc02	No	
Contoso-westus-pc03	No	

Remove connection between private clouds

- In your Azure VMware Solution private cloud, under **Manage**, select **Connectivity**.
- For the connection you want to remove, on the right, select **Delete** (trash can) and then **Yes**.

Next steps

Now that you've connected multiple private clouds in the same region, you may want to learn about:

- [Move Azure VMware Solution resources to another region](#)
- [Move Azure VMware Solution subscription to another subscription](#)

Deploy Arc for Azure VMware Solution (Preview)

12/16/2022 • 18 minutes to read • [Edit Online](#)

In this article, you'll learn how to deploy Arc for Azure VMware Solution. Once you've set up the components needed for this public preview, you'll be ready to execute operations in Azure VMware Solution vCenter Server from the Azure portal. Operations are related to Create, Read, Update, and Delete (CRUD) virtual machines (VMs) in an Arc-enabled Azure VMware Solution private cloud. Users can also enable guest management and install Azure extensions once the private cloud is Arc-enabled.

Before you begin checking off the prerequisites, verify the following actions have been done:

- You deployed an Azure VMware Solution private cluster.
- You have a connection to the Azure VMware Solution private cloud through your on-prem environment or your native Azure Virtual Network.
- There should be an isolated NSX-T Data Center segment for deploying the Arc for Azure VMware Solution Open Virtualization Appliance (OVA). If an isolated NSX-T Data Center segment doesn't exist, one will be created.

Prerequisites

The following items are needed to ensure you're set up to begin the onboarding process to deploy Arc for Azure VMware Solution (Preview).

- A jump box virtual machine (VM) with network access to the Azure VMware Solution vCenter.
 - From the jump-box VM, verify you have access to [vCenter Server and NSX-T Manager portals](#).
- Verify that your Azure subscription has been enabled or you have connectivity to Azure end points, mentioned in the [Appendices](#).
- Resource group in the subscription where you have owner or contributor role.
- A minimum of three free non-overlapping IPs addresses.
- Verify that your vCenter Server version is 6.7 or higher.
- A resource pool with minimum-free capacity of 16 GB of RAM, 4 vCPUs.
- A datastore with minimum 100 GB of free disk space that is available through the resource pool.
- On the vCenter Server, allow inbound connections on TCP port 443, so that the Arc resource bridge and VMware cluster extension can communicate with the vCenter server.

NOTE

Only the default port of 443 is supported. If you use a different port, Appliance VM creation will fail.

At this point, you should have already deployed an Azure VMware Solution private cloud. You need to have a connection from your on-prem environment or your native Azure Virtual Network to the Azure VMware Solution private cloud.

For Network planning and setup, use the [Network planning checklist - Azure VMware Solution | Microsoft Docs](#)

Registration to Arc for Azure VMware Solution feature set

The following **Register features** are for provider registration using Azure CLI.


```
az provider register --namespace Microsoft.ConnectedVMwarevSphere
az provider register --namespace Microsoft.ExtendedLocation
az provider register --namespace Microsoft.KubernetesConfiguration
az provider register --namespace Microsoft.ResourceConnector
az provider register --namespace Microsoft.AVS
```

Alternately, users can sign into their Subscription, navigate to the **Resource providers** tab, and register themselves on the resource providers mentioned previously.

For feature registration, users will need to sign into their **Subscription**, navigate to the **Preview features** tab, and search for 'Azure Arc for Azure VMware Solution'. Once registered, no other permissions are required for users to access Arc.

Users need to ensure they've registered themselves to **Microsoft.AVS/earlyAccess**. After registering, use the following feature to verify registration.

```
az feature show --name AzureArcForAVS --namespace Microsoft.AVS
```

Onboard process to deploy Azure Arc

Use the following steps to guide you through the process to onboard in Arc for Azure VMware Solution (Preview).

1. Sign into the jumpbox VM and extract the contents from the compressed file from the following [location](#). The extracted file contains the scripts to install the preview software.
2. Open the 'config_avs.json' file and populate all the variables.

Config JSON

```
{
  "subscriptionId": "",
  "resourceGroup": "",
  "applianceControlPlaneIpAddress": "",
  "privateCloud": "",
  "isStatic": true,
  "staticIpNetworkDetails": {
    "networkForApplianceVM": "",
    "networkCIDRForApplianceVM": "",
    "k8sNodeIPPoolStart": "",
    "k8sNodeIPPoolEnd": "",
    "gatewayIpAddress": ""
  }
}
```

- Populate the `subscriptionId`, `resourceGroup`, and `privateCloud` names respectively.
- `isStatic` is always true.
- `networkForApplianceVM` is the name for the segment for Arc appliance VM. One will be created if it doesn't already exist.
- `networkCIDRForApplianceVM` is the IP CIDR of the segment for Arc appliance VM. It should be unique and not affect other networks of Azure VMware Solution management IP CIDR.
- `GatewayIpAddress` is the gateway for the segment for Arc appliance VM.
- `applianceControlPlaneIpAddress` is the IP address for the Kubernetes API server that should be part of the segment IP CIDR provided. It shouldn't be part of the k8s node pool IP range.
- `k8sNodeIPPoolStart`, `k8sNodeIPPoolEnd` are the starting and ending IP of the pool of IPs to assign to

the appliance VM. Both need to be within the `networkCIDRForApplianceVM`.

- `k8sNodeIPPoolStart`, `k8sNodeIPPoolEnd`, `gatewayIPAddress`, `applianceControlPlaneIpAddress` are optional. You may choose to skip all the optional fields or provide values for all. If you choose not to provide the optional fields then you must use /28 address space for `networkCIDRForApplianceVM`

Json example

```
{
  "subscriptionId": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx",
  "resourceGroup": "test-rg",
  "privateCloud": "test-pc",
  "isStatic": true,
  "staticIpNetworkDetails": {
    "networkForApplianceVM": "arc-segment",
    "networkCIDRForApplianceVM": "10.14.10.1/28"
  }
}
```

3. Run the installation scripts. We've provided you with the option to set up this preview from a Windows or Linux-based jump box/VM.

Run the following commands to execute the installation script.

- [Windows based jump box/VM](#)
- [Linux based jump box/VM](#)

Script isn't signed so we need to bypass Execution Policy in PowerShell. Run the following commands.

```
Set-ExecutionPolicy -Scope Process -ExecutionPolicy Bypass; .\run.ps1 -Operation onboard -FilePath {config-json-path}
```

4. You'll notice more Azure Resources have been created in your resource group.

- Resource bridge
- Custom location
- VMware vCenter

IMPORTANT

You can't create the resources in a separate resource group. Make sure you use the same resource group from where the Azure VMware Solution private cloud was created to create the resources.

Discover and project your VMware infrastructure resources to Azure

When Arc appliance is successfully deployed on your private cloud, you can do the following actions.

- View the status from within the private cloud under **Operations > Azure Arc**, located in the left navigation.
- View the VMware vSphere infrastructure resources from the private cloud left navigation under **Private cloud** then select **Azure Arc vCenter resources**.
- Discover your VMware vSphere infrastructure resources and project them to Azure using the same browser experience, **Private cloud > Arc vCenter resources > Virtual Machines**.
- Similar to VMs, customers can enable networks, templates, resource pools, and data-stores in Azure.

After you've enabled VMs to be managed from Azure, you can install guest management and do the following actions.

- Enable customers to install and use extensions.
 - To enable guest management, customers will be required to use admin credentials
 - VMtools should already be running on the VM

NOTE

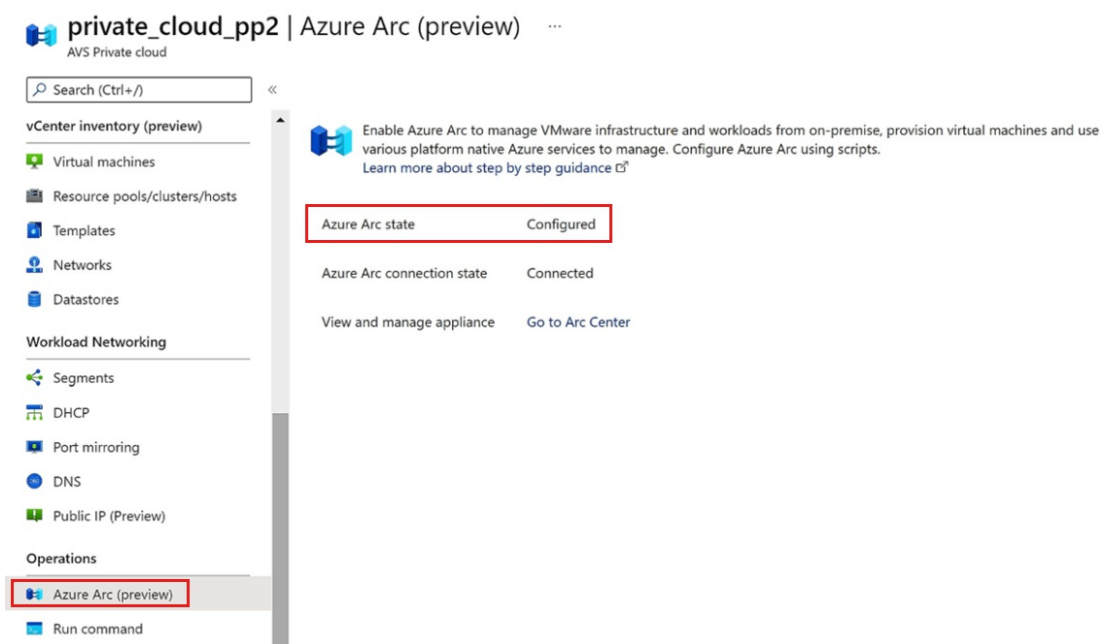
Azure VMware Solution vCenter Server will be available in global search but will NOT be available in the list of vCenter Servers for Arc for VMware.

- Customers can view the list of VM extensions available in public preview.
 - Change tracking
 - Log analytics
 - Azure policy guest configuration

Azure VMware Solution private cloud with Azure Arc

When the script has run successfully, you can check the status to see if Azure Arc has been configured. To verify if your private cloud is Arc-enabled, do the following action:

- In the left navigation, locate **Operations**.
- Choose **Azure Arc (preview)**. Azure Arc state will show as **Configured**.



Arc enabled VMware vSphere resources

After the private cloud is Arc-enabled, vCenter resources should appear under **Virtual machines**.

- From the left navigation, under **Azure Arc VMware resources (preview)**, locate **Virtual machines**.
- Choose **Virtual machines** to view the vCenter Server resources.

Manage access to VMware resources through Azure Role-Based Access Control

After your Azure VMware Solution vCenter resources have been enabled for access through Azure, there's one final step in setting up a self-service experience for your teams. You'll need to provide your teams with access to: compute, storage, networking, and other vCenter Server resources used to configure VMs.

This section will demonstrate how to use custom roles to manage granular access to VMware vSphere resources through Azure.

Arc-enabled VMware vSphere custom roles

Three custom roles are provided to meet your Role-based access control (RBAC) requirements. These roles can be applied to a whole subscription, resource group, or a single resource.

- Azure Arc VMware vSphere Administrator role
- Azure Arc VMware vSphere Private Cloud User role
- Azure Arc VMware vSphere VM Contributor role

The first role is for an Administrator. The other two roles apply to anyone who needs to deploy or manage a VM.

Azure Arc Azure VMware Solution Administrator role

This custom role gives the user permission to conduct all possible operations for the

`Microsoft.ConnectedVMwarevSphere` resource provider. This role should be assigned to users or groups who are administrators that manage Azure Arc-enabled Azure VMware Solution deployment.

Azure Arc Azure VMware Solution Private Cloud User role

This custom role gives the user permission to use the Arc-enabled Azure VMware Solutions vSphere resources that have been made accessible through Azure. This role should be assigned to any users or groups that need to deploy, update, or delete VMs.

We recommend assigning this role at the individual resource pool (host or cluster), virtual network, or template that you want the user to deploy VMs with.

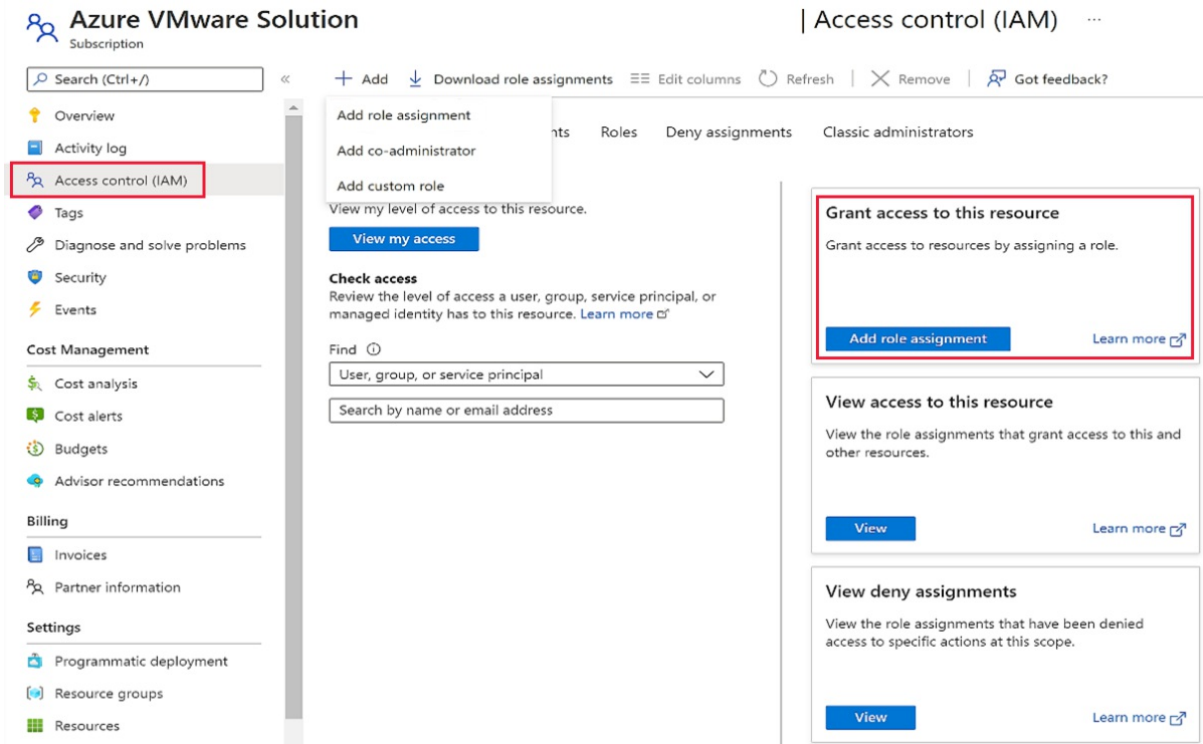
Azure Arc Azure VMware Solution VM Contributor role

This custom role gives the user permission to perform all VMware VM operations. This role should be assigned to any users or groups that need to deploy, update, or delete VMs.

We recommend assigning this role at the subscription level or resource group you want the user to deploy VMs with.

Assign custom roles to users or groups

1. Navigate to the Azure portal.
2. Locate the subscription, resource group, or the resource at the scope you want to provide for the custom role.
3. Find the Arc-enabled Azure VMware Solution vCenter Server resources.
 - a. Navigate to the resource group and select the **Show hidden types** checkbox.
 - b. Search for "Azure VMware Solution".
4. Select **Access control (IAM)** in the table of contents located on the left navigation.
5. Select **Add role assignment** from the **Grant access to this resource**.



6. Select the custom role you want to assign, Azure Arc VMware Solution: **Administrator**, **Private Cloud User**, or **VM Contributor**.
7. Search for **AAD user** or **group name** that you want to assign this role to.
8. Select the **AAD user** or **group name**. Repeat this step for each user or group you want to give permission to.
9. Repeat the above steps for each scope and role.

Create Arc-enabled Azure VMware Solution virtual machine

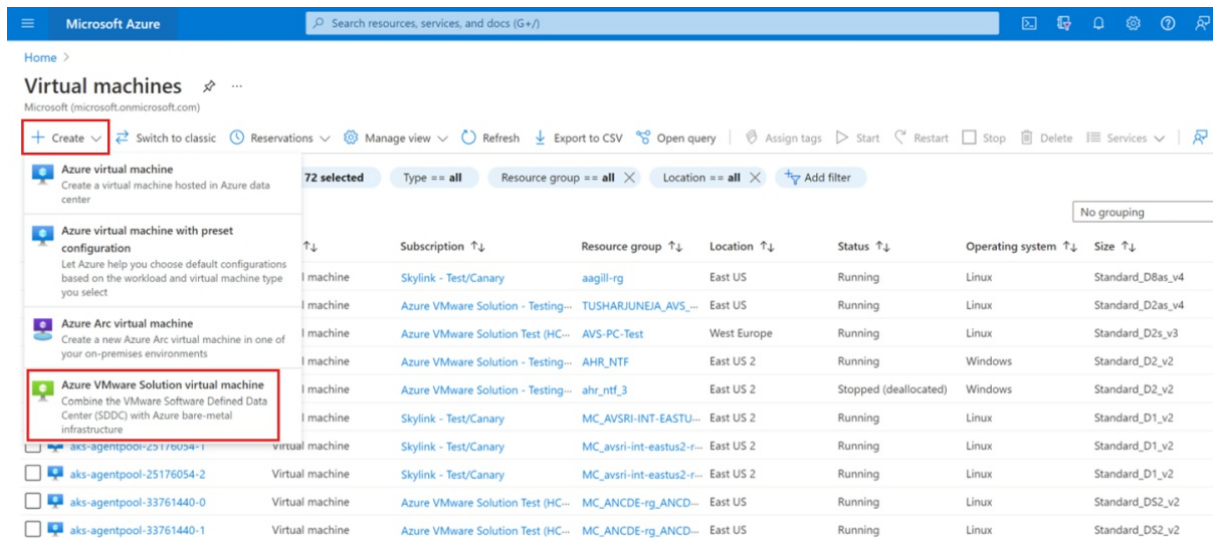
This section shows users how to create a virtual machine (VM) on VMware vCenter Server using Azure Arc. Before you begin, check the following prerequisite list to ensure you're set up and ready to create an Arc-enabled Azure VMware Solution VM.

Prerequisites

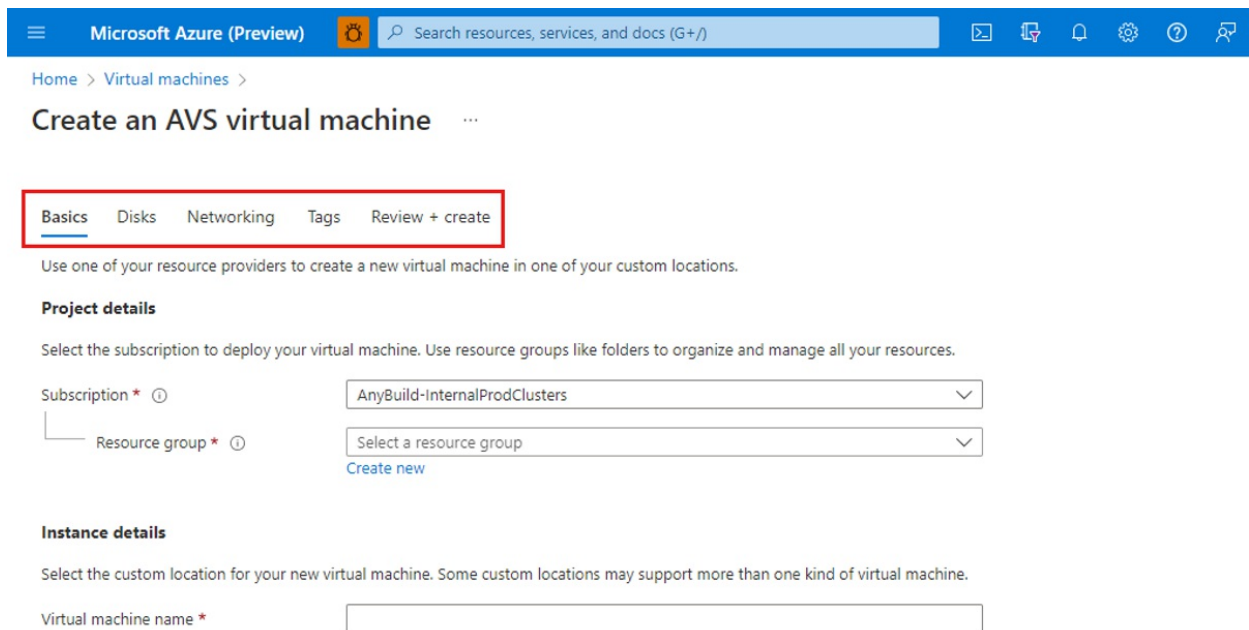
- An Azure subscription and resource group where you have an Arc VMware VM **Contributor** role.
- A resource pool resource that you have an Arc VMware private cloud resource **User** role.
- A virtual machine template resource that you have an Arc private cloud resource **User** role.
- (Optional) a virtual network resource on which you have Arc private cloud resource **User** role.

Create VM flow

- Open the [Azure portal](#)
- On the **Home** page, search for **virtual machines**. Once you've navigated to **Virtual machines**, select the **+** **Create** drop down and select **Azure VMware Solution virtual machine**.



Near the top of the **Virtual machines** page, you'll find five tabs labeled: **Basics**, **Disks**, **Networking**, **Tags**, and **Review + create**. Follow the steps or options provided in each tab to create your Azure VMware Solution virtual machine.



Basics

1. In **Project details**, select the **Subscription** and **Resource group** where you want to deploy your VM.
2. In **Instance details**, provide the **virtual machine name**.
3. Select a **Custom location** that your administrator has shared with you.
4. Select the **Resource pool/cluster/host** where the VM should be deployed.
5. For **Template details**, pick a **Template** based on the VM you plan to create.
 - Alternately, you can check the **Override template defaults** box that allows you to override the CPU and memory specifications set in the template.
 - If you chose a Windows template, you can provide a **Username** and **Password** for the **Administrator account**.
6. For **Extension setup**, the box is checked by default to **Enable guest management**. If you don't want guest management enabled, uncheck the box.
7. The connectivity method defaults to **Public endpoint**. Create a **Username**, **Password**, and **Confirm password**.

Disks

- You can opt to change the disks configured in the template, add more disks, or update existing disks. These disks will be created on the default datastore per the VMware vCenter Server storage policies.
- You can change the network interfaces configured in the template, add Network interface cards (NICs), or update existing NICs. You can also change the network that the NIC will be attached to provided you have permissions to the network resource.

Networking

- A network configuration is automatically created for you. You can choose to keep it or override it and add a new network interface instead.
- To override the network configuration, find and select **+ Add network interface** and add a new network interface.

Tags

- In this section, you can add tags to the VM resource.

Review + create

- Review the data and properties you've set up for your VM. When everything is set up how you want it, select **Create**. The VM should be created in a few minutes.

Enable guest management and extension installation

The guest management must be enabled on the VMware vSphere virtual machine (VM) before you can install an extension. Use the following prerequisite steps to enable guest management.

Prerequisite

1. Navigate to [Azure portal](#).
2. Locate the VMware vSphere VM you want to check for guest management and install extensions on, select the name of the VM.
3. Select **Configuration** from the left navigation for a VMware VM.
4. Verify **Enable guest management** has been checked.

NOTE

The following conditions are necessary to enable guest management on a VM.

- The machine must be running a [Supported operating system](#).
- The machine needs to connect through the firewall to communicate over the internet. Make sure the [URLs](#) listed aren't blocked.
- The machine can't be behind a proxy, it's not supported yet.
- If you're using Linux VM, the account must not prompt to sign in on pseudo commands.

Avoid pseudo commands by following these steps:

1. Sign into Linux VM.
2. Open terminal and run the following command: `sudo visudo`.
3. Add the line `username ALL=(ALL) NOPASSWD:ALL` at the end of the file.
4. Replace `username` with the appropriate user-name.

If your VM template already has these changes incorporated, you won't need to do the steps for the VM created from that template.

Extension installation steps

1. Go to Azure portal.
2. Find the Arc-enabled Azure VMware Solution VM that you want to install an extension on and select the VM name.
3. Navigate to **Extensions** in the left navigation, select **Add**.
4. Select the extension you want to install.
 - a. Based on the extension, you'll need to provide details. For example, `workspace Id` and `key` for LogAnalytics extension.
5. When you're done, select **Review + create**.

When the extension installation steps are completed, they trigger deployment and install the selected extension on the VM.

Change Arc appliance credential

When `cloudadmin` credentials are updated, use the following steps to update the credentials in the appliance store.

1. Log into the jumpbox VM from where onboarding was performed. Change the directory to **onboarding directory**.
2. Run the following command for Windows-based jumpbox VM.

```
./temp/.env/Scripts/activate
```

3. Run the following command.

```
az arcappliance update-infracredentials vmware --kubeconfig <kubeconfig file>
```

4. Run the following command

```
az connectedvmware vcenter connect --debug --resource-group {resource-group} --name {vcenter-name-in-azure} --location {vcenter-location-in-azure} --custom-location {custom-location-name} --fqdn {vcenter-ip} --port {vcenter-port} --username cloudadmin@vsphere.local --password {vcenter-password}
```

NOTE

Customers need to ensure kubeconfig and SSH keys remain available as they will be required for log collection, appliance Upgrade, and credential rotation. These parameters will be required at the time of upgrade, log collection, and credential update scenarios.

Parameters

Required parameters

```
-kubeconfig # kubeconfig of Appliance resource
```

Examples

The following command invokes the set credential for the specified appliance resource.

```
az arcappliance setcredential <provider> --kubeconfig <kubeconfig>
```

Manual appliance upgrade

Use the following steps to perform a manual upgrade for Arc appliance virtual machine (VM).

1. Log into vCenter Server.

2. Locate the Arc appliance VM, which should be in the resource pool that was configured during onboarding.
 - a. Power off the VM.
 - b. Delete the VM.
3. Delete the download template corresponding to the VM.
4. Delete the resource bridge ARM resource.
5. Get the previous script `config_avs` file and add the following configuration item:
 - a. `"register":false`
6. Download the latest version of the Azure VMware Solution onboarding script.
7. Run the new onboarding script with the previous `config_avs.json` from the jump box VM, without changing other config items.

Off board from Azure Arc-enabled Azure VMware Solution

This section demonstrates how to remove your VMware vSphere virtual machines (VMs) from Azure management services.

If you've enabled guest management on your Arc-enabled Azure VMware Solution VMs and onboarded them to Azure management services by installing VM extensions on them, you'll need to uninstall the extensions to prevent continued billing. For example, if you installed an MMA extension to collect and send logs to an Azure Log Analytics workspace, you'll need to uninstall that extension. You'll also need to uninstall the Azure Connected Machine agent to avoid any problems installing the agent in future.

Use the following steps to uninstall extensions from the portal.

NOTE

Steps 2-5 must be performed for all the VMs that have VM extensions installed.

1. Log into your Azure VMware Solution private cloud.
2. Select **Virtual machines** in **Private cloud**, found in the left navigation under "Arc-enabled VMware resources".
3. Search and select the virtual machine where you have **Guest management** enabled.
4. Select **Extensions**.
5. Select the extensions and select **Uninstall**.

To avoid problems onboarding the same VM to **Guest management**, we recommend you do the following steps to cleanly disable guest management capabilities.

NOTE

Steps 2-3 must be performed for all VMs that have **Guest management** enabled.

1. Sign into the virtual machine using administrator or root credentials and run the following command in the shell.
 - a. `azcmagent disconnect --force-local-only`.
2. Uninstall the `ConnectedMachine agent` from the machine.
3. Set the **identity** on the VM resource to **none**.

Remove Arc-enabled Azure VMware Solution vSphere resources from Azure

When you activate Arc-enabled Azure VMware Solution resources in Azure, a representation is created for them in Azure. Before you can delete the vCenter Server resource in Azure, you'll need to delete all of the Azure resource representations you created for your vSphere resources. To delete the Azure resource representations you created, do the following steps:

1. Go to the Azure portal.
2. Choose **Virtual machines** from Arc-enabled VMware vSphere resources in the private cloud.
3. Select all the VMs that have an Azure Enabled value as **Yes**.
4. Select **Remove from Azure**. This step will start deployment and remove these resources from Azure. The resources will remain in your vCenter Server.
 - a. Repeat steps 2, 3 and 4 for **Resourcepools/clusters/hosts, Templates, Networks, and Datastores**.
5. When the deletion completes, select **Overview**.
 - a. Note the Custom location and the Azure Arc Resource bridge resources in the Essentials section.
6. Select **Remove from Azure** to remove the vCenter resource from Azure.
7. Go to vCenter Server resource in Azure and delete it.
8. Go to the Custom location resource and select **Delete**.
9. Go to the Azure Arc Resource bridge resources and select **Delete**.

At this point, all of your Arc-enabled VMware vSphere resources have been removed from Azure.

Delete Arc resources from vCenter Server

For the final step, you'll need to delete the resource bridge VM and the VM template that were created during the onboarding process. Once that step is done, Arc won't work on the Azure VMware Solution SDDC. When you delete Arc resources from vCenter, it won't affect the Azure VMware Solution private cloud for the customer.

Preview FAQ

Is Arc supported in all the Azure VMware Solution regions?

Arc is supported in EastUS and WestEU regions however we are working to extend the regional support.

How does support work?

Standard support process for Azure VMware Solution has been enabled to support customers.

Does Arc for Azure VMware Solution support private end point?

Yes. Arc for Azure VMware Solution will support private end point for general audience. However, it's not currently supported.

Is enabling internet the only option to enable Arc for Azure VMware Solution?

Yes

Is DHCP support available?

DHCP support is not available to customers at this time, we only support static IP.

NOTE

This is Azure VMware Solution 2.0 only. It's not available for Azure VMware Solution by Cloudsimple.

Debugging tips for known issues

Use the following tips as a self-help guide.

What happens if I face an error related to Azure CLI?

- For windows jumpbox, if you have 32-bit Azure CLI installed, verify that your current version of Azure CLI has been uninstalled. Verification can be done from the Control Panel.
- To ensure it's uninstalled, try the `az` version to check if it's still installed.
- If you already installed Azure CLI using MSI, `az` installed by MSI and pip will conflict on PATH. In this case, it's recommended that you uninstall the current Azure CLI version.

My script stopped because it timed-out, what should I do?

- Retry the script for `create`. A prompt will ask you to select Y and rerun it.
- It could be a cluster extension issue that would result in adding the extension in the pending state.
- Verify you have the correct script version.
- Verify the VMware pod is running correctly on the system in running state.

Basic trouble-shooting steps if the script run was unsuccessful.

- Follow the directions provided in the [Prerequisites](#) section of this article to verify that the feature and resource providers are registered.

What happens if the Arc for VMware section shows no data?

- If the Azure Arc VMware resources in the Azure UI show no data, verify your subscription was added in the global default subscription filter.

I see the error: "`ApplianceClusterNotRunning` Appliance Cluster: `<resource-bridge-id>` expected states to be Succeeded found: Succeeded and expected status to be Running and found: Connected".

- Run the script again.

I'm unable to install extensions on my virtual machine.

- Check that **guest management** has been successfully installed.
- **VMtools** should be installed on the VM.

I'm facing Network related issues during on-boarding.

- Look for an IP conflict. You need IPs with no conflict or from free pool.
- Verify the internet is enabled for the network segment.

Where can I find more information related to Azure Arc resource bridge?

- For more information, go to [Azure Arc resource bridge \(preview\) overview](#)

Appendices

Appendix 1 shows proxy URLs required by the Azure Arc-enabled private cloud. The URLs will get pre-fixed when the script runs and can be run from the jumpbox VM to ping them.

AZURE ARC SERVICE	URL
Microsoft container registry	<code>https://mcr.microsoft.com</code>
Azure Arc Identity service	<code>https://*.his.arc.azure.com</code>

AZURE ARC SERVICE	URL
Azure Arc configuration service	<code>https://*.dp.kubernetesconfiguration.azure.com</code>
Cluster connect	<code>https://*.servicebus.windows.net</code>
Guest Notification service	<code>https://guestnotificationsservice.azure.com</code>
Resource bridge (appliance) Dataplate service	<code>https://*.dp.prod.appliances.azure.com</code>
Resource bridge (appliance) container image download	<code>https://ecpacr.azurecr.io</code>
Resource bridge (appliance) image download	<code>https://.blob.core.windows.net</code> <code>https://*.dl.delivery.mp.microsoft.com</code> <code>https://*.do.dsp.mp.microsoft.com</code>
Azure Resource Manager	<code>https://management.azure.com</code>
Azure Active Directory	<code>https://login.microsoftonline.com</code>

Additional URL resources

- [Google Container Registry](#)
- [Red Hat Quay.io](#)

Deploy disaster recovery using JetStream DR software

12/16/2022 • 14 minutes to read • [Edit Online](#)

[JetStream DR](#) is a cloud-native disaster recovery solution designed to minimize downtime of virtual machines (VMs) if there is a disaster. Instances of JetStream DR are deployed at both the protected and recovery sites.

JetStream is built on the foundation of Continuous Data Protection (CDP), using [VMware vSphere API for I/O filtering \(VAIO\) framework](#), which enables minimal or close to no data loss. JetStream DR provides the level of protection wanted for business and mission-critical applications. It also enables cost-effective DR by using minimal resources at the DR site and using cost-effective cloud storage, such as [Azure Blob Storage](#).

In this article, you'll implement JetStream DR for your Azure VMware Solution private cloud and on-premises VMware workloads.

To learn more about JetStream DR, see:

- [JetStream Solution brief](#)
- [JetStream DR on Azure Marketplace](#)
- [JetStream knowledge base articles](#)

Core components of the JetStream DR solution

ITEMS	DESCRIPTION
JetStream Management Server Virtual Appliance (MSA)	MSA enables both Day 0 and Day 2 configuration, such as primary sites, protection domains, and recovering VMs. The MSA is deployed from an OVA on a vSphere node by the cloud admin. The MSA collects and maintains statistics relevant to VM protection and implements a vCenter plugin that allows you to manage JetStream DR natively with the vSphere Client. The MSA doesn't handle replication data of protected VMs.
JetStream DR Virtual Appliance (DRVA)	Linux-based Virtual Machine appliance receives protected VMs replication data from the source ESXi host. It maintains the replication log and manages the transfer of the VMs and their data to the object store such as Azure Blob Storage. Depending on the number of protected VMs and the amount of VM data to replicate, the private cloud admin can create one or more DRVA instances.
JetStream ESXi host components (IO Filter packages)	JetStream software installed on each ESXi host configured for JetStream DR. The host driver intercepts the vSphere VMs IO and sends the replication data to the DRVA. The IO filters also monitor relevant events, such as vMotion, Storage vMotion, snapshots, etc.

ITEMS	DESCRIPTION
JetStream Protected Domain	Logical group of VMs that will be protected together using the same policies and runbook. The data for all VMs in a protection domain is stored in the same Azure Blob container instance. A single DRVA instance handles replication to remote DR storage for all VMs in a Protected Domain.
Azure Blob Storage containers	The protected VMs replicated data is stored in Azure Blobs. JetStream software creates one Azure Blob container instance for each JetStream Protected Domain.

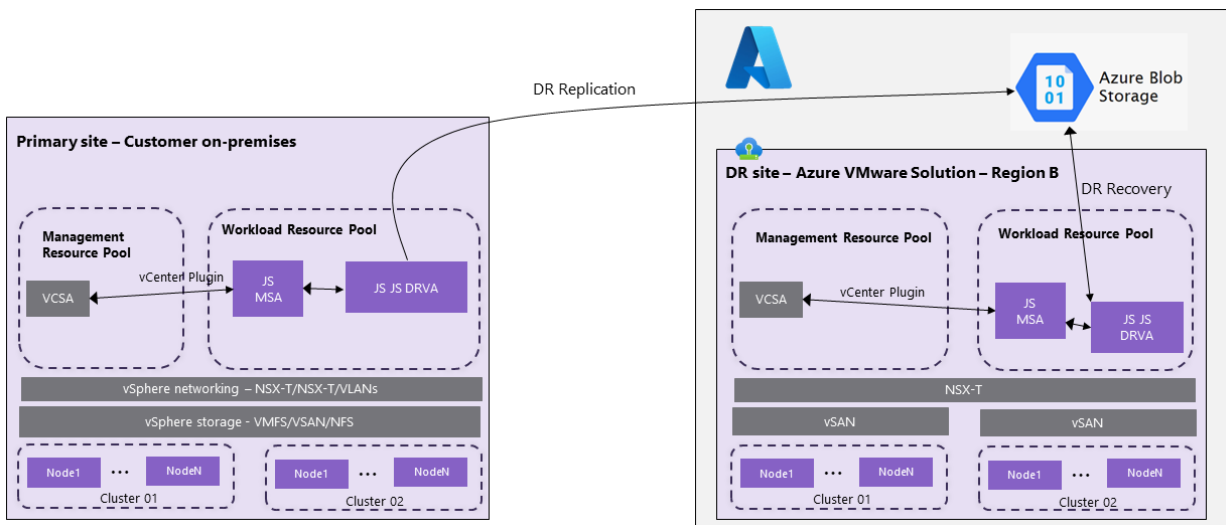
JetStream scenarios on Azure VMware Solution

You can use JetStream DR with Azure VMware Solution for the following two scenarios:

- On-premises VMware vSphere to Azure VMware Solution DR
- Azure VMware Solution to Azure VMware Solution DR

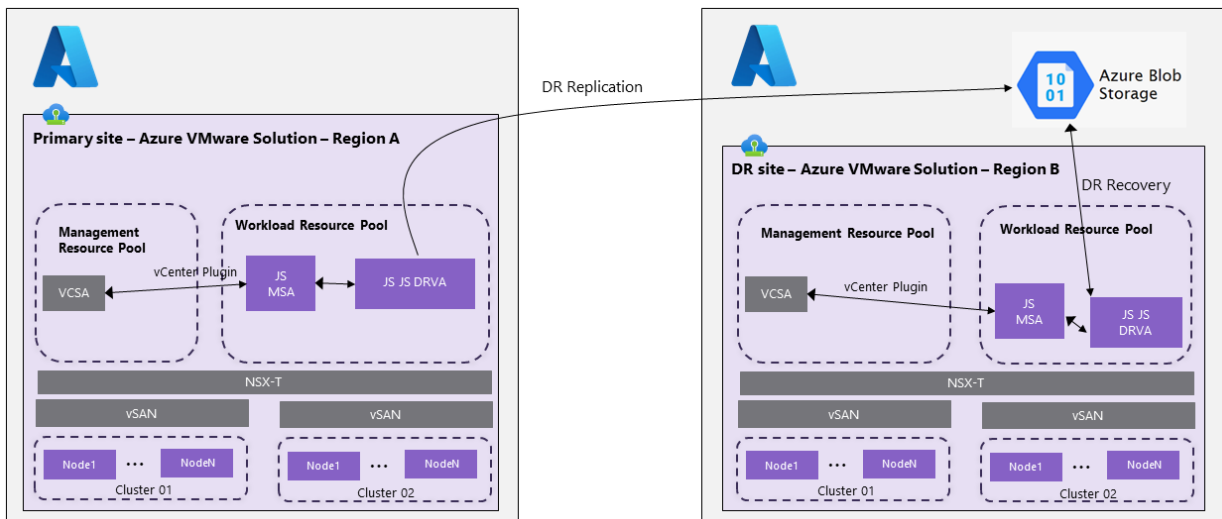
Scenario 1: On-premises VMware vSphere to Azure VMware Solution DR

In this scenario, the primary site is your on-premises VMware vSphere environment and the DR site is an Azure VMware Solution private cloud.



Scenario 2: Azure VMware Solution to Azure VMware Solution DR

In this scenario, the primary site is an Azure VMware Solution private cloud in one Azure region. The disaster recovery site is an Azure VMware Solution private cloud in a different Azure region.



Disaster Recovery with Azure NetApp Files, JetStream DR and Azure VMware Solution

Disaster Recovery to cloud is a resilient and cost-effective way of protecting the workloads against site outages and data corruption events like ransomware. Leveraging the VMware VAIO framework, on-premises VMware workloads can be replicated to Azure Blob storage and recovered with minimal or close to no data loss and near-zero Recovery Time Objective (RTO). JetStream DR can seamlessly recover workloads replicated from on-premises to Azure VMware Solution and specifically to Azure NetApp Files.

JetStream DR enables cost-effective disaster recovery by consuming minimal resources at the DR site and using cost-effective cloud storage. JetStream DR automates recovery to Azure NetApp Files (ANF) datastores using Azure Blob Storage. It can recover independent VMs or groups of related VMs into the recovery site infrastructure according to runbook settings. It also provides point-in-time recovery for ransomware protection.

Install JetStream DR

To install JetStream DR in the on-premises data center and in the Azure VMware Solution private cloud:

- Install JetStream DR in the on-premises data center:
 - Download the JetStream DR bundle from Azure Marketplace (ZIP) and deploy the JetStream DR MSA (OVA) in the designated cluster.
 - Configure the cluster with the IO filter package (install JetStream VIB).
 - Provision Azure Blob (Azure Storage Account) in the same region as the DR Azure VMware Solution cluster.
 - Deploy the disaster recovery virtual appliance (DRVA) and assign a replication log volume (VMDK from existing datastore or shared iSCSI storage).
 - Create Protected Domains (groups of related VMs) and assign DRVAs and the Azure Blob Storage/ANF.
 - Start protection.
- Install JetStream DR in the Azure VMware Solution private cloud:
 - Use the Run command to install and configure JetStream DR.
 - Add the same Azure Blob container and discover domains using the Scan Domain option.
 - Deploy the DRVA appliance.
 - Create a replication log volume using an available vSAN or ANF datastore.
 - Import protected domains and configure RocVA (recovery VA) to use ANF datastore for VM placements.
 - Select the appropriate failover option and start continuous rehydration for near-zero RTO

domains/VMs.

- During a disaster event, trigger failover to Azure NetApp Files datastores in the designated Azure VMware Solution DR site.
- Invoke failback to the protected site after the protected site has been recovered.

Ransomware recovery

Recovering from ransomware can be a daunting task. Specifically, it can be hard for IT organizations to pinpoint what the “safe point of return is”. After that safe point is determined, how to ensure that recovered workloads are safeguarded from the attacks re-occurring by sleeping malware or through vulnerable applications.

JetStream DR for Azure VMware Solution together with Azure NetApp Files datastores can address these concerns by allowing organizations to recover from an available point-in-time. It ensures that workloads are recovered to a functional and isolated network if required, allows the applications to function and communicate with each other without exposing them to any North-South traffic. It also gives security teams a safe place to perform forensics, and conduct other recovery measures.

For full details, refer to the article: [Disaster Recovery with Azure NetApp Files, JetStream DR and Azure VMware Solution](#).

Prerequisites

Scenario 1: On-premises VMware vSphere to Azure VMware Solution DR

- Azure VMware Solution private cloud deployed with a minimum of three nodes in the target DR region.

Azure VMware Solution

Microsoft

+ Create Manage view Refresh Export to CSV Open query Assign tags Feedback

Filter for any field... Subscription == all Resource group == all Location == all Add filter

Showing 1 to 2 of 2 records.

<input type="checkbox"/>	Name ↑↓	Type ↑↓	Resource group ↑↓	Location ↑↓
<input type="checkbox"/>	AVS_EUW_001	AVS Private cloud	AVS_EUW_001	West Europe
<input type="checkbox"/>	AVS_UKS_001	AVS Private cloud	AVS_UKS_001	UK South

- Network connectivity configured between the primary site JetStream appliances and the Azure Storage blob instance.
- [Setup and Subscribe to JetStream DR](#) from the Azure Marketplace to download the JetStream DR software.
- [Azure Blob Storage account](#) created using either Standard or Premium Performance tier. For [access tier](#), select **Hot**.

NOTE

The **Enable hierarchical namespace** option on the blob isn't supported.

- An NSX-T network segment configured on Azure VMware Solution private cloud with DHCP enabled on the segment for the transient JetStream Virtual appliances employed during recovery or failover.
- A DNS server configured to resolve the IP addresses of Azure VMware Solution vCenter Server, Azure VMware Solution ESXi hosts, Azure Storage account, and the JetStream Marketplace service for the JetStream virtual appliances.
- (Optional) Azure NetApp Files volume(s) are created and attached to the Azure VMware Solution private

cloud for recovery or failover of protected VMs to Azure NetApp Files backed datastores.

- [Attach Azure NetApp Files datastores to Azure VMware Solution hosts](#)
- [Disaster Recovery with Azure NetApp Files, JetStream DR and AVS \(Azure VMware Solution\)](#)

Scenario 2: Azure VMware Solution to Azure VMware Solution DR

- Azure VMware Solution private cloud deployed with a minimum of three nodes in both the primary and secondary regions.
- Network connectivity configured between the primary site JetStream appliances and the Azure Storage blob instance.
- [Setup and Subscribe to JetStream DR](#) from the Azure Marketplace to download the JetStream DR software.
- [Azure Blob Storage account](#) created using either Standard or Premium Performance tier. For [access tier](#), select **Hot**.

NOTE

The **Enable hierarchical namespace** option on the blob isn't supported.

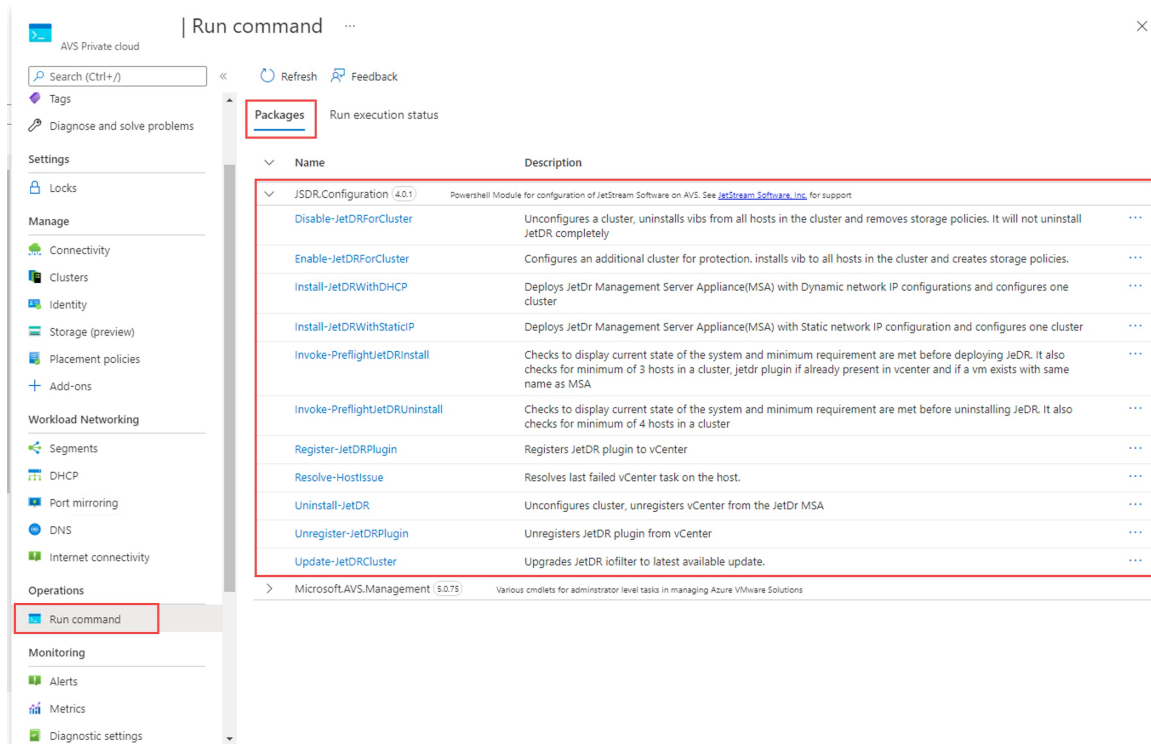
- An NSX-T network segment configured on Azure VMware Solution private cloud with DHCP enabled on the segment for the transient JetStream Virtual appliances employed during recovery or failover.
- DNS configured on both the primary and DR sites to resolve the IP addresses of Azure VMware Solution vCenter Server, Azure VMware Solution ESXi hosts, Azure Storage account, the JetStream DR Management Server Appliance (MSA) and the JetStream Marketplace service for the JetStream virtual appliances.
- (Optional) Azure NetApp Files volume(s) are created and attached to the Azure VMware Solution private cloud for recovery or failover of protected VMs to Azure NetApp Files backed datastores.
 - [Attach Azure NetApp Files datastores to Azure VMware Solution hosts](#)
 - [Disaster Recovery with Azure NetApp Files, JetStream DR and AVS \(Azure VMware Solution\)](#)

For more on-premises JetStream DR prerequisites, see the [JetStream Pre-Installation Guide](#).

Install JetStream DR on Azure VMware Solution

You can follow these steps for both supported scenarios.

1. In your on-premises data center, install JetStream DR following the [JetStream documentation](#).
2. In your Azure VMware Solution private cloud, install JetStream DR using a Run command. From the [Azure portal](#), select **Run command** > **Packages** > **JSDR.Configuration**.



NOTE

The default CloudAdmin user in Azure VMware Solution doesn't have sufficient privileges to install JetStream DR. Azure VMware Solution enables simplified and automated installation of JetStream DR by invoking the Azure VMware Solution Run command for JetStream DR.

3. Run the **Invoke-PreflightJetDRInstall** cmdlet, which checks if the prerequisites for installing JetStream DR have been met. For example, it validates the required number of hosts, cluster names, and unique VM names.
4. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
Network	Name of the NSX-T Data Center network segment where you must deploy the JetStream MSA.
Datastore	Name of the datastore where you will deploy the JetStream MSA.
ProtectedCluster	Name of the Azure VMware Solution private cloud cluster to be protected, for example, Cluster-1 . You can only provide one cluster name.
Cluster	Name of the Azure VMware Solution private cluster where the JetStream MSA will be deployed, for example, Cluster-1 .
VMName	Name of JetStream MSA VM, for example, jetstreamServer .
Specify name for execution	Alphanumeric name of the execution, for example, Invoke-PreflightJetDRInstall-Exec1 . It's used to verify if the cmdlet ran successfully.

FIELD	VALUE
Timeout	The period after which a cmdlet exits if taking too long to finish.

5. [View the status of the execution.](#)

Install the JetStream DR MSA

Azure VMware Solution supports the installation of JetStream using either static IP addresses or using DHCP-based IP addresses.

Static IP address

1. Select **Run command** > **Packages** > **Install-JetDRWithStaticIP**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
ProtectedCluster	Name of the Azure VMware Solution private cloud cluster to be protected, for example, Cluster-1 . You can only provide one cluster name during the install.
Datastore	Name of the datastore where the JetStream MSA will be deployed.
VMName	Name of JetStream MSA VM, for example, jetstreamServer .
Cluster	Name of the Azure VMware Solution private cluster where the JetStream MSA will be deployed, for example, Cluster-1 .
Netmask	Netmask of the MSA to be deployed, for example, 255.255.255.0 .
MSIp	IP address of the JetStream MSA VM.
Dns	DNS IP that the JetStream MSA VM should use.
Gateway	IP address of the network gateway for the JetStream MSA VM.
Credential	Credentials of the root user of the JetStream MSA VM.
HostName	Hostname (FQDN) of the JetStream MSA VM.
Network	Name of the NSX-T Data Center network segment where the JetStream MSA will be deployed.
Specify name for execution	Alphanumeric name of the execution, for example, Install-JetDRWithStaticIP-Exec1 . It's used to verify if the cmdlet ran successfully and should be unique for each run.

3. [View the status of the execution.](#)

DHCP-based IP address

This step also installs JetStream vSphere Installation Bundle (VIB) on the clusters that need DR protection.

1. Select **Run command** > **Packages** > **Install-JetDRWithDHCP**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
ProtectedCluster	Name of the Azure VMware Solution private cloud cluster to be protected, for example, Cluster-1 . You can only provide one cluster name during the install.
Datastore	Name of the datastore where the JetStream MSA will be deployed.
VMName	Name of JetStream MSA VM, for example, jetstreamServer .
Cluster	Name of the Azure VMware Solution private cluster where the JetStream MSA will be deployed, for example, Cluster-1 .
Credential	Credentials of the root user of the JetStream MSA VM.
HostName	Hostname (FQDN) of the JetStream MSA VM.
Network	Name of the NSX-T Data Center network segment where the JetStream MSA will be deployed.
Specify name for execution	Alphanumeric name of the execution, for example, Install-JetDRWithDHCP-Exec1 . It's used to verify if the cmdlet ran successfully and should be unique for each run.

3. [View the status of the execution](#).

Add JetStream DR to new Azure VMware Solution clusters

1. Select **Run command** > **Packages** > **Enable-JetDRForCluster**.
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
ProtectedCluster	Name of the Azure VMware Solution private cloud cluster to be protected, for example, Cluster-1 . You can only provide one cluster name during the install.
Credential	Credentials of the root user of the JetStream MSA VM.
MSIP	IP address of the JetStream MSA VM.
Specify name for execution	Alphanumeric name of the execution, for example, Enable-JetDRForCluster-Exec1 . It's used to verify if the cmdlet ran successfully and should be unique for each run.

3. [View the status of the execution.](#)

Configure JetStream DR

This section only covers an overview of the steps required for configuring JetStream DR. For detailed descriptions and steps, see the [Configuring JetStream DR](#) documentation.

Once JetStream DR MSA and JetStream VIB are installed on the Azure VMware Solution clusters, use the JetStream portal to complete the remaining configuration steps.

1. Access the JetStream portal from the vCenter appliance.
2. [Add an external storage site.](#)
3. [Deploy a JetStream DRVA appliance.](#)
4. [Create a JetStream replication log store volume](#) using one of the datastores available to the Azure VMware Solution cluster.

TIP

Fast local storage, such as vSAN datastore, is preferred for the replication log volume.

5. [Create a JetStream protected domain.](#) You'll provide the Azure Blob Storage site, JetStream DRVA instance, and replication log volume created in previous steps.
6. [Select the VMs](#) you want to protect and then [start VM protection.](#)

For remaining configuration steps for JetStream DR, such as creating a failover runbook, invoking failover to the DR site, and invoking failback to the primary site, see the [JetStream Admin Guide documentation](#).

Disable JetStream DR on an Azure VMware Solution cluster

This cmdlet disables JetStream DR only on one of the clusters and doesn't completely uninstall JetStream DR.

1. Select **Run command > Packages > Disable-JetDRForCluster.**
2. Provide the required values or change the default values, and then select **Run.**

FIELD	VALUE
ProtectedCluster	Name of the Azure VMware Solution private cloud cluster currently protected by JetStream DR, for example, Cluster-1 . You can only provide one cluster name to be disabled.
Credential	Credentials of the root user of the JetStream MSA VM.
MSIp	IP address of the JetStream MSA VM.
Specify name for execution	Alphanumeric name of the execution, for example, Disable-JetDRForCluster-Exec1 . It's used to verify if the cmdlet ran successfully and should be unique for each run.

3. [View the status of the execution.](#)

Uninstall JetStream DR

1. Select **Run command > Packages > Invoke-PreflightJetDRUninstall**. This cmdlet checks if the cluster has at least four hosts (minimum required).
2. Provide the required values or change the default values, and then select **Run**.

FIELD	VALUE
ProtectedCluster	Name of the Azure VMware Solution private cloud cluster currently protected by JetStream DR, for example, Cluster-1 . You can only provide one cluster name during uninstall.
Credential	Credentials of the root user of the JetStream MSA VM.
MSIp	IP address of the JetStream MSA VM.
Specify name for execution	Alphanumeric name of the execution, for example, Invoke-PreflightJetDRUninstall-Exec1 . It's used to verify if the cmdlet ran successfully and should be unique for each run.

3. [View the status of the execution.](#)
4. After the preflight cmdlet completes successfully, select **Uninstall-JetDR**, provide the required values or change the default values, and select **Run**.

FIELD	VALUE
ProtectedCluster	Name of the Azure VMware Solution private cloud cluster currently protected by JetStream DR, for example, Cluster-1 . You can only provide one cluster name during uninstall.
Credential	Credentials of the root user of the JetStream MSA VM.
MSIp	IP address of the JetStream MSA VM.
Specify name for execution	Alphanumeric name of the execution, for example, Uninstall-JetDR-Exec1 . It's used to verify if the cmdlet ran successfully and should be unique for each run.

5. [View the status of the execution.](#)

Support

JetStream DR is a solution that [JetStream Software](#) supports. For any product or support issues with JetStream, contact support-avs@jetstreamsoft.com.

Azure VMware Solution uses the Run command to automate both the install and uninstall of JetStream DR. Contact Microsoft support for any issue with the run commands. For issues with JetStream install and uninstall cmdlets, contact JetStream for support.

Next steps

- [Infrastructure Setup: JetStream DR for Azure VMware Solution](#)

- [JetStream DR for Azure VMware Solution \(Full demo\)](#)
 - [Get started with JetStream DR for Azure VMware Solution](#)
 - [Configure and protect VMs](#)
 - [Failover to Azure VMware Solution](#)
 - [Failback to on-premises](#)

Deploy disaster recovery using VMware HCX

12/16/2022 • 5 minutes to read • [Edit Online](#)

In this article, you'll deploy disaster recovery of your virtual machines (VMs) with VMware HCX solution and using an Azure VMware Solution private cloud as the recovery or target site.

IMPORTANT

Although part of HCX, VMware HCX Disaster Recovery (DR) is not recommended for large deployments. The disaster recovery orchestration is 100% manual, and Azure VMware Solution currently doesn't have runbooks or features to support manual HCX DR failover. For enterprise-class disaster recovery, refer to VMware Site Recovery Manager (SRM) or VMware business continuity and disaster recovery (BCDR) solutions.

VMware HCX provides various operations that provide fine control and granularity in replication policies.

Available Operations include:

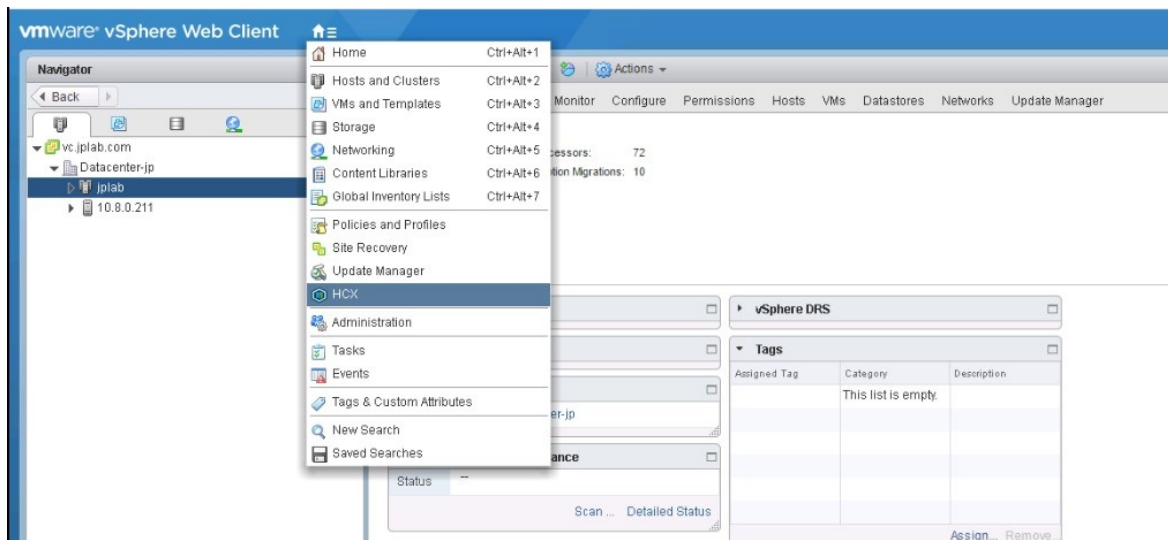
- **Reverse** – After a disaster has occurred. Reverse helps make Site B the source site and Site A, where the protected VM now lives.
- **Pause** – Pause the current replication policy associated with the VM selected.
- **Resume** - Resume the current replication policy associated with the VM selected.
- **Remove** - Remove the current replication policy associated with the VM selected.
- **Sync Now** – Out of bound sync source VM to the protected VM.

This guide covers the following replication scenarios:

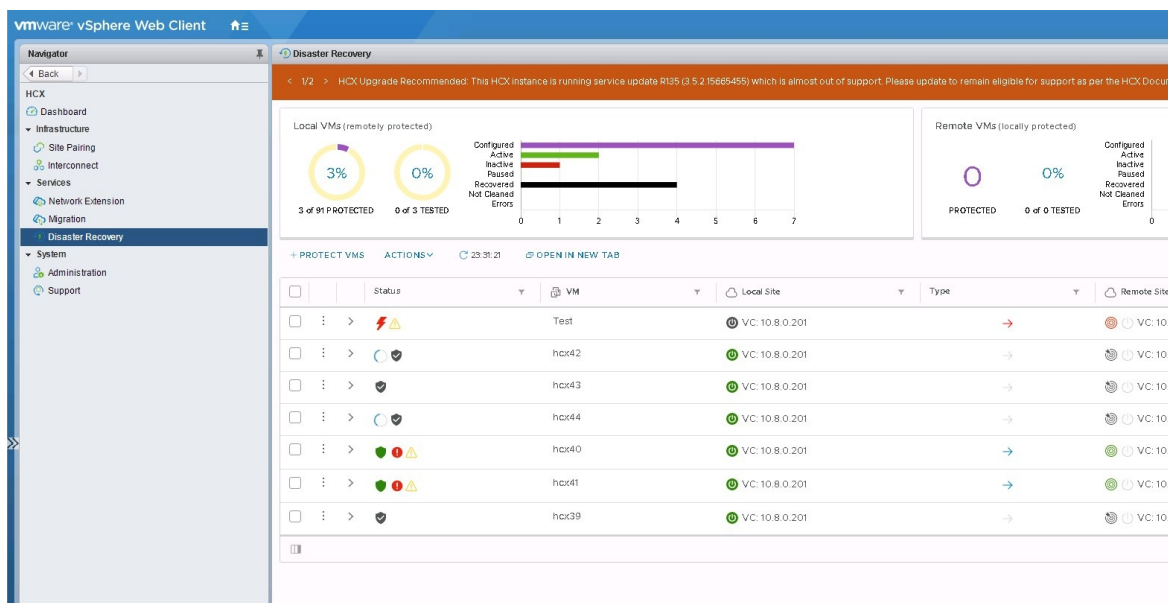
- Protect a VM or a group of VMs.
- Complete a Test Recover of a VM or a group of VMs.
- Recover a VM or a group of VMs.
- Reverse Protection of a VM or a group of VMs.

Protect VMs

1. Log into **vSphere Client** on the source site and access **HCX plugin**.



2. Enter the Disaster Recovery area and select PROTECT VMS.



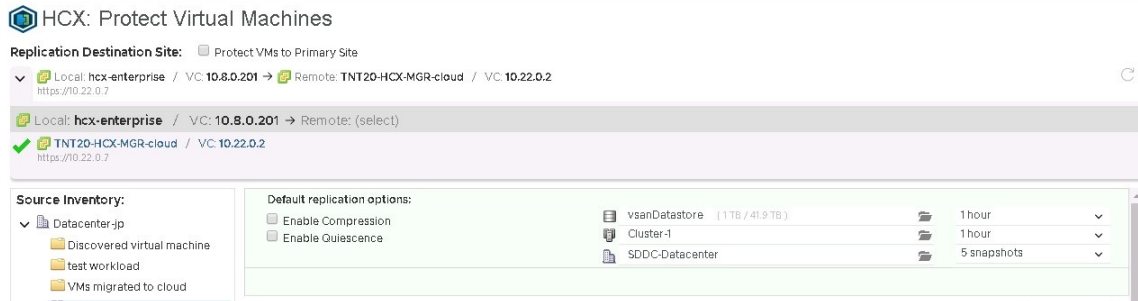
3. Select the Source and the Remote sites. The Remote site in this case should be the Azure VMware Solution private cloud.



4. If needed, select the Default replication options:

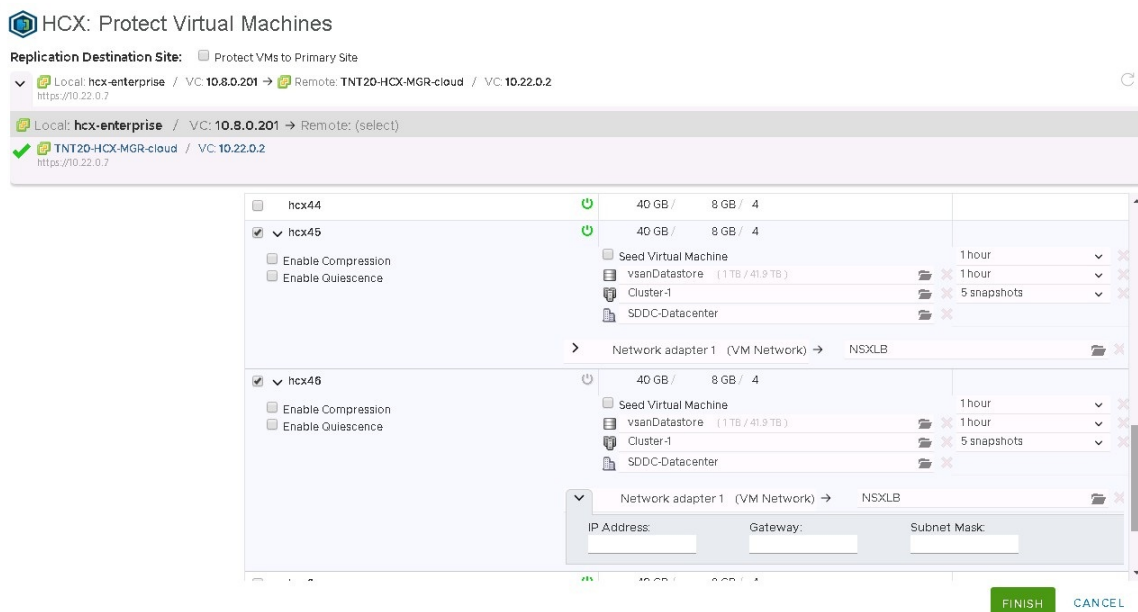
- **Enable Compression:** Recommended for low throughput scenarios.
- **Enable Quiescence:** Pauses the VM to ensure a consistent copy is synced to the remote site.
- **Destination Storage:** Remote datastore for the protected VMs, and in an Azure VMware Solution private cloud, which should be the vSAN datastore.
- **Compute Container:** Remote vSphere Cluster or Resource Pool.

- **Destination Folder:** Remote destination folder, which is optional, and if no folder is selected, the VMs are placed directly under the selected cluster.
- **RPO:** Synchronization interval between the source VM and the protected VM. It can be anywhere from 5 minutes to 24 hours.
- **Snapshot interval:** Interval between snapshots.
- **Number of Snapshots:** Total number of snapshots within the configured snapshot interval.

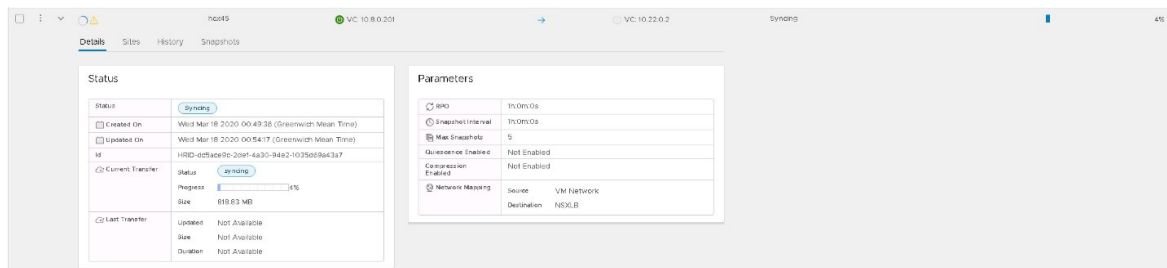


5. Select one or more VMs from the list and configure the replication options as needed.

By default, the VMs inherit the Global Settings Policy configured in the Default replication options. For each network interface in the selected VM, configure the remote **Network Port Group** and select **Finish** to start the protection process.



6. Monitor the process for each of the selected VMs in the same disaster recovery area.



7. After the VM has been protected, you can view the different snapshots in the Snapshots tab.

Replica Snapshot	Transfer Bytes	Duration	Tested On	Test Status
Wed Mar 18 2020 09:18:22 (Greenwich Mean Time)	1.7 MB	0h:0m:3s	-	⚠
Wed Mar 18 2020 07:18:29 (Greenwich Mean Time)	1.65 MB	0h:0m:2s	-	⚠
Wed Mar 18 2020 06:18:36 (Greenwich Mean Time)	1.66 MB	0h:0m:2s	-	⚠
Wed Mar 18 2020 05:19:01 (Greenwich Mean Time)	1.68 MB	0h:0m:2s	-	⚠
Wed Mar 18 2020 04:19:08 (Greenwich Mean Time)	49.54 MB	0h:0m:5s	-	⚠
Wed Mar 18 2020 03:19:12 (Greenwich Mean Time)	1.73 MB	0h:0m:3s	-	⚠

The yellow triangle means the snapshots and the virtual machines haven't been tested in a Test Recovery operation.

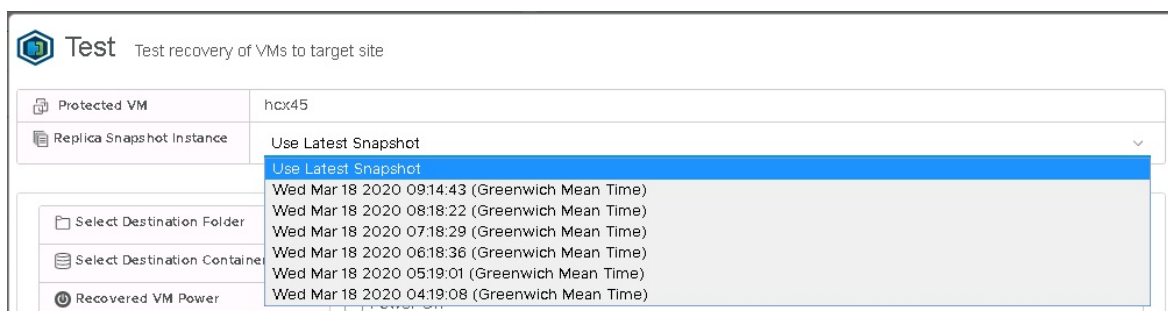
There are key differences between a VM that is powered off and one powered on. The image shows the syncing process for a powered-on VM. It starts the syncing process until it finishes the first snapshot, which is a full copy of the VM, and then completes the next ones in the configured interval. It syncs a copy for a powered off VM, and then the VM appears as inactive, and protection operation shows as completed. When the VM is powered on, it starts the syncing process to the remote site.

Complete a test recover of VMs

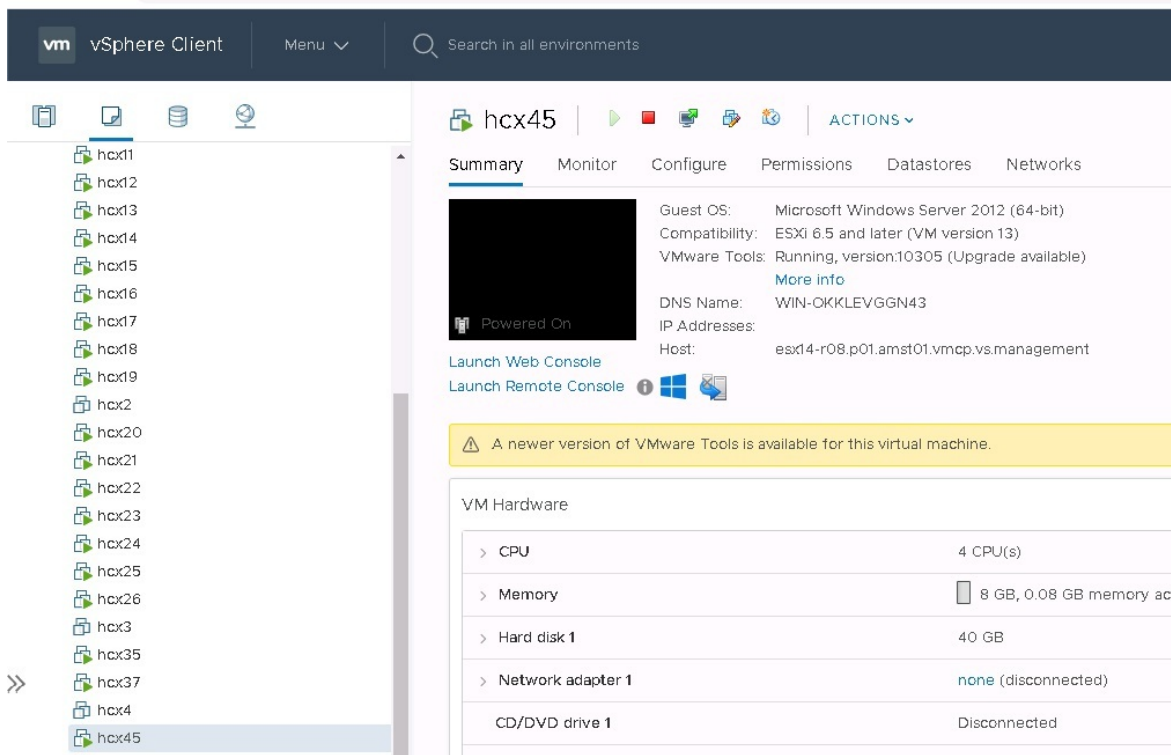
1. Log into **vSphere Client** on the remote site, which is the Azure VMware Solution private cloud.
2. Within the **HCX plugin**, in the Disaster Recovery area, select the vertical ellipses on any VM to display the operations menu and then select **Test Recover VM**.



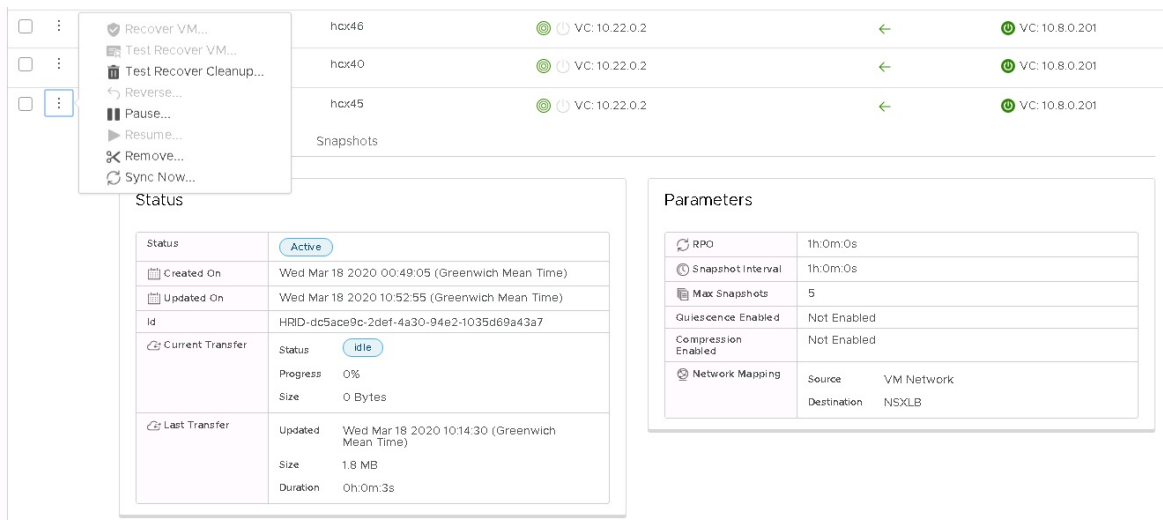
3. Select the options for the test and the snapshot you want to use to test different states of the VM.



4. After selecting **Test**, the recovery operation begins.
5. When finished, you can check the new VM in the Azure VMware Solution private cloud vCenter Server.



6. After testing has been done on the VM or any application running on it, do a cleanup to delete the test instance.

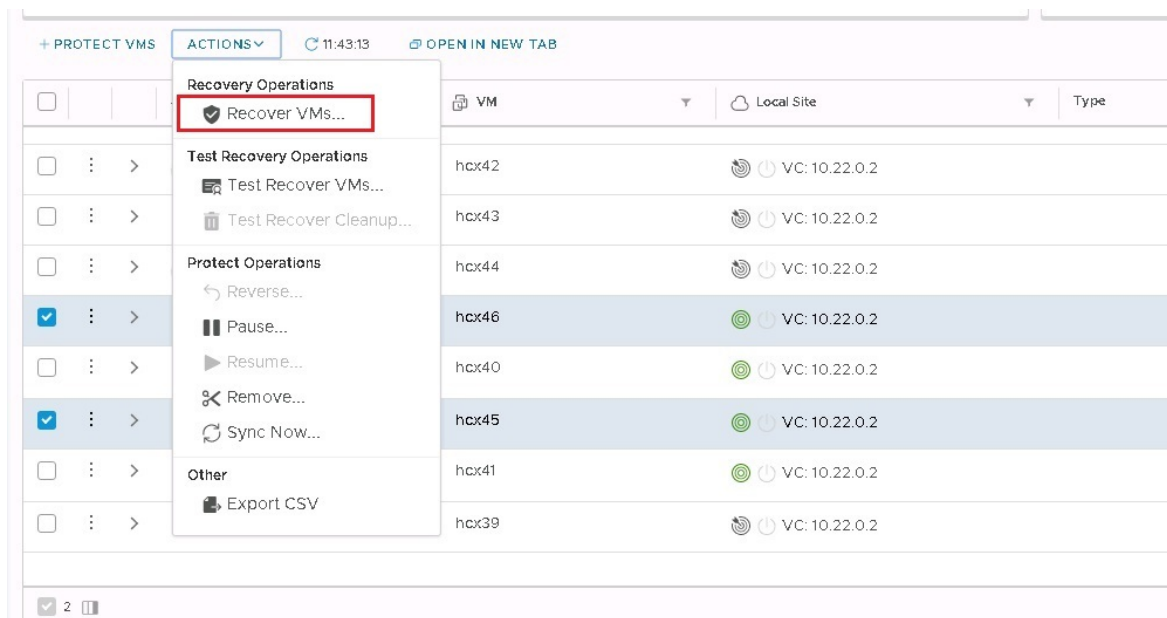


Recover VMs

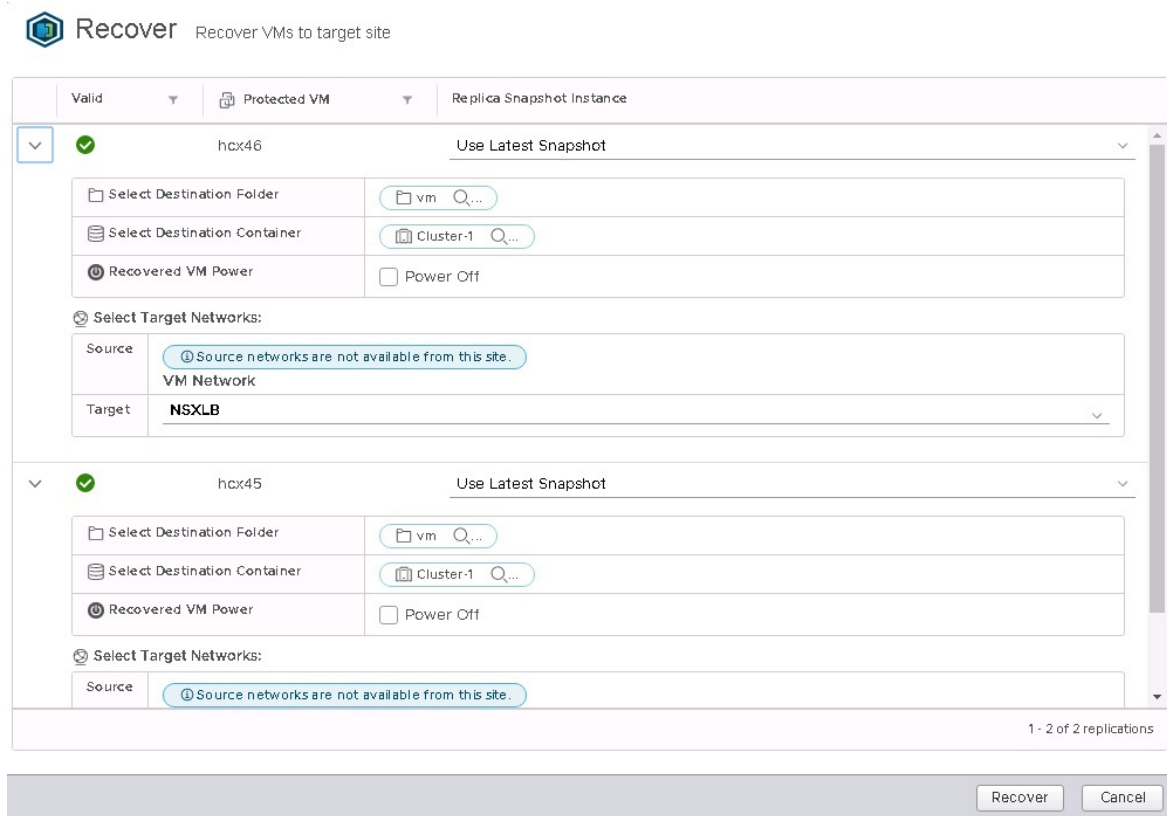
1. Log into vSphere Client on the remote site, which is the Azure VMware Solution private cloud, and access the HCX plugin.

For the recovery scenario, a group of VMs used for this example.

2. Select the VM to be recovered from the list, open the ACTIONS menu, and select Recover VMs.



3. Configure the recovery options for each instance and select **Recover** to start the recovery operation.



4. After the recovery operation is completed, the new VMs appear in the remote vCenter Server inventory.

Complete a reverse replication on VMs

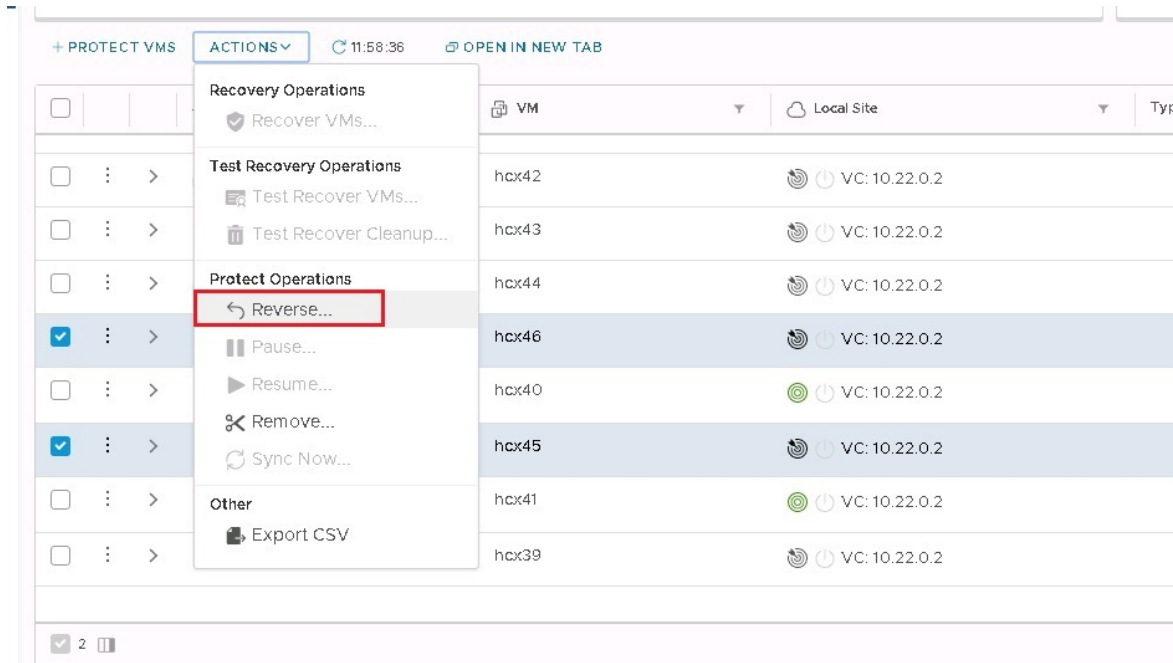
1. Log into vSphere Client on your Azure VMware Solution private cloud, and access HCX plugin.

NOTE

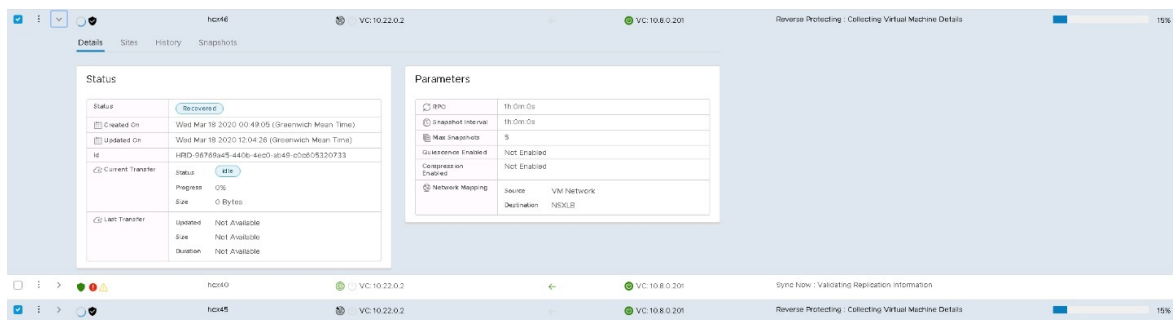
Ensure the original VMs on the source site are powered off before you start the reverse replication. The operation fails if the VMs aren't powered off.

2. From the list, select the VMs to be replicated back to the source site, open the **ACTIONS** menu, and select **Reverse**.

3. Select **Reverse** to start the replication.



4. Monitor on the details section of each VM.



Disaster recovery plan automation

VMware HCX currently doesn't have a built-in mechanism to create and automate a disaster recovery plan. However, VMware HCX provides a set of REST APIs, including APIs for the Disaster Recovery operation. The API specification can be accessed within VMware HCX Manager in the URL.

These APIs cover the following operations in Disaster Recovery.

- Protect
- Recover
- Test Recover
- Planned Recover
- Reverse
- Query
- Test Cleanup
- Pause
- Resume
- Remove Protection

- Reconfigure

An example of a recover operation payload in JSON is shown below.

```
[
  {
    "replicationId": "string",
    "needPowerOn": true,
    "instanceId": "string",
    "source": {
      "endpointType": "string",
      "endpointId": "string",
      "endpointName": "string",
      "resourceType": "string",
      "resourceId": "string",
      "resourceName": "string"
    },
    "destination": {
      "endpointType": "string",
      "endpointId": "string",
      "endpointName": "string",
      "resourceType": "string",
      "resourceId": "string",
      "resourceName": "string"
    },
    "placement": [
      {
        "containerType": "string",
        "containerId": "string"
      }
    ],
    "resourceId": "string",
    "forcePowerOff": true,
    "isTest": true,
    "forcePowerOffAfterTimeout": true,
    "isPlanned": true
  }
]
```

```
}  
]
```

With these APIs, you can build a custom mechanism to automate a disaster recovery plan's creation and execution.

Deploy disaster recovery with VMware Site Recovery Manager

12/16/2022 • 10 minutes to read • [Edit Online](#)

This article explains how to implement disaster recovery for on-premises VMware virtual machines (VMs) or Azure VMware Solution-based VMs. The solution in this article uses [VMware Site Recovery Manager \(SRM\)](#) and vSphere Replication with Azure VMware Solution. Instances of SRM and replication servers are deployed at both the protected and the recovery sites.

SRM is a disaster recovery solution designed to minimize downtime of the virtual machines in an Azure VMware Solution environment if there was a disaster. SRM automates and orchestrates failover and failback, ensuring minimal downtime in a disaster. Also, built-in non-disruptive testing ensures your recovery time objectives are met. Overall, SRM simplifies management through automation and ensures fast and highly predictable recovery times.

vSphere Replication is VMware's hypervisor-based replication technology for vSphere VMs. It protects VMs from partial or complete site failures. In addition, it simplifies DR protection through storage-independent, VM-centric replication. vSphere Replication is configured on a per-VM basis, allowing more control over which VMs are replicated.

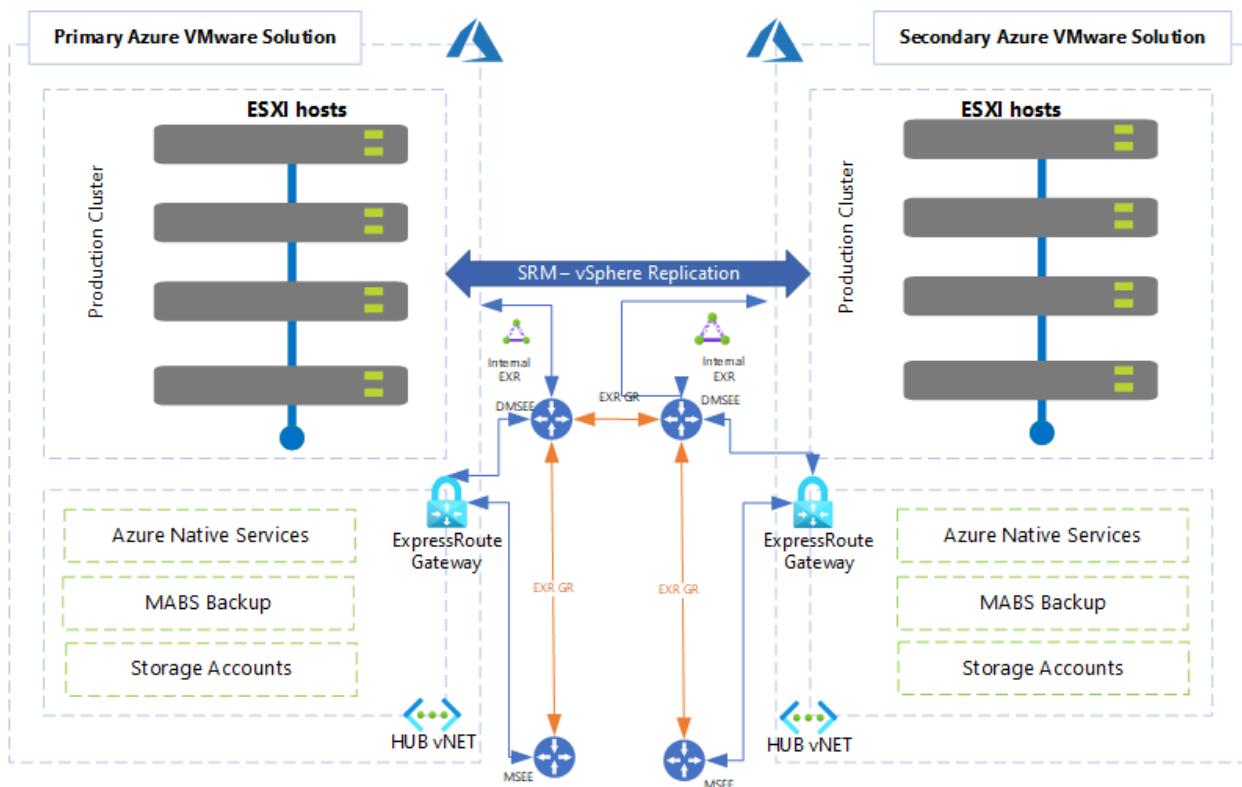
In this article, you'll implement disaster recovery for on-premises VMware virtual machines (VMs) or Azure VMware Solution-based VMs.

Supported scenarios

SRM helps you plan, test, and run the recovery of VMs between a protected vCenter Server site and a recovery vCenter Server site. You can use SRM with Azure VMware Solution with the following two DR scenarios:

- On-premises VMware to Azure VMware Solution private cloud disaster recovery
- Primary Azure VMware Solution to Secondary Azure VMware Solution private cloud disaster recovery

The diagram shows the deployment of the primary Azure VMware Solution to secondary Azure VMware Solution scenario.



You can use SRM to implement different types of recovery, such as:

- **Planned migration** commences when both primary and secondary Azure VMware Solution sites are running and fully functional. It's an orderly migration of virtual machines from the protected site to the recovery site where no data loss is expected when migrating workloads in an orderly fashion.
- **Disaster recovery** using SRM can be invoked when the protected Azure VMware Solution site goes offline unexpectedly. Site Recovery Manager orchestrates the recovery process with the replication mechanisms to minimize data loss and system downtime.

In Azure VMware Solution, only individual VMs can be protected on a host by using SRM in combination with vSphere Replication.

- **Bidirectional Protection** uses a single set of paired SRM sites to protect VMs in both directions. Each site can simultaneously be a protected site and a recovery site, but for a different set of VMs.

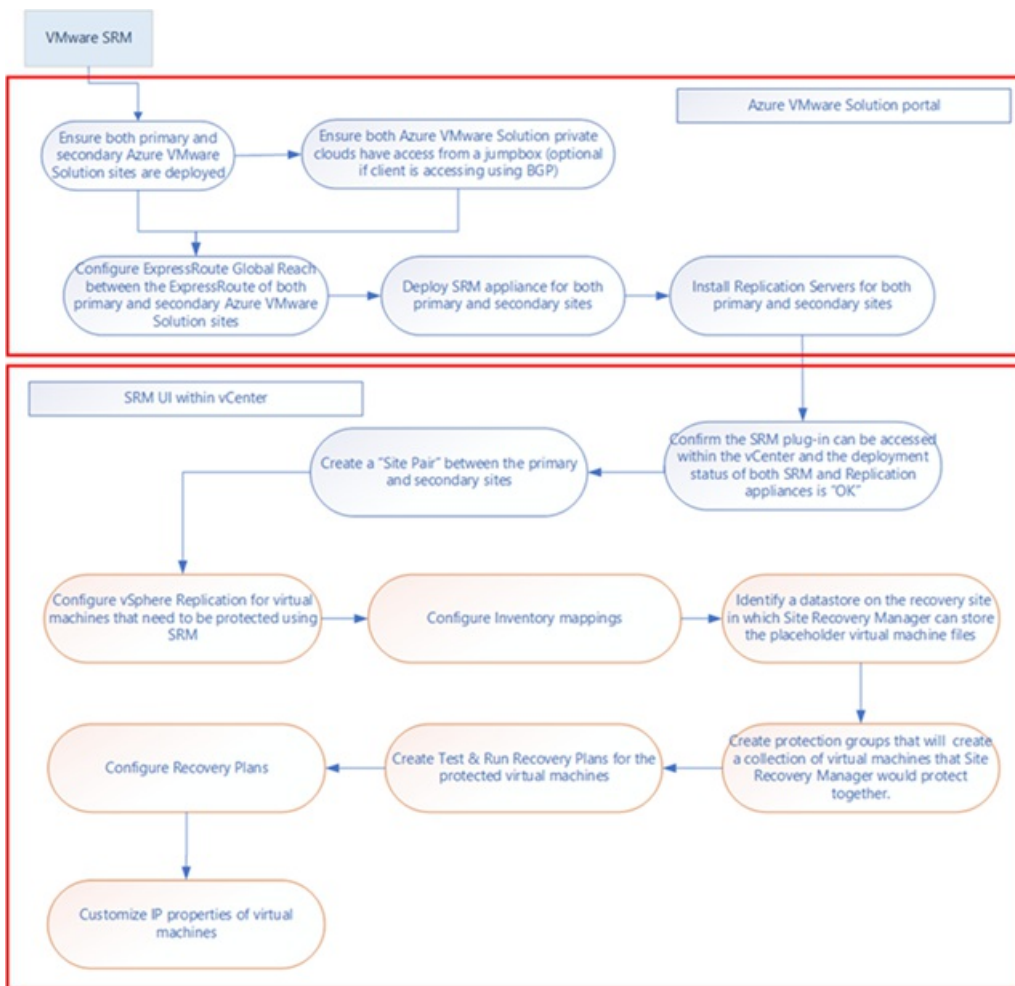
IMPORTANT

Azure VMware Solution doesn't support:

- Array-based replication and storage policy protection groups
- VVOLs Protection Groups
- SRM IP customization using SRM command-line tools
- One-to-Many and Many-to-One topology
- Custom SRM plug-in identifier or extension ID

Deployment workflow

The workflow diagram shows the Primary Azure VMware Solution to secondary workflow. In addition, it shows steps to take within the Azure portal and the VMware vSphere environments of Azure VMware Solution to achieve the end-to-end protection of VMs.



Prerequisites

Make sure you've explicitly provided the remote user the VRM administrator and SRM administrator roles in the remote vCenter Server.

Scenario: On-premises to Azure VMware Solution

- Azure VMware Solution private cloud deployed as a secondary region.
- [DNS resolution](#) to on-premises SRM and virtual cloud appliances.

NOTE

For private clouds created on or after July 1, 2021, you can configure private DNS resolution. For private clouds created before July 1, 2021, that need a private DNS resolution, open a [support request](#) to request **Private DNS configuration**.

- ExpressRoute connectivity between on-premises and Azure VMware Solution - 2 Gbps.

Scenario: Primary Azure VMware Solution to secondary

- Azure VMware Solution private cloud must be deployed in the primary and secondary region.

Azure VMware Solution

Microsoft

+ Create Manage view Refresh Export to CSV Open query Assign tags Feedback

Filter for any field... Subscription == all Resource group == all Location == all Add filter

Showing 1 to 2 of 2 records.

Name	Type	Resource group	Location
AVS_EUW_001	AVS Private cloud	AVS_EUW_001	West Europe
AVS_UKS_001	AVS Private cloud	AVS_UKS_001	UK South

- Connectivity, like ExpressRoute Global Reach, between the source and target Azure VMware Solution private cloud.

The screenshot shows the 'ExpressRoute Global Reach' settings page for the AVS_UKS_001 private cloud. The page includes a search bar, navigation tabs (Settings, ExpressRoute, HCX, Public IP, ExpressRoute Global Reach), and a list of connections. A table at the bottom shows a single connection with the following details:

Subscription	Resource group	ExpressRoute circuit	Authorization key	State
0b4f9382-c705-4cd4-9dd...	tnt86-cust-p01-westeurop...	/subscriptions/0b4f...	*****	Connected

Install SRM in Azure VMware Solution

- In your on-premises datacenter, install VMware SRM and vSphere Replication.

NOTE

Use the [Two-site Topology with one vCenter Server instance per PSC](#) deployment model. Also, make sure that the [required vSphere Replication Network ports](#) are opened.

- In your Azure VMware Solution private cloud, under **Manage**, select **Add-ons > Disaster recovery**.

The default CloudAdmin user in the Azure VMware Solution private cloud doesn't have sufficient privileges to install VMware SRM or vSphere Replication. The installation process involves multiple steps outlined in the [Prerequisites](#) section. Instead, you can install VMware SRM with vSphere Replication as an add-on service from your Azure VMware Solution private cloud.

Home > Microsoft.AVS-20200914145430 > Azure private cloud

Azure private cloud | Add-ons

AVS Private cloud

Search (Ctrl+/) << Overview Disaster recovery

- Overview
- Activity log
- Access control (IAM)
- Tags
- Diagnose and solve problems

Settings

- Locks
- Export template

Manage

- Connectivity
- Identity
- Clusters
- + Add-ons**

Monitoring

- Alerts
- Metrics

Support + troubleshooting

- Resource health
- New support request

Enhance your private cloud with these optional features

Disaster recovery
[Start >](#)

- From the Disaster Recovery Solution drop-down, select VMware Site Recovery Manager (SRM) – vSphere Replication.

Home > Microsoft.AVS-20200914145430 > Azure private cloud

Azure private cloud | Add-ons

AVS Private cloud

Search (Ctrl+/) << Overview **Disaster recovery**

Disaster recovery solution *
VMware Site Recovery Manager (SRM) - vSphere replication

License key
Enter your license key.

I agree with terms and conditions.
By checking this, you allow Microsoft Azure to install this software on your behalf. Microsoft does not manage your license. You are responsible for maintaining your license with VMware.

Installation would take approximately 30 minutes to complete. You can track progress using Azure notifications. Once complete, go to vCenter and complete configuration and site pairing.

Install

- Provide the License key, select agree with terms and conditions, and then select **Install**.

NOTE

If you don't provide the license key, SRM is installed in an Evaluation mode. The license is used only to enable VMware SRM.

The screenshot shows the Azure portal interface for configuring a disaster recovery solution. The breadcrumb path is 'Home > Microsoft.AVS-20200914145430 > Azure private cloud'. The page title is 'Azure private cloud | Add-ons'. The left sidebar contains navigation options: Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Settings (Locks, Export template), and Manage (Connectivity, Identity, Clusters, Add-ons). The main content area is titled 'Disaster recovery' and shows the 'Disaster recovery solution' dropdown set to 'VMware Site Recovery Manager (SRM) - vSphere replication'. Below this, the 'License key' field is highlighted with a red box and contains a long alphanumeric string. A checkbox labeled 'I agree with terms and conditions.' is checked. Below the checkbox, there is explanatory text about the license and installation time. At the bottom of the configuration area, there is a blue 'Install' button.

Install the vSphere Replication appliance

After the SRM appliance installs successfully, you'll need to install the vSphere Replication appliances. Each replication server accommodates up to 200 protected VMs. Scale in or scale out as per your needs.

1. From the **Replication using** drop-down, on the **Disaster recovery** tab, select **vSphere Replication**.

Setup replication

Replication using *

vSphere Replication

Each vSphere replication server accommodates up to 200 protected VMs. Scale in or scale out as per your needs.

vSphere servers * ⓘ

3

Install

2. Move the vSphere server slider to indicate the number of replication servers you want based on the number of VMs to be protected. Then select **Install**.

Setup replication

Replication using *

vSphere Replication

Each vSphere replication server accomodates up to 200 protected VMs. Scale in or scale out as per your needs.

vSphere servers * ⓘ

3

Install

3. Once installed, verify that both SRM and the vSphere Replication appliances are installed.

TIP

The Uninstall button indicates that both SRM and the vSphere Replication appliances are currently installed.

Overview Disaster recovery

Disaster recovery solution

VMware Site Recovery Manager (SRM)

i Disaster recovery with Site Recovery Manager is a [preview feature](#).

License key ⓘ

Enter your license key.

Save

or

Completely remove and uninstall SRM cloud appliance.

This will remove the software. All site pairs should be deleted before uninstalling.

Uninstall VMware Site Recovery Manager (SRM)

Replication using

vSphere Replication

Each vSphere replication server accomodates up to 200 protected VMs. Scale in or scale out as per your needs.

vSphere servers * ⓘ

3

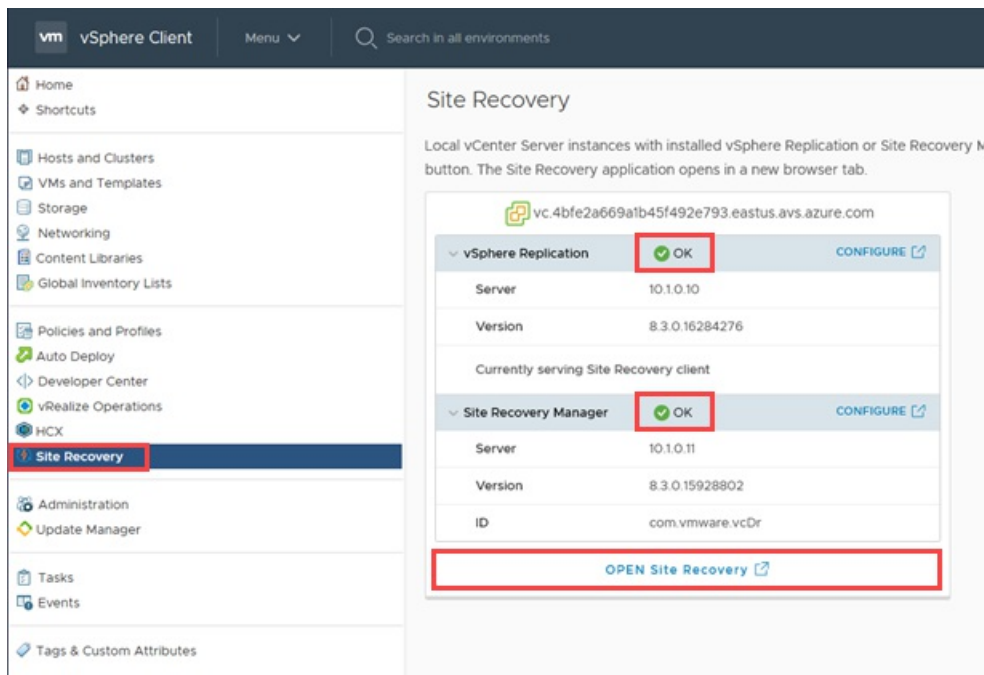
Save

Uninstall vSphere Replication

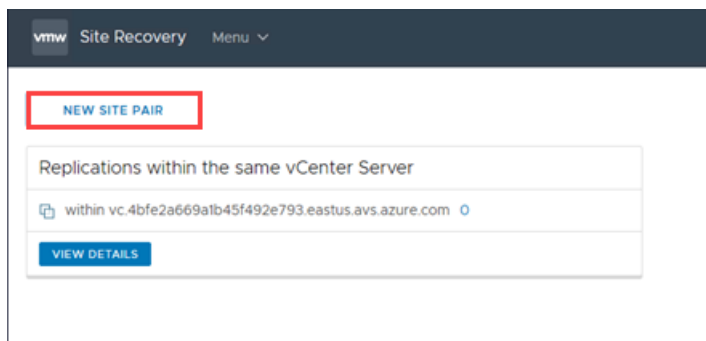
Configure site pairing in vCenter Server

After installing VMware SRM and vSphere Replication, you need to complete the configuration and site pairing in vCenter Server.

1. Sign in to vCenter Server as cloudadmin@vsphere.local.
2. Navigate to **Site Recovery**, check the status of both vSphere Replication and VMware SRM, and then select **OPEN Site Recovery** to launch the client.



3. Select **NEW SITE PAIR** in the Site Recovery (SR) client in the new tab that opens.

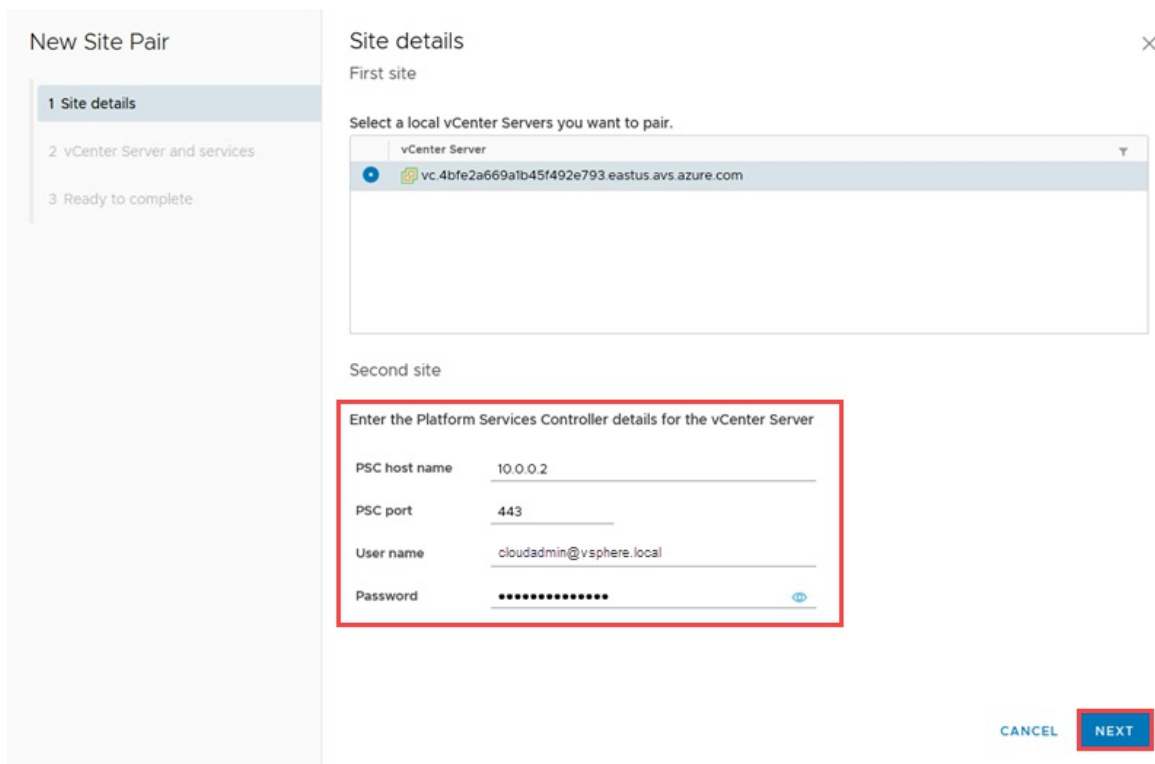


4. Enter the remote site details, and then select **NEXT**.

NOTE

An Azure VMware Solution private cloud operates with an embedded Platform Services Controller (PSC), so only one local vCenter can be selected. If the remote vCenter Server is using an embedded Platform Service Controller (PSC), use the vCenter Server's FQDN (or its IP address) and port to specify the PSC.

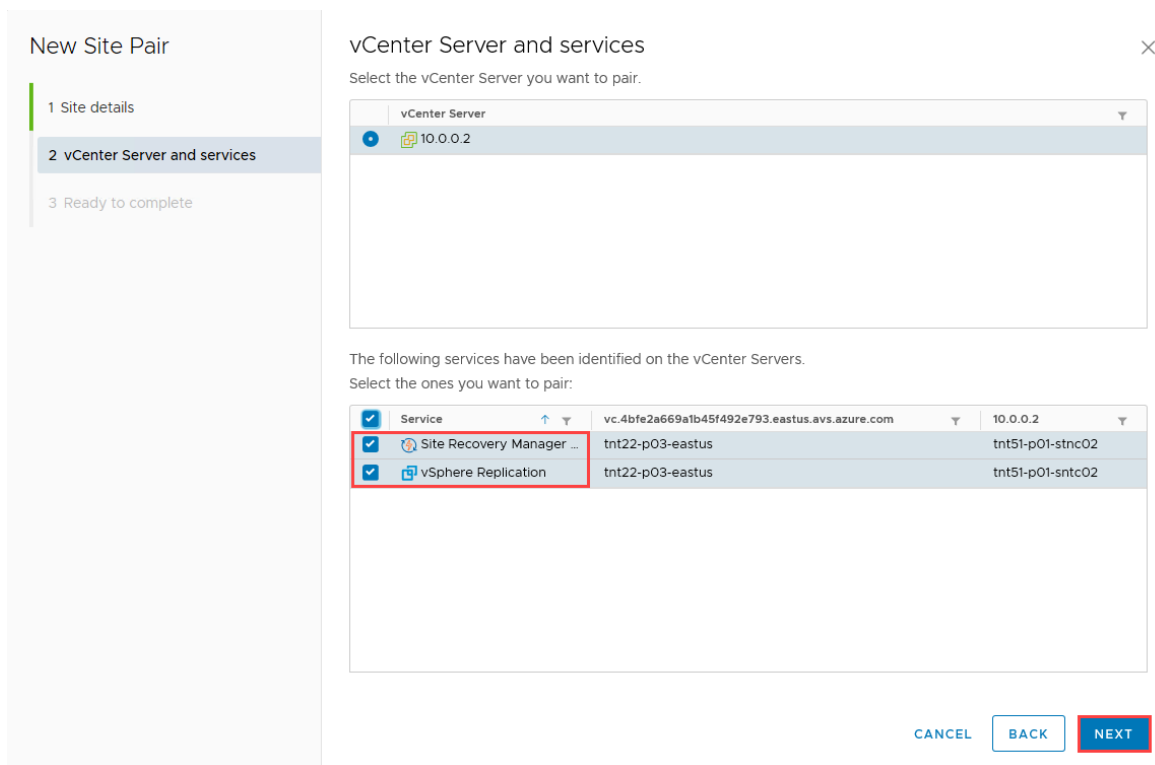
The remote user must have sufficient permissions to perform the pairings. An easy way to ensure this is to give that user the VRM administrator and SRM administrator roles in the remote vCenter Server. For a remote Azure VMware Solution private cloud, cloudadmin is configured with those roles.



5. Select **CONNECT** to accept the certificate for the remote vCenter Server.

At this point, the client should discover the VRM and SRM appliances on both sides as services to pair.

6. Select the appliances to pair and then select **NEXT**.



7. Select **CONNECT** to accept the certificates for the remote VMware SRM and the remote vCenter Server (again).

8. Select **CONNECT** to accept the certificates for the local VMware SRM and the local vCenter Server.

9. Review the settings and then select **FINISH**.

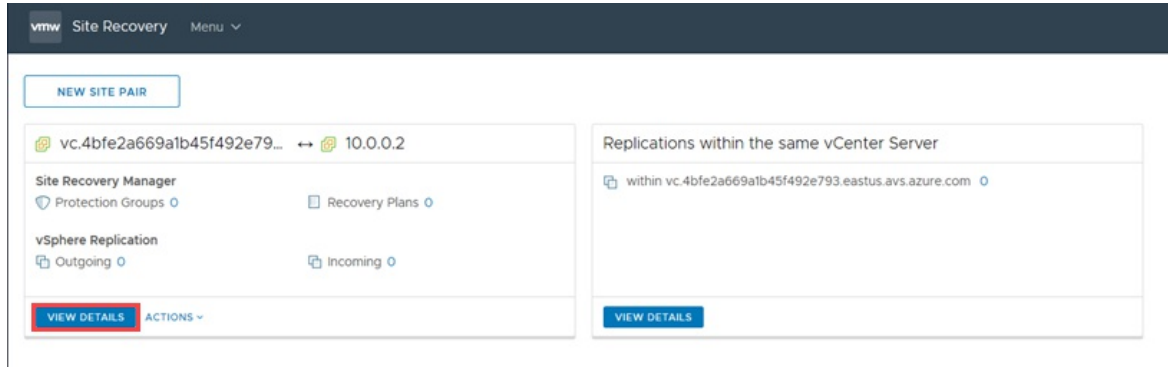
If successful, the client displays another panel for the pairing. However, if unsuccessful, an alarm will be reported.

- At the bottom, in the right corner, select the double-up arrow to expand the panel to show **Recent Tasks** and **Alarms**.

NOTE

The SR client sometimes takes a long time to refresh. If an operation seems to take too long or appears "stuck", select the refresh icon on the menu bar.

- Select **VIEW DETAILS** to open the panel for remote site pairing, which opens a dialog to sign in to the remote vCenter Server.



- Enter the username with sufficient permissions to do replication and site recovery and then select **LOG IN**.

For pairing, the login, which is often a different user, is a one-time action to establish pairing. The SR client requires this login every time the client is launched to work with the pairing.

NOTE

The user with sufficient permissions should have **VRM administrator** and **SRM administrator** roles given to them in the remote vCenter Server. The user should also have access to the remote vCenter Server inventory, like folders and datastores. For a remote Azure VMware Solution private cloud, the cloudadmin user has the appropriate permissions and access.

Log In Site



Enter vCenter Server credentials

vCenter Server

10.0.0.2

User name

cloudadmin@vsphere.local

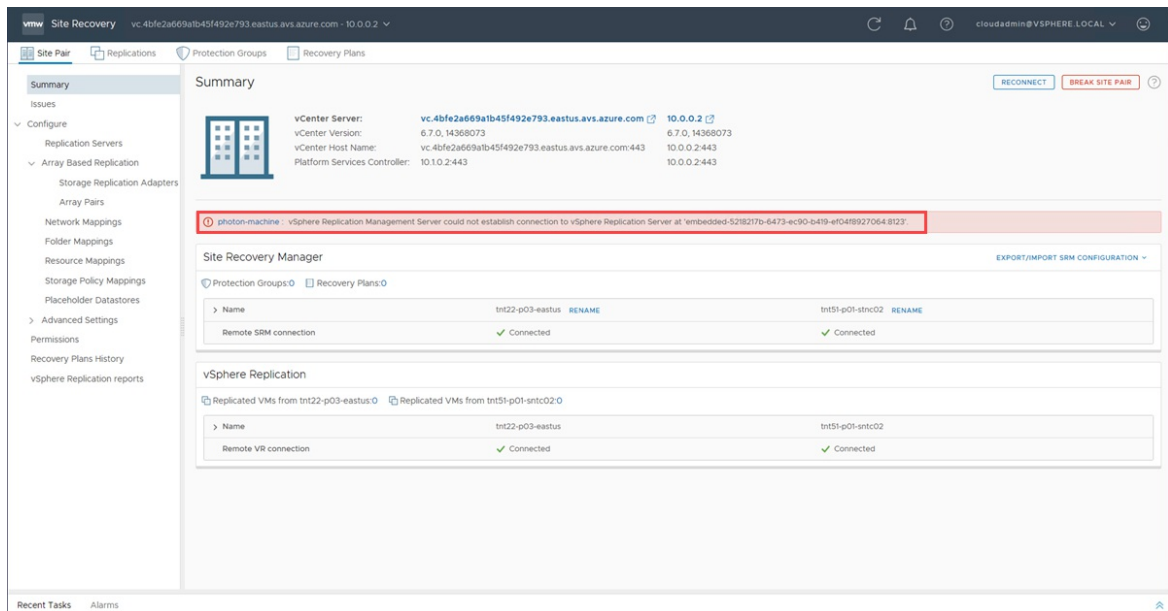
Password

●●●●●●●●●●

CANCEL

LOG IN

You'll see a warning message indicating that the embedded VRS in the local VRM isn't running. This is because Azure VMware Solution doesn't use the embedded VRS in an Azure VMware Solution private cloud. Instead, it uses VRS appliances.



SRM protection, re-protection, and failback

After you've created the site pairing, follow the VMware documentation mentioned below for end-to-end protection of VMs from the Azure portal.

- [Using vSphere Replication with Site Recovery Manager \(vmware.com\)](#)
- [Inventory Mappings for Array-Based Replication Protection Groups and vSphere Replication Protection Groups \(vmware.com\)](#)
- [About Placeholder Virtual Machines \(vmware.com\)](#)
- [vSphere Replication Protection Groups \(vmware.com\)](#)
- [Creating, Testing, and Running Recovery Plans \(vmware.com\)](#)
- [Configuring a Recovery Plan \(vmware.com\)](#)
- [Customizing IP Properties for Virtual Machines \(vmware.com\)](#)
- [How Site Recovery Manager Reprotects Virtual Machines with vSphere Replication \(vmware.com\)](#)
- [Perform a Failback \(vmware.com\)](#)

NOTE

If IP Customization Rules have been defined for network mappings between the AVS environment and the on-premises environment, these rules will not be applied on failback from the AVS environment to the on-premises environment due to a [known issue](#) with SRM 8.3.0. You can work around this limitation by removing protection from all VMs in the Protection Group and then reconfiguring protection on them prior to initiating the failback.

Ongoing management of your SRM solution

While Microsoft aims to simplify VMware SRM and vSphere Replication installation on an Azure VMware Solution private cloud, you are responsible for managing your license and the day-to-day operation of the disaster recovery solution.

Scale limitations

To learn about the limits for the VMware Site Recovery Manager Add-On with the Azure VMware Solution, check

the [Azure subscription and service limits, quotas, and constraints](#).

SRM licenses

You can install VMware SRM using an evaluation license or a production license. The evaluation license is valid for 60 days. After the evaluation period, you'll be required to obtain a production license of VMware SRM.

You can't use pre-existing on-premises VMware SRM licenses for your Azure VMware Solution private cloud. Work with your sales teams and VMware to acquire a new term-based production license of VMware SRM.

Once a production license of SRM is acquired, you'll be able to use the Azure VMware Solution portal to update SRM with the new production license.

Uninstall SRM

If you no longer require SRM, you must uninstall it in a clean manner. Before you uninstall SRM, you must remove all SRM configurations from both sites in the correct order. If you do not remove all configurations before uninstalling SRM, some SRM components, such as placeholder VMs, might remain in the Azure VMware Solution infrastructure.

1. In the vSphere Client or the vSphere Web Client, select **Site Recovery** > **Open Site Recovery**.
2. On the **Site Recovery** home tab, select a site pair and select **View Details**.
3. Select the **Recovery Plans** tab, right-click on a recovery plan and select **Delete**.

NOTE

You cannot delete recovery plans that are running.

4. Select the **Protection Groups** tab, select a protection group, and select the **Virtual Machines** tab.
5. Highlight all virtual machines, right-click, and select **Remove Protection**.

Removing protection from a VM deletes the placeholder VM from the recovery site. Repeat this operation for all protection groups.

6. In the **Protection Groups** tab, right-click a protection group and select **Delete**.

NOTE

You cannot delete a protection group that is included in a recovery plan. You cannot delete vSphere Replication protection groups that contain virtual machines on which protection is still configured.

7. Select **Site Pair** > **Configure** and remove all inventory mappings.
 - a. Select each of the **Network Mappings**, **Folder Mappings**, and **Resource Mappings** tabs.
 - b. In each tab, select a site, right-click a mapping, and select **Delete**.
8. For both sites, select **Placeholder Datastores**, right-click the placeholder datastore, and select **Remove**.
9. Select **Site Pair** > **Summary**, and select **Break Site Pair**.

NOTE

Breaking the site pairing removes all information related to registering Site Recovery Manager with Site Recovery Manager, vCenter Server, and the Platform Services Controller on the remote site.

10. In your private cloud, under **Manage**, select **Add-ons > Disaster recovery**, and then select **Uninstall the replication appliances**.
11. Once replication appliances are uninstalled, from the **Disaster recovery** tab, select **Uninstall for the Site Recovery Manager**.
12. Repeat these steps on the secondary Azure VMware Solution site.

Support

VMware Site Recovery Manager (SRM) is a Disaster Recovery solution from VMware.

Microsoft only supports install/uninstall of SRM and vSphere Replication Manager and scale up/down of vSphere Replication appliances within Azure VMware Solution.

For all other issues, such as configuration and replication, contact VMware for support.

VMware and Microsoft support teams will engage each other as needed to troubleshoot SRM issues on Azure VMware Solution.

References

- [VMware Site Recovery Manager Documentation](#)
- [Compatibility Matrices for VMware Site Recovery Manager 8.3](#)
- [VMware SRM 8.3 release notes](#)
- [VMware vSphere Replication Documentation](#)
- [Compatibility Matrices for vSphere Replication 8.3](#)
- [Operational Limits of Site Recovery Manager 8.3](#)
- [Operational Limits of vSphere Replication 8.3](#)
- [Calculate bandwidth for vSphere Replication](#)
- [SRM installation and configuration](#)
- [vSphere Replication administration](#)
- [Pre-requisites and Best Practices for SRM installation](#)
- [Network ports for SRM](#)
- [Network ports for vSphere Replication](#)

Deploy Zerto disaster recovery on Azure VMware Solution

12/16/2022 • 5 minutes to read • [Edit Online](#)

In this article, you'll learn how to implement disaster recovery for on-premises VMware or Azure VMware Solution-based virtual machines (VMs). The solution in this article uses [Zerto disaster recovery](#). Instances of Zerto are deployed at both the protected and the recovery sites.

Zerto is a disaster recovery solution designed to minimize downtime of VMs should a disaster occur. Zerto's platform is built on the foundation of Continuous Data Protection (CDP) that enables minimal or close to no data loss. The platform provides the level of protection wanted for many business-critical and mission-critical enterprise applications. Zerto also automates and orchestrates failover and failback to ensure minimal downtime in a disaster. Overall, Zerto simplifies management through automation and ensures fast and highly predictable recovery times.

Core components of the Zerto platform

COMPONENT	DESCRIPTION
Zerto Virtual Manager (ZVM)	Management application for Zerto implemented as a Windows service installed on a Windows VM. The private cloud administrator installs and manages the Windows VM. The ZVM enables Day 0 and Day 2 disaster recovery configuration. For example, configuring primary and disaster recovery sites, protecting VMs, recovering VMs, and so on. However, it doesn't handle the replication data of the protected customer VMs.
Virtual Replication appliance (vRA)	Linux VM to handle data replication from the source to the replication target. One instance of vRA is installed per ESXi host, delivering a true scale architecture that grows and shrinks along with the private cloud's hosts. The vRA manages data replication to and from protected VMs to its local or remote target, storing the data in the journal.
Zerto ESXi host driver	Installed on each VMware ESXi host configured for Zerto disaster recovery. The host driver intercepts a vSphere VM's IO and sends the replication data to the chosen vRA for that host. The vRA is then responsible for replicating the VM's data to one or more disaster recovery targets.

COMPONENT	DESCRIPTION
Zerto Cloud Appliance (ZCA)	<p>Windows VM only used when Zerto is used to recover vSphere VMs as Azure Native IaaS VMs. The ZCA is composed of:</p> <ul style="list-style-type: none"> • ZVM: A Windows service that hosts the UI and integrates with the native APIs of Azure for management and orchestration. • VRA: A Windows service that replicates the data from or to Azure. <p>The ZCA integrates natively with the platform it's deployed on, allowing you to use Azure Blob storage within a storage account on Microsoft Azure. As a result, it ensures the most cost-efficient deployment on each of these platforms.</p>
Virtual Protection Group (VPG)	<p>Logical group of VMs created on the ZVM. Zerto allows configuring disaster recovery, Backup, and Mobility policies on a VPG. This mechanism enables a consistent set of policies to be applied to a group of VMs.</p>

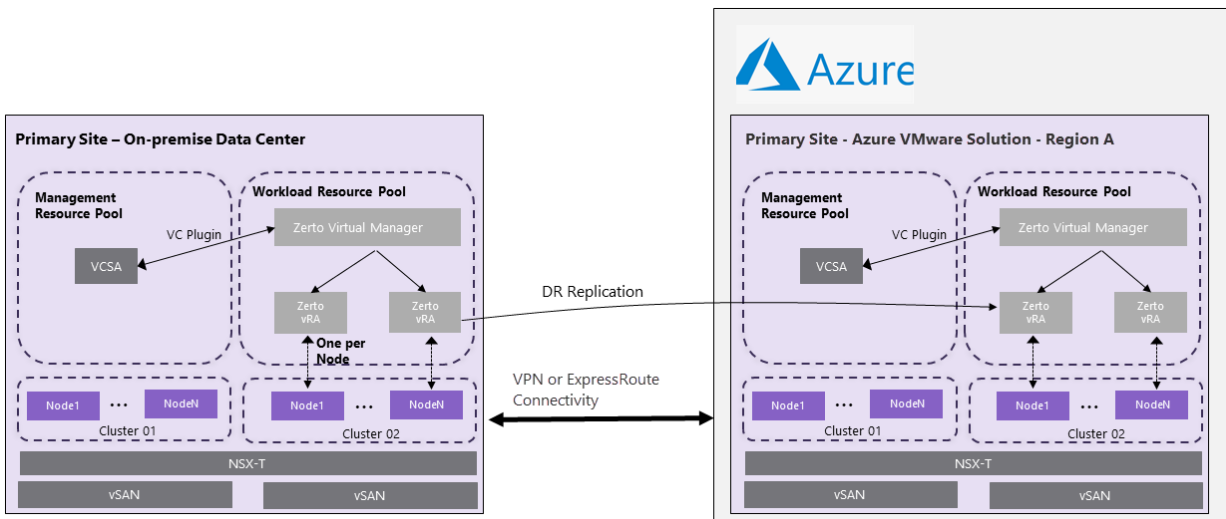
To learn more about Zerto platform architecture, see the [Zerto Platform Architecture Guide](#).

Supported Zerto scenarios

You can use Zerto with Azure VMware Solution for the following three scenarios.

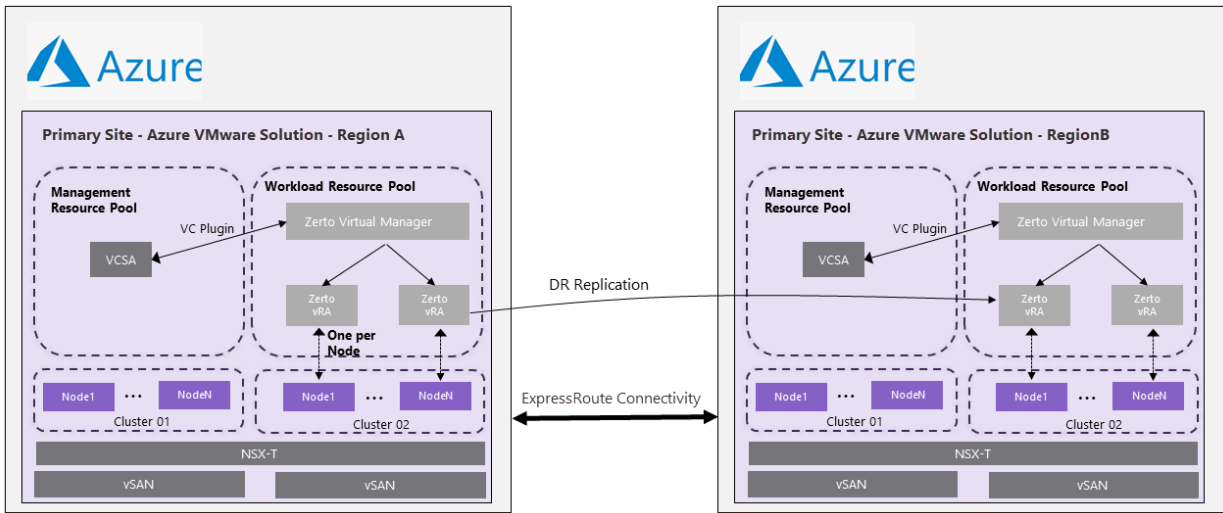
Scenario 1: On-premises VMware vSphere to Azure VMware Solution disaster recovery

In this scenario, the primary site is an on-premises vSphere-based environment. The disaster recovery site is an Azure VMware Solution private cloud.



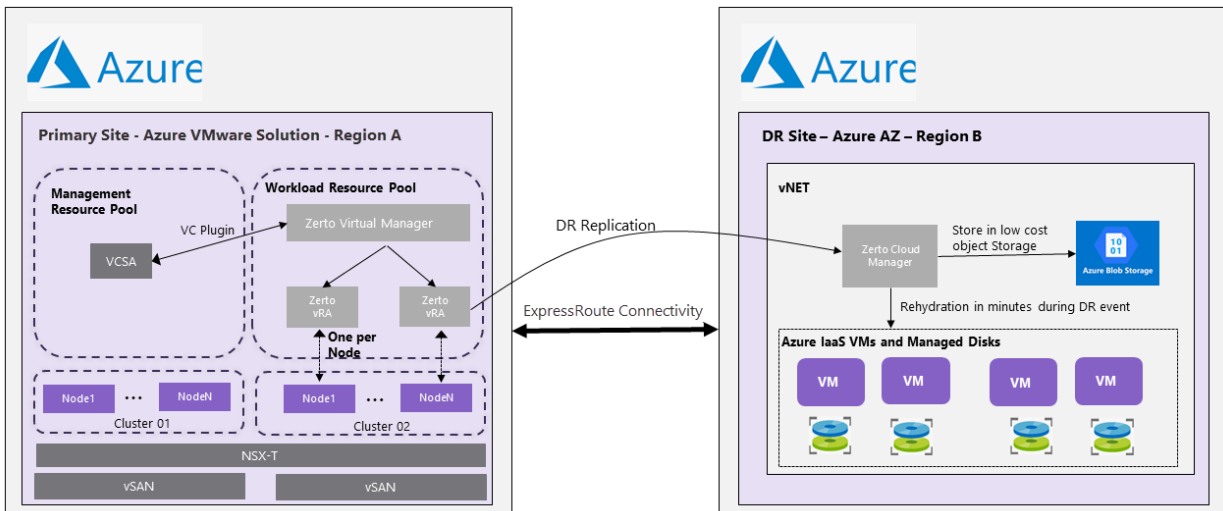
Scenario 2: Azure VMware Solution to Azure VMware Solution cloud disaster recovery

In this scenario, the primary site is an Azure VMware Solution private cloud in one Azure Region. The disaster recovery site is an Azure VMware Solution private cloud in a different Azure Region.



Scenario 3: Azure VMware Solution to IaaS VMs cloud disaster recovery

In this scenario, the primary site is an Azure VMware Solution private cloud in one Azure Region. Azure Blobs and Azure IaaS (Hyper-V based) VMs are used in times of Disaster.



Prerequisites

On-premises VMware to Azure VMware Solution disaster recovery

- Azure VMware Solution private cloud deployed as a secondary region.
- VPN or ExpressRoute connectivity between on-premises and Azure VMware Solution.

Azure VMware Solution to Azure VMware Solution cloud disaster recovery

- Azure VMware Solution private cloud must be deployed in the primary and secondary regions.

Azure VMware Solution ✎ ...

Microsoft

[+ Create](#) [Manage view](#) [Refresh](#) [Export to CSV](#) [Open query](#) [Assign tags](#) [Feedback](#)

Filter for any field... [Subscription == all](#) [Resource group == all](#) [Location == all](#) [Add filter](#)

Showing 1 to 2 of 2 records.

<input type="checkbox"/> Name ↑↓	Type ↑↓	Resource group ↑↓	Location ↑↓
<input type="checkbox"/> AVS_EUW_001	AVS Private cloud	AVS_EUW_001	West Europe
<input type="checkbox"/> AVS_UKS_001	AVS Private cloud	AVS_UKS_001	UK South

- Connectivity, like ExpressRoute Global Reach, between the source and target Azure VMware Solution private cloud.

Azure VMware Solution IaaS VMs cloud disaster recovery

- Network connectivity, ExpressRoute based, from Azure VMware Solution to the vNET used for disaster recovery.
- Follow the [Zerto Virtual Replication Azure Quickstart Guide](#) for the rest of the prerequisites.

Install Zerto on Azure VMware Solution

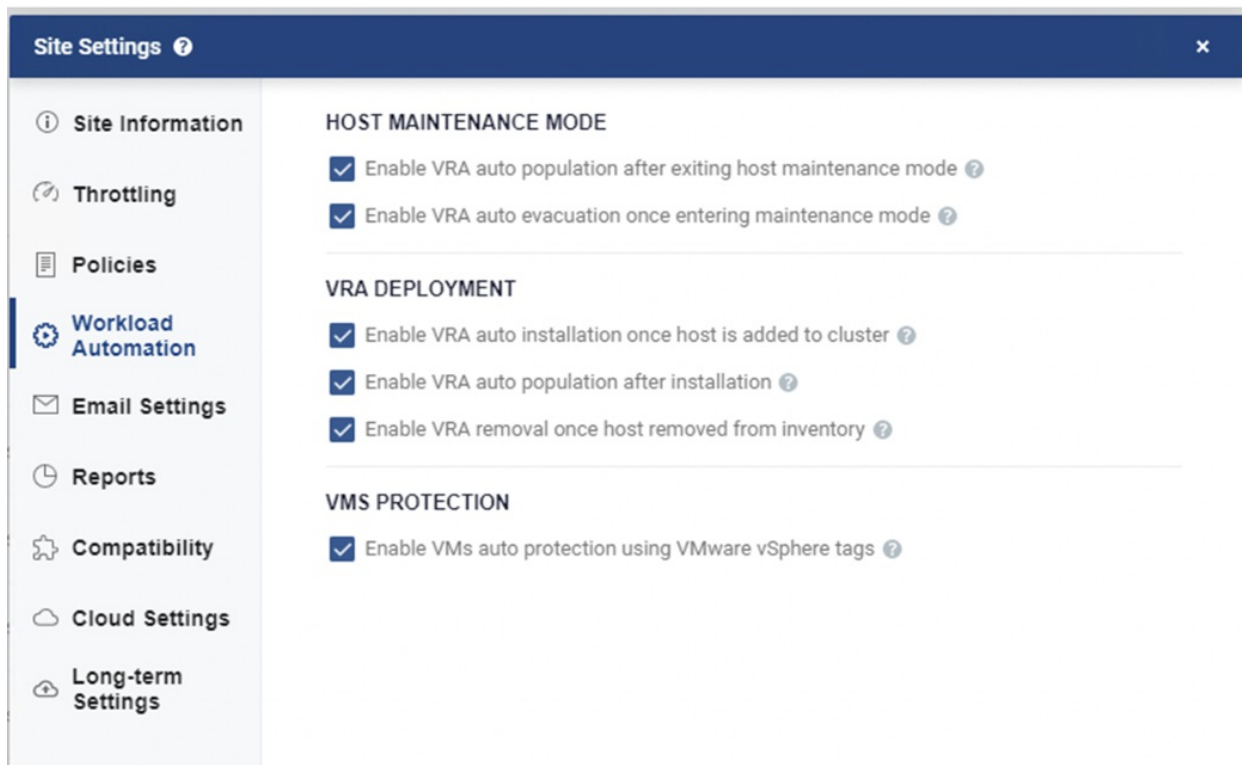
Currently, Zerto disaster recovery on Azure VMware Solution is in an Initial Availability (IA) phase. In the IA phase, you must contact Microsoft to request and qualify for IA support.

To request IA support for Zerto on Azure VMware Solution, submit this [Install Zerto on AVS form](#) with the required information. In the IA phase, Azure VMware Solution only supports manual installation and onboarding of Zerto. However, Microsoft will work with you to ensure that you can manually install Zerto on your private cloud.

NOTE

As part of the manual installation, Microsoft creates a new vCenter user account for Zerto. This user account is only for Zerto Virtual Manager (ZVM) to perform operations on the Azure VMware Solution vCenter. When installing ZVM on Azure VMware Solution, don't select the "Select to enforce roles and permissions using Zerto vCenter privileges" option.

After the ZVM installation, select the options below from the Zerto Virtual Manager **Site Settings**.



NOTE

General Availability of Azure VMware Solution will enable self-service installation and Day 2 operations of Zerto on Azure VMware Solution.

Configure Zerto for disaster recovery

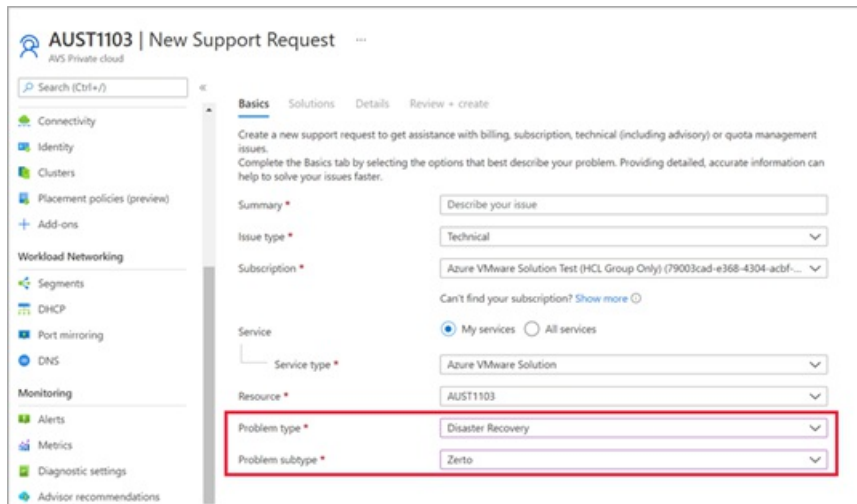
To configure Zerto for the on-premises VMware to Azure VMware Solution disaster recovery and Azure VMware

Solution to Azure VMware Solution Cloud disaster recovery scenarios, see the [Zerto Virtual Manager Administration Guide vSphere Environment](#).

For more information, see the [Zerto technical documentation](#).

Ongoing management of Zerto

- As you scale your Azure VMware Solution private cloud operations, you might need to add new Azure VMware Solution hosts for Zerto protection or configure Zerto disaster recovery to new Azure VMware Solution vSphere Clusters. In both these scenarios, you'll be required to open a Support Request with the Azure VMware Solution team in the Initial Availability phase. You can open the [support ticket](#) from the Azure portal for these Day 2 configurations.



The screenshot shows the 'New Support Request' form in the Azure portal for subscription AUST1103. The form is titled 'AUST1103 | New Support Request' and is for the 'AUS Private cloud'. The 'Basics' tab is selected, and the form is partially filled out. The 'Problem type' is set to 'Disaster Recovery' and the 'Problem subtype' is set to 'Zerto'. These two fields are highlighted with a red box. Other fields include 'Summary', 'Issue type' (Technical), 'Subscription' (Azure VMware Solution Test (HCL Group Only) (79003cad-e368-4304-acbf-...)), 'Service' (My services), 'Service type' (Azure VMware Solution), and 'Resource' (AUST1103).

- In the GA phase, all the above operations will be enabled in an automated self-service fashion.

FAQs

Can I use a pre-existing Zerto product license on Azure VMware Solution?

You can reuse pre-existing Zerto product licenses for Azure VMware Solution environments. If you need new Zerto licenses, email Zerto at info@zerto.com to acquire new licenses.

How is Zerto supported?

Zerto disaster recovery is a solution that is sold and supported by Zerto. For any support issue with Zerto disaster recovery, always contact [Zerto support](#).

Zerto and Microsoft support teams will engage each other as needed to troubleshoot Zerto disaster recovery issues on Azure VMware Solution.

Deploy Horizon on Azure VMware Solution

12/16/2022 • 11 minutes to read • [Edit Online](#)

NOTE

This document focuses on the VMware Horizon product, formerly known as Horizon 7. Horizon is a different solution than Horizon Cloud on Azure, although there are some shared components. Key advantages of the Azure VMware Solution include both a more straightforward sizing method and the integration of VMware Cloud Foundation management into the Azure portal.


VMware Horizon®, a virtual desktop and applications platform, runs in the data center and provides simple and centralized management. It delivers virtual desktops and applications on any device, anywhere. Horizon lets you create, and broker connections to Windows and Linux virtual desktops, Remote Desktop Server (RDS) hosted applications, desktops, and physical machines.

Here, we focus specifically on deploying Horizon on Azure VMware Solution. For general information on VMware Horizon, refer to the Horizon production documentation:

- [What is VMware Horizon?](#)
- [Learn more about VMware Horizon](#)
- [Horizon Reference Architecture](#)

With Horizon's introduction on Azure VMware Solution, there are now two Virtual Desktop Infrastructure (VDI) solutions on the Azure platform. The following diagram summarizes the key differences at a high level.

Differences between VMware Horizon on Azure VMware Solution and VMware Horizon Cloud on Azure



Horizon 2006 and later versions on the Horizon 8 release line supports both on-premises and Azure VMware Solution deployment. There are a few Horizon features that are supported on-premises but not on Azure VMware Solution. Other products in the Horizon ecosystem are also supported. For more information, see [feature parity and interoperability](#).

Deploy Horizon in a hybrid cloud

You can deploy Horizon in a hybrid cloud environment by using Horizon Cloud Pod Architecture (CPA) to interconnect on-premises and Azure data centers. CPA scales up your deployment, builds a hybrid cloud, and provides redundancy for Business Continuity and Disaster Recovery. For more information, see [Expanding Existing Horizon 7 Environments](#).

IMPORTANT

CPA is not a stretched deployment; each Horizon pod is distinct, and all Connection Servers that belong to each of the individual pods are required to be located in a single location and run on the same broadcast domain from a network perspective.

Like on-premises or private data centers, you can deploy Horizon in an Azure VMware Solution private cloud. We'll discuss key differences in deploying Horizon on-premises and Azure VMware Solution in the following sections.

The *Azure private cloud* is conceptually the same as the *VMware SDDC*, a term typically used in Horizon documentation. The rest of this document uses both terms interchangeably.

The Horizon Cloud Connector is required for Horizon on Azure VMware Solution to manage subscription licenses. You can deploy Cloud Connector in Azure Virtual Network alongside Horizon Connection Servers.

IMPORTANT

Horizon Control Plane support for Horizon on Azure VMware Solution is not yet available. Be sure to download the VHD version of Horizon Cloud Connector.

vCenter Server Cloud Admin role

Since Azure VMware Solution is an SDDC service and Azure manages the lifecycle of the SDDC on Azure VMware Solution, the vCenter Server permission model on Azure VMware Solution is limited by design.

Customers are required to use the Cloud Admin role, which has a limited set of vCenter Server permissions. The Horizon product was modified to work with the Cloud Admin role on Azure VMware Solution, specifically:

- Instant clone provisioning was modified to run on Azure VMware Solution.
- A specific vSAN policy (VMware_Horizon) was created on Azure VMware Solution to work with Horizon, which must be available and used in the SDDCs deployed for Horizon.
- vSphere Content-Based Read Cache (CBRC), also known as View Storage Accelerator, is disabled when running on the Azure VMware Solution.

IMPORTANT

CBRC must not be turned back on.

NOTE

Azure VMware Solution automatically configures specific Horizon settings as long as you deploy Horizon 2006 (aka Horizon 8) and above on the Horizon 8 branch and select the **Azure** option in the Horizon Connection Server installer.

Horizon on Azure VMware Solution deployment architecture

A typical Horizon architecture design uses a pod and block strategy. A block is a single vCenter Server, while multiple blocks combined make a pod. A Horizon pod is a unit of organization determined by Horizon scalability limits. Each Horizon pod has a separate management portal, and so a standard design practice is to minimize the number of pods.

Every cloud has its own network connectivity scheme. Combined with VMware SDDC networking / NSX-T Data Center, the Azure VMware Solution network connectivity presents unique requirements for deploying Horizon that is different from on-premises.

Each Azure private cloud and SDDC can handle 4,000 desktop or application sessions, assuming:

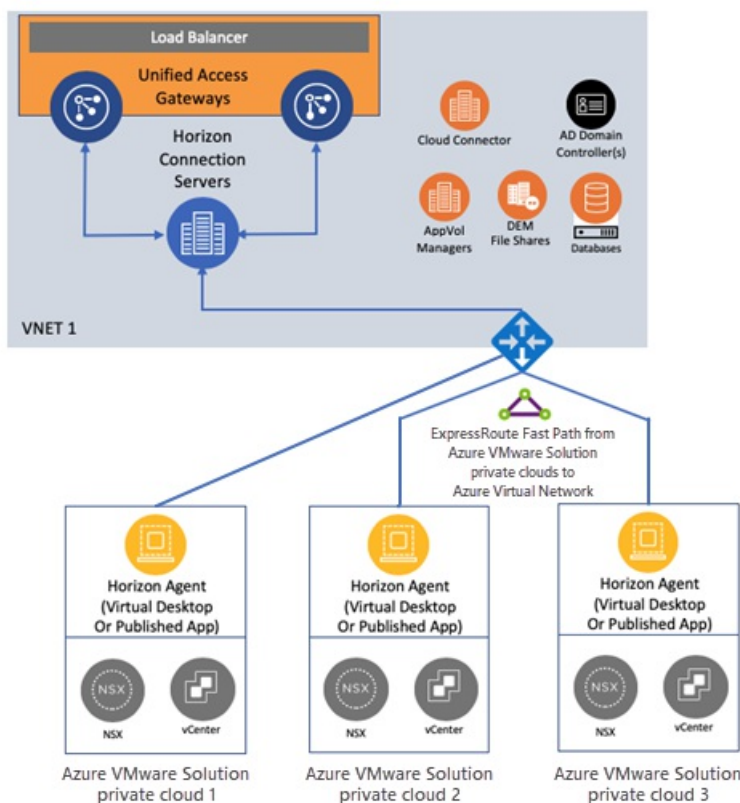
- The workload traffic aligns with the LoginVSI task worker profile.
- Only protocol traffic is considered, no user data.
- NSX Edge is configured to be large.

NOTE

Your workload profile and needs may be different, and therefore results may vary based on your use case. User Data volumes may lower scale limits in the context of your workload. Size and plan your deployment accordingly. For more information, see the sizing guidelines in the [Size Azure VMware Solution hosts for Horizon deployments](#) section.

Given the Azure private cloud and SDDC max limit, we recommend a deployment architecture where the Horizon Connection Servers and VMware Unified Access Gateways (UAGs) are running inside the Azure Virtual Network. It effectively turns each Azure Desktop private cloud and SDDC into a block. In turn, maximizing the scalability of Horizon running on Azure VMware Solution.

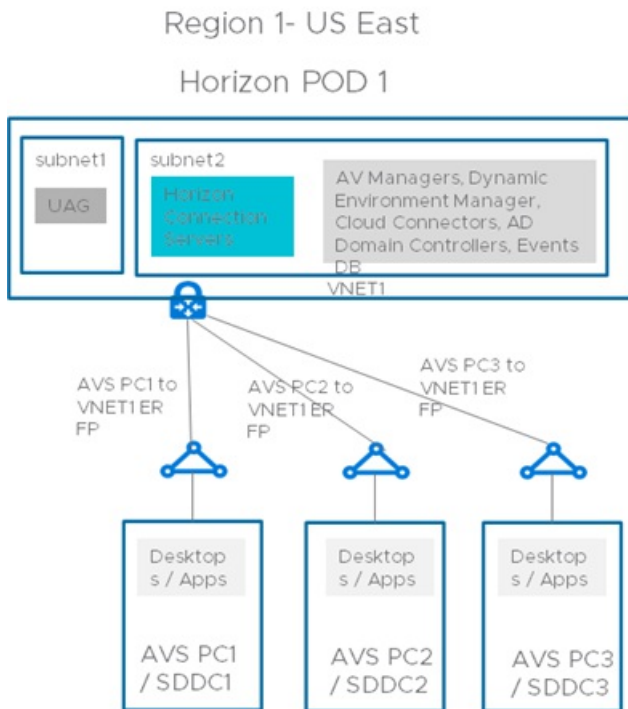
The connection from Azure Virtual Network to the Azure private clouds / SDDCs should be configured with ExpressRoute FastPath. The following diagram shows a basic Horizon pod deployment.



Network connectivity to scale Horizon on Azure VMware Solution

This section lays out the network architecture at a high level with some common deployment examples to help you scale Horizon on Azure VMware Solution. The focus is specifically on critical networking elements.

Single Horizon pod on Azure VMware Solution



A single Horizon pod is the most straight forward deployment scenario because you deploy just one Horizon pod in the US East region. Since each private cloud and SDDC is estimated to handle 4,000 desktop sessions, you deploy the maximum Horizon pod size. You can plan the deployment of up to three private clouds/SDDCs.

With the Horizon infrastructure virtual machines (VMs) deployed in Azure Virtual Network, you can reach the 12,000 sessions per Horizon pod. The connection between each private cloud and SDDC to the Azure Virtual Network is ExpressRoute Fast Path. No east-west traffic between private clouds is needed.

Key assumptions for this basic deployment example include that:

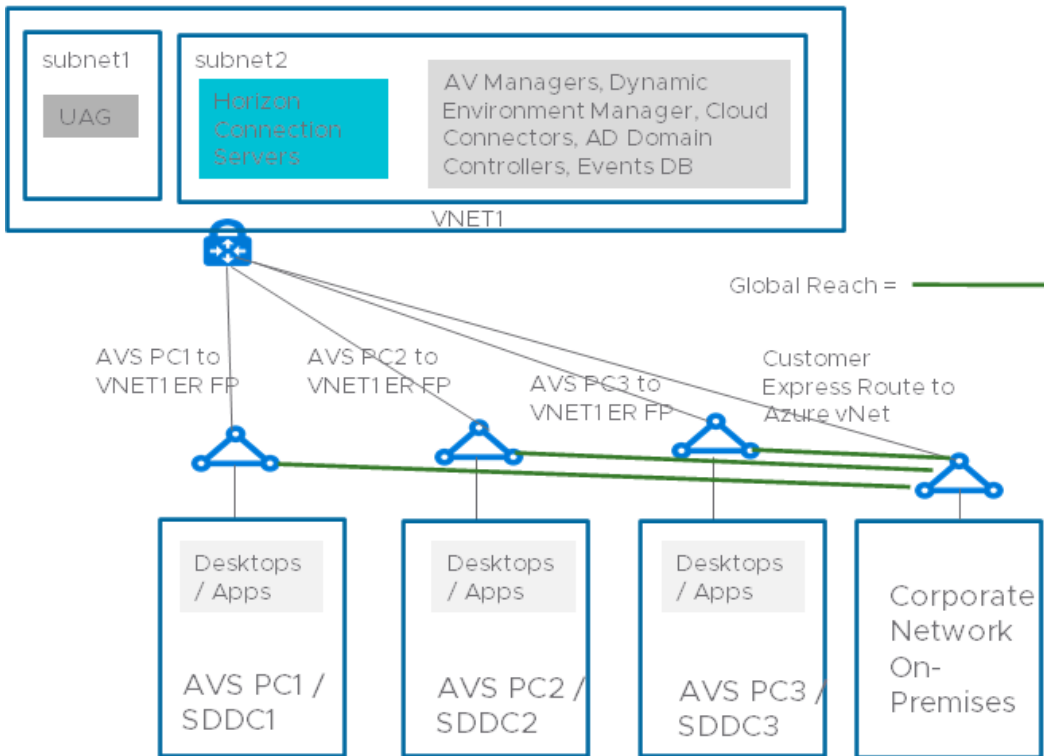
- You don't have an on-premises Horizon pod that you want to connect to this new pod using Cloud Pod Architecture (CPA).
- End users connect to their virtual desktops through the internet (vs. connecting via an on-premises datacenter).

You connect your AD domain controller in Azure Virtual Network with your on-premises AD through VPN or ExpressRoute circuit.

A variation on the basic example might be to support connectivity for on-premises resources. For example, users access desktops and generate virtual desktop application traffic or connect to an on-premises Horizon pod using CPA.

The diagram shows how to support connectivity for on-premises resources. To connect to your corporate network to the Azure Virtual Network, you'll need an ExpressRoute circuit. You'll also need to connect your corporate network with each of the private cloud and SDDCs using ExpressRoute Global Reach. It allows the connectivity from the SDDC to the ExpressRoute circuit and on-premises resources.

Region 1- US East Horizon POD 1




Multiple Horizon pods on Azure VMware Solution across multiple regions

Another scenario is scaling Horizon across multiple pods. In this scenario, you deploy two Horizon pods in two different regions and federate them using CPA. It's similar to the network configuration in the previous example, but with some additional cross-regional links.

You'll connect the Azure Virtual Network in each region to the private clouds/SDDCs in the other region. It allows Horizon connection servers part of the CPA federation to connect to all desktops under management. Adding extra private clouds/SDDCs to this configuration would allow you to scale to 24,000 sessions overall.

The same principles apply if you deploy two Horizon pods in the same region. Make sure to deploy the second Horizon pod in a *separate Azure Virtual Network*. Just like the single pod example, you can connect your corporate network and on-premises pod to this multi-pod/region example using ExpressRoute and Global Reach.



Size Azure VMware Solution hosts for Horizon deployments

Horizon's sizing methodology on a host running in Azure VMware Solution is simpler than Horizon on-premises. That's because the Azure VMware Solution host is standardized. Exact host sizing helps determine the number of hosts needed to support your VDI requirements. It's central to determining the cost-per-desktop.

Sizing tables

Specific vCPU/vRAM requirements for Horizon virtual desktops depend on the customer's specific workload profile. Work with your MSFT and VMware sales team to help determine your vCPU/vRAM requirements for your virtual desktops.

V C P U P E R F O R M A N C E	V R A M P E R F O R M A N C E (G B)	IN S T A N C E	10	20	30	40	50	60	70	80	90	100	200	300	400	500	600	6400
			V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S
2	3.5	A V S	3	3	4	4	5	6	6	7	8	9	17	25	33	41	49	53
2	4	A V S	3	3	4	5	6	6	7	8	9	9	18	26	34	42	51	54
2	6	A V S	3	4	5	6	7	9	10	11	12	13	26	38	50	62	74	79
2	8	A V S	3	5	6	8	9	11	12	14	16	18	34	51	67	84	100	106
2	12	A V S	4	6	9	11	13	16	19	22	26	30	55	82	109	136	163	171
2	16	A V S	5	8	11	14	18	22	27	32	37	43	76	113	150	187	224	234
4	3.5	A V S	3	3	4	5	6	7	8	9	10	11	22	33	44	55	66	70
4	4	A V S	3	3	4	5	6	7	8	9	10	11	22	33	44	55	66	70
4	6	A V S	3	4	5	6	7	9	10	11	12	13	26	38	50	62	74	79

V C P U P E R V M	V R A M P E R V M (G B)	I N S T A N C E	10	20	30	40	50	60	70	80	90	10	20	30	40	50	60	64
			0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	00 V M S	00 V M S	00 V M S	00 V M S	00 V M S
4	8	A V S	3	5	6	8	9	1 1	1 2	1 4	1 6	1 8	3 4	5 1	6 7	8 4	1 0 0	1 0 6
4	1 2	A V S	4	6	9	1 1	1 3	1 6	1 9	2 1	2 3	2 6	5 1	7 5	1 0 0	1 2 4	1 4 9	1 5 8
4	1 6	A V S	5	8	1 1	1 4	1 8	2 1	2 4	2 7	3 0	3 4	6 7	1 0 0	1 3 3	1 6 5	1 9 8	2 1 1
6	3. 5	A V S	3	4	5	6	7	9	1 0	1 1	1 3	1 4	2 7	4 1	5 4	6 8	8 1	8 6
6	4	A V S	3	4	5	6	7	9	1 0	1 1	1 3	1 4	2 7	4 1	5 4	6 8	8 1	8 6
6	6	A V S	3	4	5	6	7	9	1 0	1 1	1 3	1 4	2 7	4 1	5 4	6 8	8 1	8 6
6	8	A V S	3	5	6	8	9	1 1	1 2	1 4	1 6	1 8	3 4	5 1	6 7	8 4	1 0 0	1 0 6
6	1 2	A V S	4	6	9	1 1	1 3	1 6	1 9	2 1	2 3	2 6	5 1	7 5	1 0 0	1 2 4	1 4 9	1 5 8
6	1 6	A V S	5	8	1 1	1 4	1 8	2 1	2 4	2 7	3 0	3 4	6 7	1 0 0	1 3 3	1 6 5	1 9 8	2 1 1
8	3. 5	A V S	3	4	6	7	9	1 0	1 2	1 4	1 5	1 7	3 3	4 9	6 6	8 2	9 8	1 0 5
8	4	A V S	3	4	6	7	9	1 0	1 2	1 4	1 5	1 7	3 3	4 9	6 6	8 2	9 8	1 0 5
8	6	A V S	3	4	6	7	9	1 0	1 2	1 4	1 5	1 7	3 3	4 9	6 6	8 2	9 8	1 0 5

V C P U P E R V M	R A M P E S (G B)	I N S T A L L E D	10	20	30	40	50	60	70	80	90	10	20	30	40	50	60	64
			0	0	0	0	0	0	0	0	0	0	0	00	00	00	00	00
8	8	A V S	3	5	6	8	9	1	1	1	1	1	3	5	6	8	1	1
8	1 2	A V S	4	6	9	1	1	1	1	2	2	2	5	7	1	1	1	1
8	1 6	A V S	5	8	1	1	1	2	2	2	3	3	6	1	1	1	1	2

Horizon sizing inputs

Here's what you'll need to gather for your planned workload:

- Number of concurrent desktops
- Required vCPU per desktop
- Required vRAM per desktop
- Required storage per desktop

In general, VDI deployments are either CPU or RAM constrained, which determines the host size. Let's take the following example for a LoginVSI Knowledge Worker type of workload, validated with performance testing:

- 2,000 concurrent desktop deployment
- 2vCPU per desktop.
- 4-GB vRAM per desktop.
- 50 GB of storage per desktop

For this example, the total number of hosts factors out to 18, yielding a VM-per-host density of 111.

IMPORTANT

Customer workloads will vary from this example of a LoginVSI Knowledge Worker. As a part of planning your deployment, work with your VMware EUC SEs for your specific sizing and performance needs. Be sure to run your own performance testing using the actual, planned workload before finalizing host sizing and adjust accordingly.

Horizon on Azure VMware Solution licensing

There are four components to the overall costs of running Horizon on Azure VMware Solution.

Azure VMware Solution Capacity Cost

For information on the pricing, see the [Azure VMware Solution pricing](#) page

Horizon Licensing Cost

There are two available licenses for use with the Azure VMware Solution, which can be either Concurrent User (CCU) or Named User (NU):

- Horizon Subscription License
- Horizon Universal Subscription License

If only deploying Horizon on Azure VMware Solution for the foreseeable future, then use the Horizon Subscription License as it is a lower cost.

If deployed on Azure VMware Solution and on-premises, choose the Horizon Universal Subscription License as a disaster recovery use case. However, it includes a vSphere license for on-premises deployment, so it has a higher cost.

Work with your VMware EUC sales team to determine the Horizon licensing cost based on your needs.

Azure Instance Types

To understand the Azure virtual machine sizes that are required for the Horizon Infrastructure, see [Horizon Installation on Azure VMware Solution](#).

References

[System Requirements For Horizon Agent for Linux](#)

Next steps

To learn more about VMware Horizon on Azure VMware Solution, read the [VMware Horizon FAQ](#).

Deploy Citrix on Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

Citrix Virtual Apps and Desktop service supports Azure VMware Solution. Azure VMware Solution provides cloud infrastructure containing vSphere clusters created by Azure infrastructure. You can leverage the Citrix Virtual Apps and Desktop Service to use Azure VMware Solution for provisioning your [Virtual Delivery Agent \(VDA\)](#) workload in the same way you would using vSphere in on-premises environments.

[Learn more about Citrix virtual apps and desktops](#)

[Deployment guide](#)

[Solution brief](#)

FAQ (review Q&As)

- Q. Can I migrate my existing Citrix desktops and apps to Azure VMware Solution, or operate a hybrid environment that consists of on-premises and Azure VMware Solution-based Citrix workloads?

A. Yes. You can use the same machine images, application packages, and processes you currently use. You're able to seamlessly link on-premises and Azure VMware Solution-based environments together for a migration.
- Q. Can Citrix be deployed as a standalone environment within Azure VMware Solution?

A. Yes. You're free to migrate, operate a hybrid environment, or deploy a standalone directly into Azure VMware Solution.
- Q. Does Azure VMware Solution support both PVS and MCS?

A. Yes.
- Q. Are GPU-based workloads supported in Citrix on Azure VMware Solution?

A. Not at this time. However, Citrix workloads on Microsoft Azure support GPU if that use case is important to you.
- Q. Is Azure VMware Solution supported with on-premises Citrix deployments or LTSR?

A. No. Azure VMware Solution is only supported with the Citrix Virtual Apps and Desktops service offerings.
- Q. Who do I call for support?

A. Customers should contact Citrix support www.citrix.com/support for assistance.
- Q. Can I use my Azure Virtual Desktop benefit from Microsoft with Citrix on Azure VMware Solution?

A. No. Azure Virtual Desktop benefits are applicable to native Microsoft Azure workloads only. Citrix Virtual Apps and Desktops service, as a native Azure offering, can apply your Azure Virtual Desktop benefit alongside your Azure VMware Solution deployment.
- Q. How do I purchase Citrix Virtual Apps and Desktops service to use Azure VMware Solution?

A. You can purchase Citrix offerings via your Citrix partner or directly from the Azure Marketplace.

Deploy vSAN stretched clusters (Preview)

12/16/2022 • 9 minutes to read • [Edit Online](#)

In this article, you'll learn how to implement a vSAN stretched cluster for an Azure VMware Solution private cloud.

Background

Azure's global infrastructure is broken up into Regions. Each region supports the services for a given geography. Within each region, Azure builds isolated, and redundant islands of infrastructure called availability zones (AZ). An AZ acts as a boundary for resource management. The compute and other resources available to an AZ are finite and may become exhausted by customer demands. An AZ is built to be independently resilient, meaning failures in one AZ doesn't affect other AZs.

With Azure VMware Solution, ESXi hosts deployed in a standard vSphere cluster traditionally reside in a single Azure Availability Zone (AZ) and are protected by vSphere high availability (HA). However, it doesn't protect the workloads against an Azure AZ failure. To protect against an AZ failure, a single vSAN cluster can be enabled to span two separate availability zones, called a **vSAN stretched cluster**.

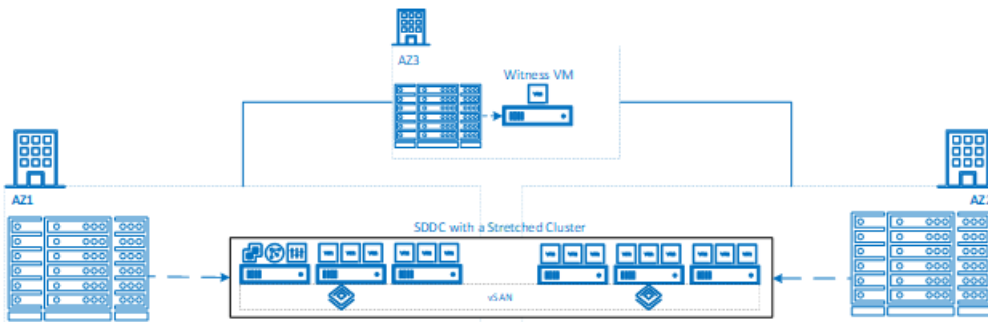
Stretched clusters allow the configuration of vSAN Fault Domains across two AZs to notify vCenter Server that hosts reside in each Availability Zone (AZ). Each fault domain is named after the AZ it resides within to increase clarity. When you stretch a vSAN cluster across two AZs within a region, should an AZ go down, it's treated as a vSphere HA event and the virtual machine is restarted in the other AZ.

Stretched cluster benefits:

- Improve application availability.
- Provide a zero recovery point objective (RPO) capability for enterprise applications without needing to redesign them, or deploy expensive disaster recovery (DR) solutions.
- A private cloud with stretched clusters is designed to provide 99.99% availability due to its resilience to AZ failures.
- Enable customers to focus on core application requirements and features, instead of infrastructure availability.

To protect against split-brain scenarios and help measure site health, a managed vSAN Witness is created in a third AZ. With a copy of the data in each AZ, vSphere HA attempts to recover from any failure using a simple restart of the virtual machine.

The following diagram depicts a vSAN cluster stretched across two AZs.

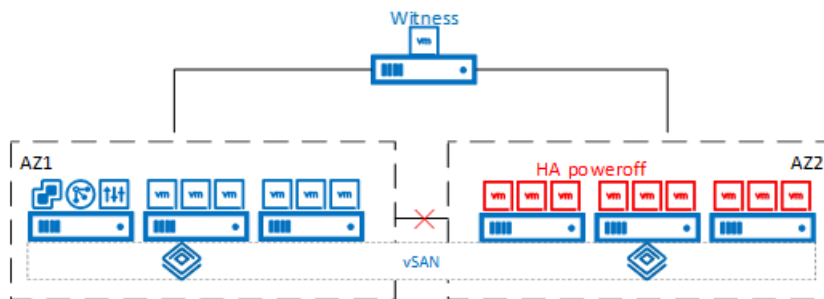


In summary, stretched clusters simplify protection needs by providing the same trusted controls and capabilities in addition to the scale and flexibility of the Azure infrastructure.

It's important to understand that stretched cluster private clouds only offer an extra layer of resiliency, and they don't address all failure scenarios. For example, stretched cluster private clouds:

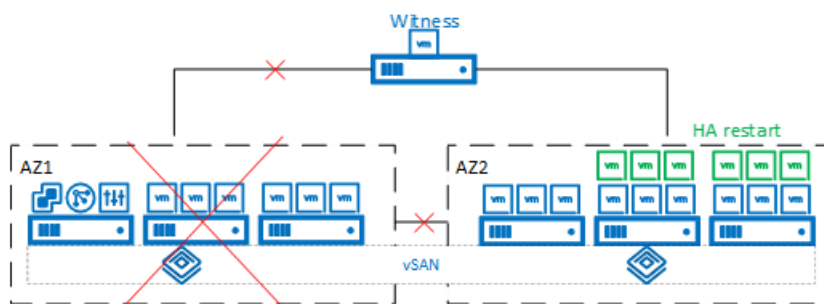
- Don't protect against region-level failures within Azure or data loss scenarios caused by application issues or poorly planned storage policies.
- Provides protection against a single zone failure but aren't designed to provide protection against double or progressive failures. For example:
 - Despite various layers of redundancy built into the fabric, if an inter-AZ failure results in the partitioning of the secondary site, vSphere HA starts powering off the workload VMs on the secondary site.

The following diagram shows the secondary site partitioning scenario.



- If the secondary site partitioning progressed into the failure of the primary site instead, or resulted in a complete partitioning, vSphere HA would attempt to restart the workload VMs on the secondary site. If vSphere HA attempted to restart the workload VMs on the secondary site, it would put the workload VMs in an unsteady state.

The following diagram shows the preferred site failure or complete partitioning scenario.



It should be noted that these types of failures, although rare, fall outside the scope of the protection offered by a stretched cluster private cloud. Because of those types of rare failures, a stretched cluster solution should be regarded as a multi-AZ high availability solution reliant upon vSphere HA. It's important you understand that a stretched cluster solution isn't meant to replace a comprehensive multi-region Disaster Recovery strategy that can be employed to ensure application availability. The reason is because a Disaster Recovery solution typically has separate management and control planes in separate Azure regions. Azure VMware Solution stretched clusters have a single management and control plane stretched across two availability zones within the same Azure region. For example, one vCenter Server, one NSX-T Manager cluster, one NSX-T Data Center Edge VM pair.

Deploy a stretched cluster private cloud

Currently, Azure VMware Solution stretched clusters is in the (preview) phase. While in the (preview) phase, you must contact Microsoft to request and qualify for support.

Prerequisites

To request support, send an email request to avsStretchedCluster@microsoft.com with the following details:

- Company name
- Point of contact (email)
- Subscription (a new, separate subscription is required)
- Region requested (West Europe, UK South, Germany West Central)
- Number of nodes in first stretched cluster (minimum 6, maximum 16 - in multiples of two)
- Estimated provisioning date (used for billing purposes)

When the request support details are received, quota will be reserved for a stretched cluster environment in the region requested. The subscription gets enabled to deploy a stretched cluster SDDC through the Azure portal. A confirmation email will be sent to the designated point of contact within two business days upon which you should be able to [self-deploy a stretched cluster private cloud via the Azure portal](#). Be sure to select **Hosts in two availability zones** to ensure that a stretched cluster gets deployed in the region of your choice.



Create a private cloud ...

Prerequisites * **Basics** Tags Review and Create

Project details

Subscription * ⓘ

Resource group * ⓘ [Create new](#)

Private cloud details

Resource name * ⓘ ✓

Location * ⓘ

Size of host * ⓘ

Host location *

All hosts in one availability zone

Hosts in two availability zones
Hosts will be equally divided across 2 availability zones. Since there will be two availability zones, the number of hosts you can select are in multiples of 2 only.

Number of hosts * ⓘ 6

[Find out how many hosts you need if you need more hosts, request a quota increase](#)

estimated monthly total

CIDR address block

Provide IP address for private cloud for cluster management. Make sure these are unique and do not overlap with any other Azure vnets or on-premise networks.

Address block for private cloud * ⓘ ✓

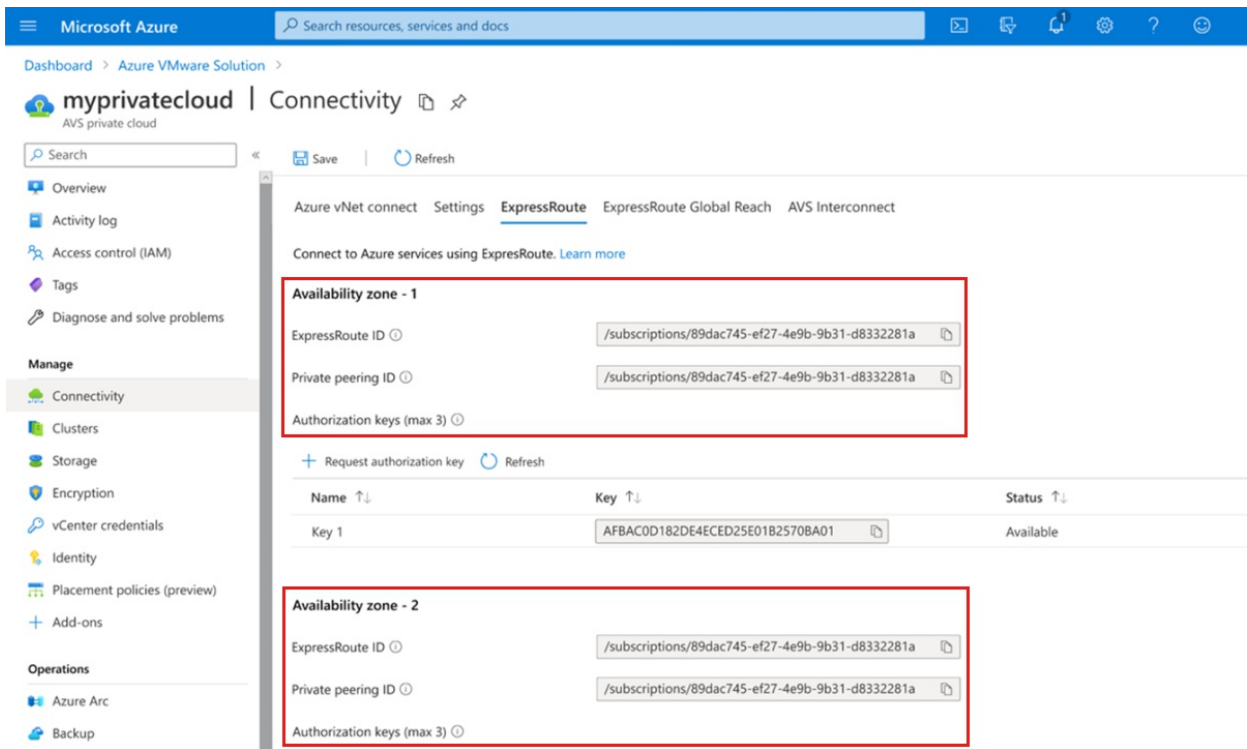
- ⓘ The address block must fall within the following allowed network blocks: 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16
- ⓘ The address block cannot overlap any of the following restricted network blocks: 172.17.0.0/16
- ⓘ The address block cannot be smaller than a /22 network.

[Review and Create](#)

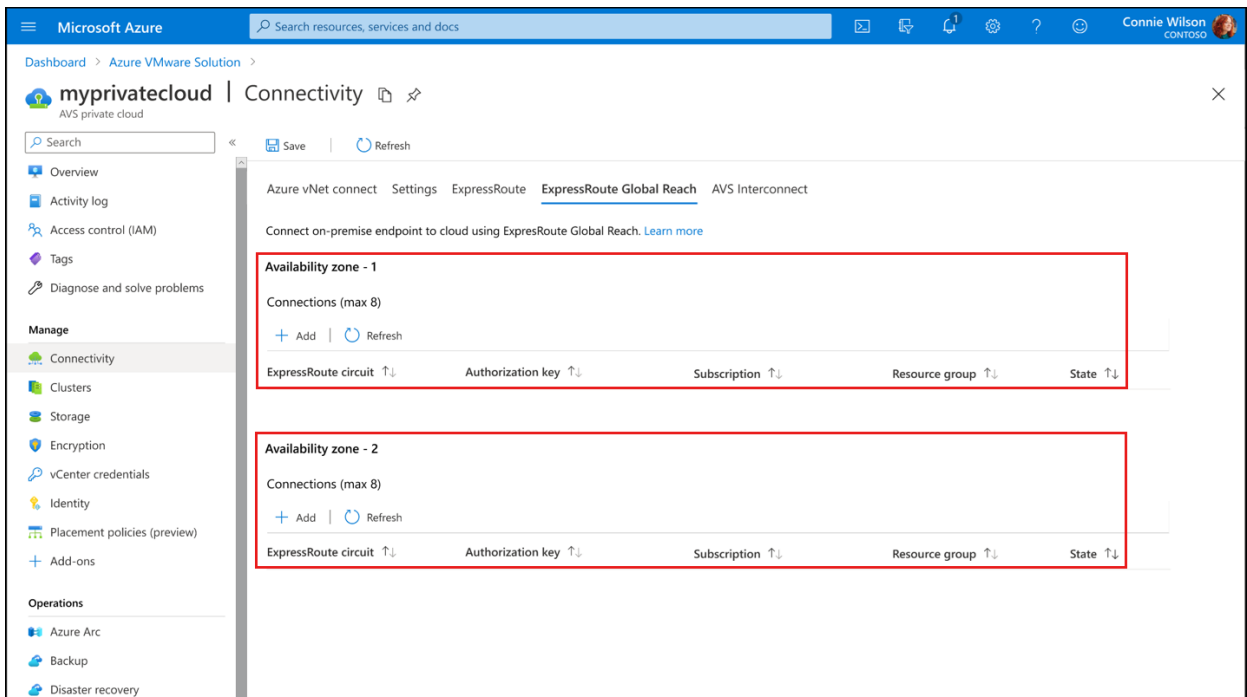
[Previous](#)

[Next : Tags >](#)

Once the private cloud is created, you can peer both availability zones (AZs) to your on-premises ExpressRoute circuit with Global Reach that helps connect your on-premises data center to the private cloud. Peering both the AZs will ensure that an AZ failure doesn't result in a loss of connectivity to your private cloud. Since an ExpressRoute Auth Key is valid for only one connection, repeat the [Create an ExpressRoute auth key in the on-premises ExpressRoute circuit](#) process to generate another authorization.



Next, repeat the process to [peer ExpressRoute Global Reach](#) two availability zones to the on-premises ExpressRoute circuit.



Supported scenarios

The following scenarios are supported:

- Workload connectivity to internet from both AZs via Customer vWAN or On-premises data center
- Private DNS resolution
- Placement policies (except for VM-AZ affinity)
- Cluster scale out and scale in
- The following SPBM policies are supported, with a PFTT of "Dual Site Mirroring" and SFTT of "RAID 1 (Mirroring)" enabled as the default policies for the cluster:
 - Site disaster tolerance settings (PFTT):
 - Dual site mirroring

- Dual-site mirroring
 - o None - keep data on preferred
 - o None - keep data on non-preferred
- o Local failures to tolerate (SFTT):
 - o 1 failure – RAID 1 (Mirroring)
 - o 1 failure – RAID 5 (Erasure coding), requires a minimum of 4 hosts in each AZ
 - o 2 failures – RAID 1 (Mirroring)
 - o 2 failures – RAID 6 (Erasure coding), requires a minimum of 6 hosts in each AZ
 - o 3 failures – RAID 1 (Mirroring)

In this phase, while the creation of the private cloud and the first stretched cluster is enabled via the Azure portal, open a [support ticket](#) from the Azure portal for other supported scenarios and configurations listed below. While doing so, make sure you select **Stretched Clusters** as a Problem Type.

Once stretched clusters are made generally available, it's expected that all the following supported scenarios will be enabled in an automated self-service fashion.

- HCX installation, deployment, removal, and support for migration
- Connect a private cloud in another region to a stretched cluster private cloud
- Connect two stretched cluster private clouds in a single region
- Configure Active Directory as an identity source for vCenter Server
- A PFTT of "Keep data on preferred" or "Keep data on non-preferred" requires keeping VMs on either one of the availability zones. For such VMs, open a support ticket to ensure that those VMs are pinned to an availability zone.
- Cluster addition
- Cluster deletion
- Private cloud deletion

Supported regions

Azure VMware Solution stretched clusters are available in the following regions:

- UK South
- West Europe
- Germany West Central

FAQ

Are any other regions planned?

As of now, the only 3 regions listed above are planned for support of stretched clusters.

What kind of SLA does Azure VMware Solution provide with the stretched clusters (preview) release?

A private cloud created with a vSAN stretched cluster is designed to offer a 99.99% infrastructure availability commitment when the following conditions exist:

- A minimum of 6 nodes are deployed in the cluster (3 in each availability zone)
- When a VM storage policy of PFTT of "Dual-Site Mirroring" and an SFTT of 1 is used by the workload VMs
- Compliance with the **Additional Requirements** captured in the [SLA details of Azure VMware Solution](#) is required to achieve the availability goals

Do I get to choose the availability zone in which a private cloud is deployed?

No. A stretched cluster is created between two availability zones, while the third zone is used for deploying the witness node. Because all of the zones are effectively used for deploying a stretched cluster environment, a

choice isn't provided to the customer. Instead, the customer chooses to deploy hosts in multiple AZs at the time of private cloud creation.

What are the limitations I should be aware of?

- Once a private cloud has been created with a stretched cluster, it can't be changed to a standard cluster private cloud. Similarly, a standard cluster private cloud can't be changed to a stretched cluster private cloud after creation.
- Scale out and scale-in of stretched clusters can only happen in pairs. A minimum of 6 nodes and a maximum of 16 nodes are supported in a stretched cluster environment.
- Customer workload VMs are restarted with a medium vSphere HA priority. Management VMs have the highest restart priority.
- The solution relies on vSphere HA and vSAN for restarts and replication. Recovery time objective (RTO) is determined by the amount of time it takes vSphere HA to restart a VM on the surviving AZ after the failure of a single AZ.
- Preview and recent GA features for standard private cloud environments aren't supported in a stretched cluster environment.
- Disaster recovery add-ons like, VMware SRM, Zerto, and JetStream are currently not supported in a stretched cluster environment.

What kind of latencies should I expect between the availability zones (AZs)?

vSAN stretched clusters operate within a 5-millisecond round trip time (RTT) and 10 Gb/s or greater bandwidth between the AZs that host the workload VMs. The Azure VMware Solution stretched cluster deployment follows that guiding principle. Consider that information when deploying applications (with SFTT of dual site mirroring, which uses synchronous writes) that have stringent latency requirements.

Can I mix stretched and standard clusters in my private cloud?

No. A mix of stretched and standard clusters aren't supported within the same private cloud. A stretched or standard cluster environment is selected when you create the private cloud. Once a private cloud has been created with a stretched cluster, it's assumed that all clusters created within that private cloud are stretched in nature.

How much does the solution cost?

Customers will be charged based on the number of nodes deployed within the private cloud.

Will I be charged for the witness node and for inter-AZ traffic?

No. While in (preview), customers won't see a charge for the witness node and the inter-AZ traffic. The witness node is entirely service managed, and Azure VMware Solution provides the required lifecycle management of the witness node. As the entire solution is service managed, the customer only needs to identify the appropriate SPBM policy to set for the workload virtual machines. The rest is managed by Microsoft.

Which SKUs are available?

Stretched clusters will solely be supported on the AV36 SKU.

Move Azure VMware Solution subscription to another subscription

12/16/2022 • 2 minutes to read • [Edit Online](#)

This article describes how to move an Azure VMware Solution subscription to another subscription. You might move your subscription for various reasons, like billing.

Prerequisites

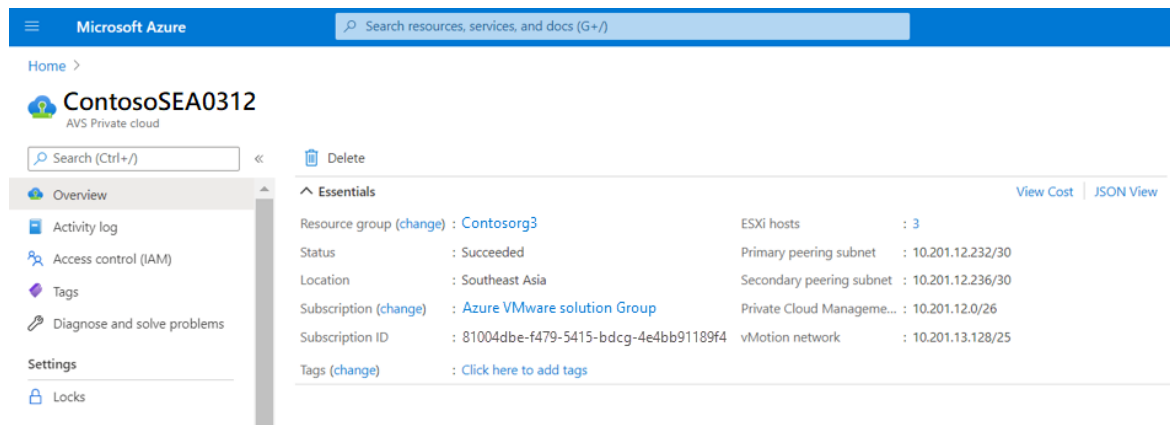
You should have at least contributor rights on both **source** and **target** subscriptions.

IMPORTANT

VNet and VNet gateway can't be moved from one subscription to another. Additionally, moving your subscriptions has no impact on the management and workloads, like the vCenter, NSX, and workload virtual machines.

Prepare and move

1. In the Azure portal, select the private cloud you want to move.



The screenshot shows the Azure portal interface for a private cloud. The top navigation bar includes the Microsoft Azure logo and a search bar. Below the navigation bar, the page title is "ContosoSEA0312" with a sub-label "AVS Private cloud". A search bar and a "Delete" button are visible. The left sidebar contains navigation options: Overview, Activity log, Access control (IAM), Tags, Diagnose and solve problems, Settings, and Locks. The main content area displays the "Essentials" section, which includes a table of key properties:

Property	Value
Resource group (change)	Contosorg3
Status	Succeeded
Location	Southeast Asia
Subscription (change)	Azure VMware solution Group
Subscription ID	81004dbe-f479-5415-bdcg-4e4bb91189f4
Tags (change)	Click here to add tags
ESXi hosts	3
Primary peering subnet	10.201.12.232/30
Secondary peering subnet	10.201.12.236/30
Private Cloud Manageme...	10.201.12.0/26
vMotion network	10.201.13.128/25

2. From a command prompt, ping the components and workloads to verify that they are pinging from the same subscription.

```
C:\Users\contoso\contoso>ping 10.201.12.2

Pinging 10.201.12.2 with 32 bytes of data:
Reply from 10.201.12.2: bytes=32 time=222ms TTL=60
Reply from 10.201.12.2: bytes=32 time=222ms TTL=60
Reply from 10.201.12.2: bytes=32 time=222ms TTL=60
Reply from 10.201.12.2: bytes=32 time=221ms TTL=60

Ping statistics for 10.201.12.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 221ms, Maximum = 222ms, Average = 221ms

C:\Users\contoso\contoso>ping 10.201.12.3

Pinging 10.201.12.3 with 32 bytes of data:
Reply from 10.201.12.3: bytes=32 time=224ms TTL=60
Reply from 10.201.12.3: bytes=32 time=224ms TTL=60
Reply from 10.201.12.3: bytes=32 time=223ms TTL=60
Reply from 10.201.12.3: bytes=32 time=223ms TTL=60

Ping statistics for 10.201.12.3:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 223ms, Maximum = 224ms, Average = 223ms

C:\Users\contoso\contoso>ping 192.168.12.11

Pinging 192.168.12.11 with 32 bytes of data:
Reply from 192.168.12.11: bytes=32 time=227ms TTL=123
Reply from 192.168.12.11: bytes=32 time=231ms TTL=123
Reply from 192.168.12.11: bytes=32 time=227ms TTL=123
```


3. Select the Subscription (change) link.

The screenshot shows the Microsoft Azure portal interface. At the top, there's a search bar and the text 'Microsoft Azure'. Below that, the breadcrumb 'Home > ContosoSEA0312' is visible. The main content area is titled 'Essentials' and shows details for the resource group 'Contosorg3'. The 'Subscription (change)' link is highlighted with a red box. Other details include: Status: Succeeded, Location: Southeast Asia, Subscription ID: 81004dbe-f479-5415-bdcg-4e4bb91189f4, and Tags (change) link.

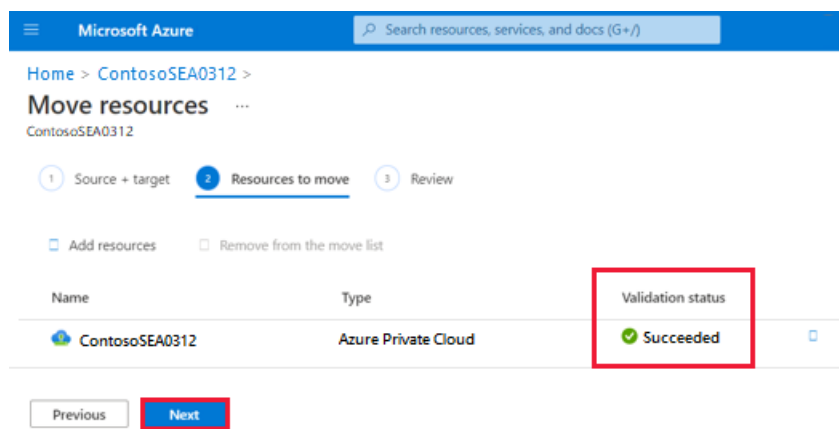
4. Provide the subscription details for Target and select Next.

The screenshot shows the 'Move resources' wizard in the Microsoft Azure portal. The wizard has three steps: 1. Source + target, 2. Resources to move, and 3. Review. The 'Source' section shows 'Subscription' as 'Azure VMware solution Group' and 'Resource group' as 'Contosorg3'. The 'Target' section is highlighted with a red box and shows 'Subscription' as 'AzureVsolution' and 'Resource group' as 'Migration-RG'. At the bottom, there are 'Previous' and 'Next' buttons, with 'Next' being highlighted.

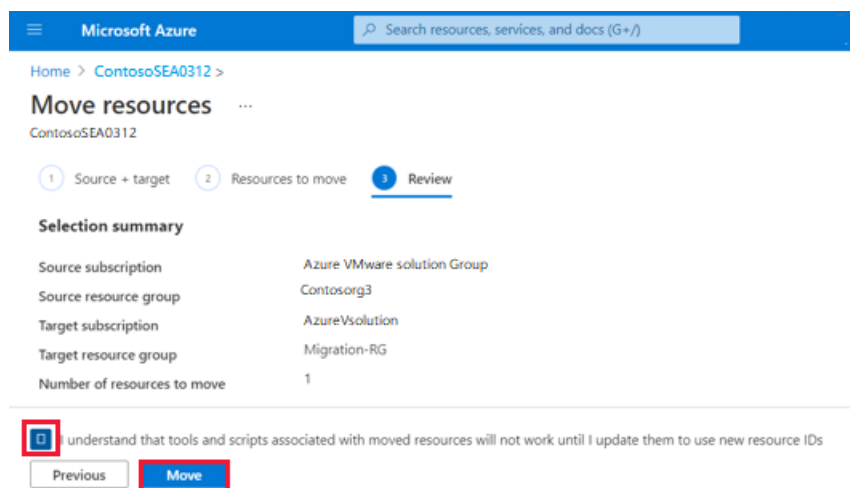
5. Confirm the validation of the resources you selected to move. During the validation, you'll see *Pending validation* under **Validation status**.



6. Once the validation is successful, select **Next** to start the migration of your private cloud.

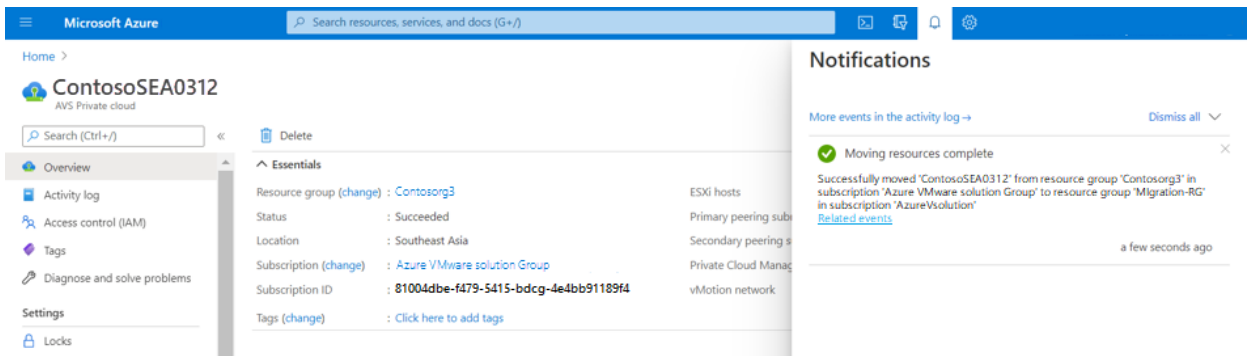


7. Select the check box indicating you understand that the tools and scripts associated won't work until you update them to use the new resource IDs. Then select **Move**.

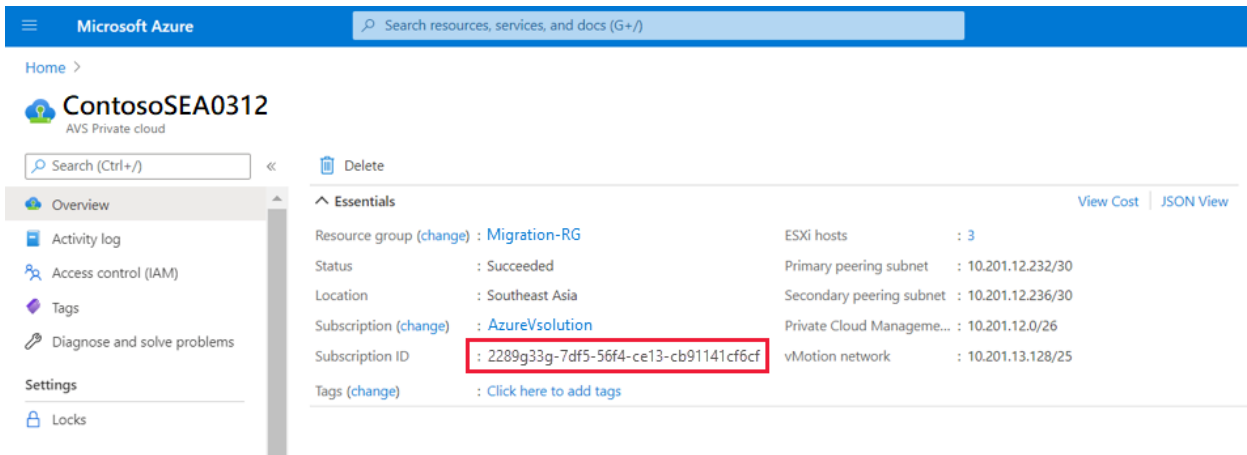


Verify the move

A notification appears once the resource move is complete.



The new subscription appears in the private cloud Overview.



Next steps

Learn more about:

- [Move Azure VMware Solution across regions](#)
- [Move guidance for networking resources](#)
- [Move guidance for virtual machines](#)
- [Move guidance for App Service resources](#)

Move Azure VMware Solution resources to another region

12/16/2022 • 10 minutes to read • [Edit Online](#)

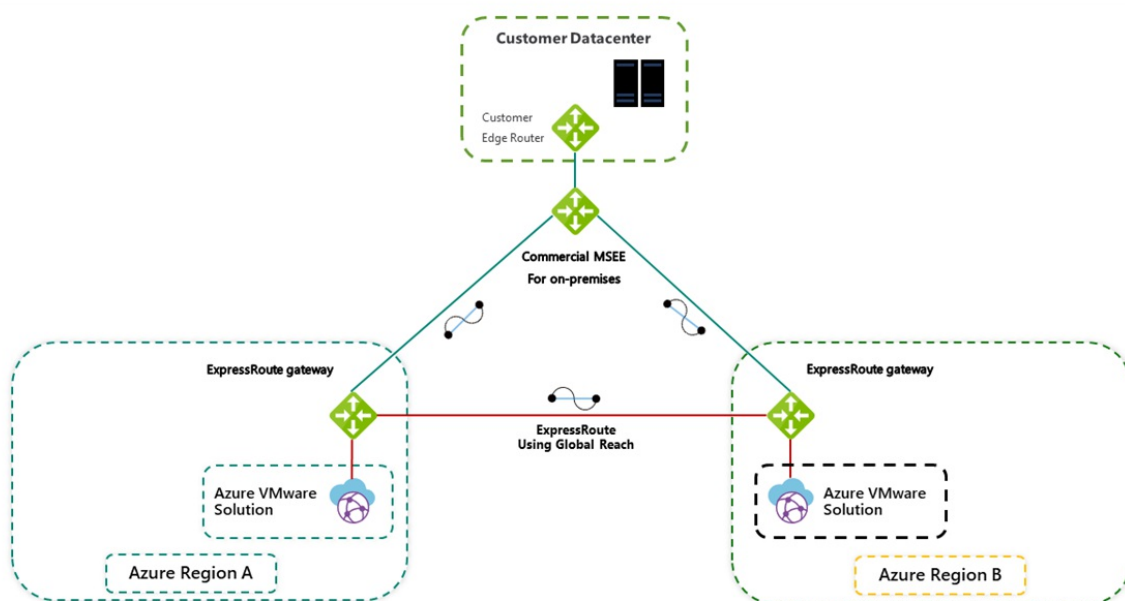
IMPORTANT

The steps in this article are strictly for moving Azure VMware Solution (source) in one region to Azure VMware Solution (target) in another region.

You can move Azure VMware Solution resources to a different region for several reasons. For example, deploy features or services available in specific regions only, meet policy and governance requirements, or respond to capacity planning requirements.

This article helps you plan and migrate Azure VMware Solution from one Azure region to another, such as Azure region A to Azure region B.

The diagram shows the recommended ExpressRoute connectivity between the two Azure VMware Solution environments. An HCX site pairing and service mesh are created between the two environments. The HCX migration traffic and Layer-2 extension moves (depicted by the red line) between the two environments. For VMware recommended HCX planning, see [Planning an HCX Migration](#).




NOTE

You don't need to migrate any workflow back to on-premises because the traffic will flow between the private clouds (source and target):

Azure VMware Solution private cloud (source) > ExpressRoute gateway (source) > ExpressRoute gateway (target) > Azure VMware Solution private cloud (target)

The diagram shows the connectivity between both Azure VMware Solution environments.



In this article, we'll walk you through the steps to:

- Prepare and plan the move to another Azure region
- Establish network connectivity between the two Azure VMware Solution private clouds
- Export the configuration from the Azure VMware Solution source environment
- Reapply the supported configuration elements to the Azure VMware Solution target environment
- Migrate workloads using VMware HCX

Prerequisites

- [VMware HCX appliance is upgraded to the latest patch](#) to avoid migration issues if any.
- Source's local content library is a [published content library](#).

Prepare

The following steps show how to prepare your Azure VMware Solution private cloud to move to another Azure VMware Solution private cloud.

Export the source configuration

1. From the source, [export the extended segments, firewall rules, port details, and route tables](#).
2. [Export the contents of an inventory list view to a CSV file](#).
3. [Sort workloads into migration groups \(migration wave\)](#).

Deploy the target environment

Before you can move the source configuration, you'll need to [deploy the target environment](#).

Back up the source configuration

Back up the Azure VMware Solution (source) configuration that includes vCenter Server, NSX-T Data Center, and firewall policies and rules.

- **Compute:** Export existing inventory configuration. For Inventory backup, you can use RVtools (an open-source app).
- **Network and firewall policies and rules:** On the Azure VMware Solution target, create the same network segments as the source environment.

Azure VMware Solution supports all backup solutions. You'll need CloudAdmin privileges to install, backup data, and restore backups. For more information, see [Backup solutions for Azure VMware Solution VMs](#).

- VM workload backup using the Commvault solution:
 - [Create a VMware client](#) from the Command center for Azure VMware Solution vCenter.
 - [Create a VM group](#) with the required VMs for backups.
 - [Run backups on VM groups](#).
 - [Restore VMs](#).
- VM workload backup using [Veritas NetBackup solution](#).

TIP

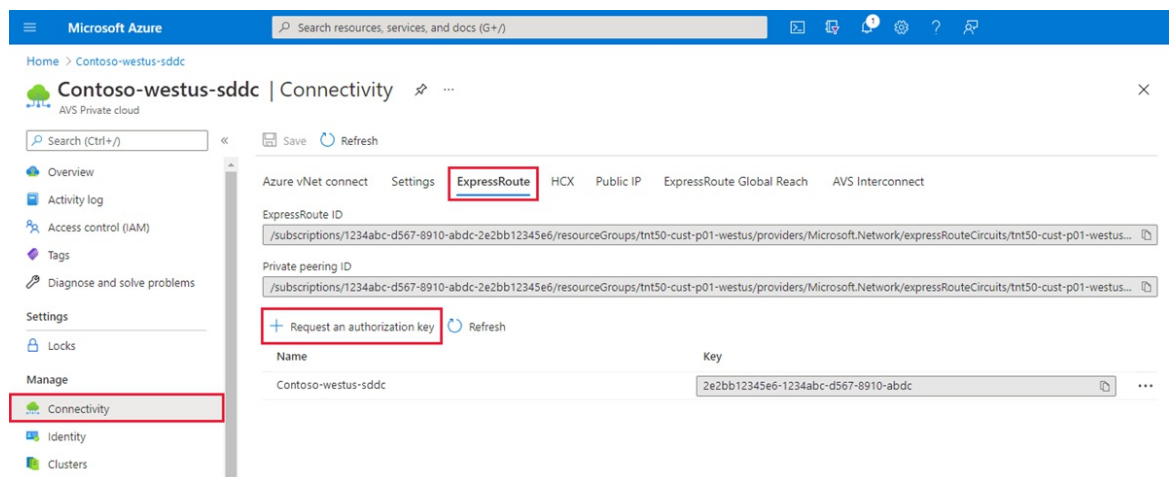
You can use [Azure Resource Mover](#) to verify and migrate the list of supported resources to move across regions, which are dependent on Azure VMware Solution.

Locate the source ExpressRoute circuit ID

1. From the source, sign in to the [Azure portal](#).
2. Select **Manage > Connectivity > ExpressRoute**.
3. Copy the source's ExpressRoute ID. You'll need it to peer between the private clouds.

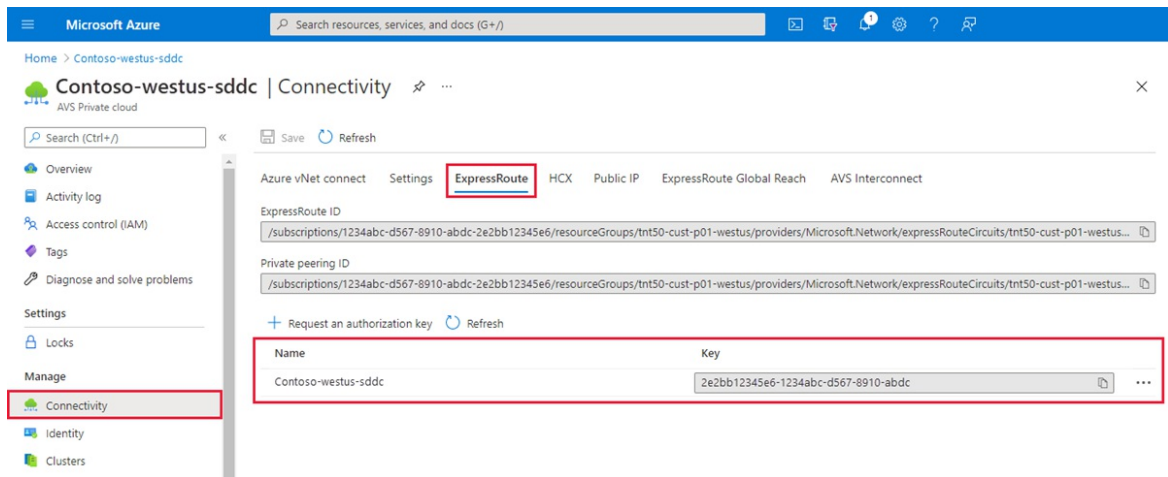
Create the target's authorization key

1. From the target, sign in to the [Azure portal](#).
2. Select **Manage > Connectivity > ExpressRoute** and then select **+ Request an authorization key**.



3. Provide a name for it and select **Create**.

It may take about 30 seconds to create the key. Once created, the new key appears in the list of authorization keys for the private cloud.

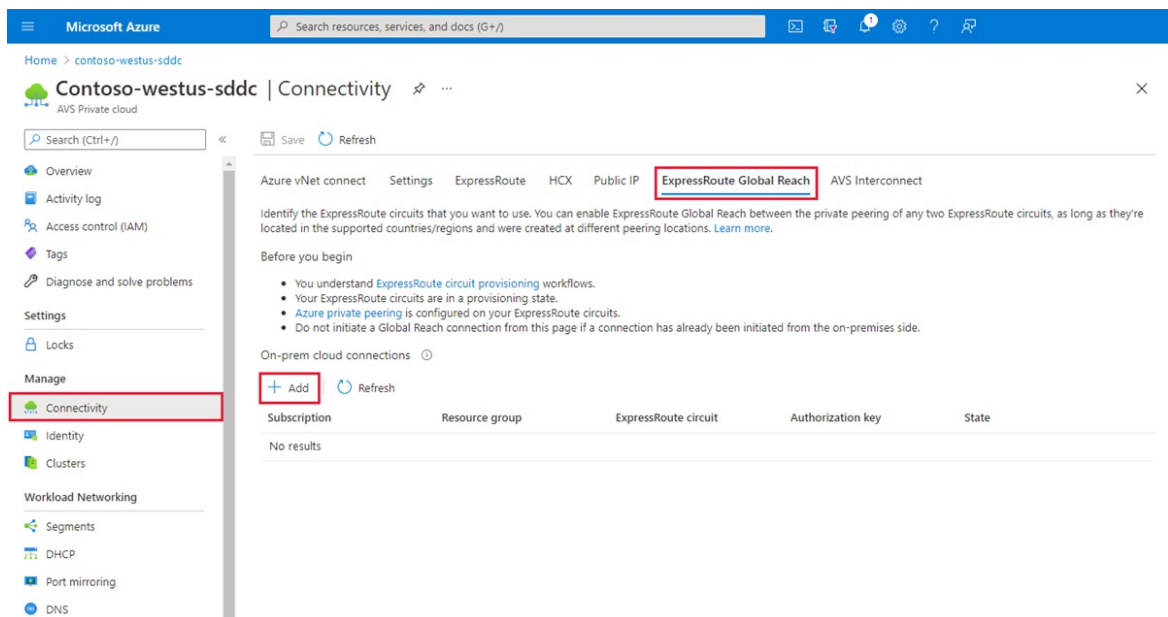


4. Copy the authorization key and ExpressRoute ID. You'll need them to complete the peering. The authorization key disappears after some time, so copy it as soon as it appears.

Peer between private clouds

Now that you have the ExpressRoute circuit IDs and authorization keys for both environments, you can peer the source to the target. You'll use the resource ID and authorization key of your private cloud ExpressRoute circuit to finish the peering.

1. From the target, sign in to the [Azure portal](#) using the same subscription as the source's ExpressRoute circuit.
2. Under Manage, select **Connectivity > ExpressRoute Global Reach > Add**.



3. Paste the ExpressRoute circuit ID and target's authorization key you created in the previous step. Then select **Create**:

On-prem cloud connections



Subscription
Contoso

Resource group
contoso-westus-rg

ExpressRoute circuit * ⓘ
[Empty dropdown]

or

If you have a circuit ID, copy/paste below ⓘ
Enter an ExpressRoute circuit ID

Authorization key
[Empty text box]

Create Cancel

Create a site pairing between private clouds

After you establish connectivity, you'll create a VMware HCX site pairing between the private clouds to facilitate the migration of your VMs. You can connect or pair the VMware HCX Cloud Manager in Azure VMware Solution with the VMware HCX Connector in your data center.

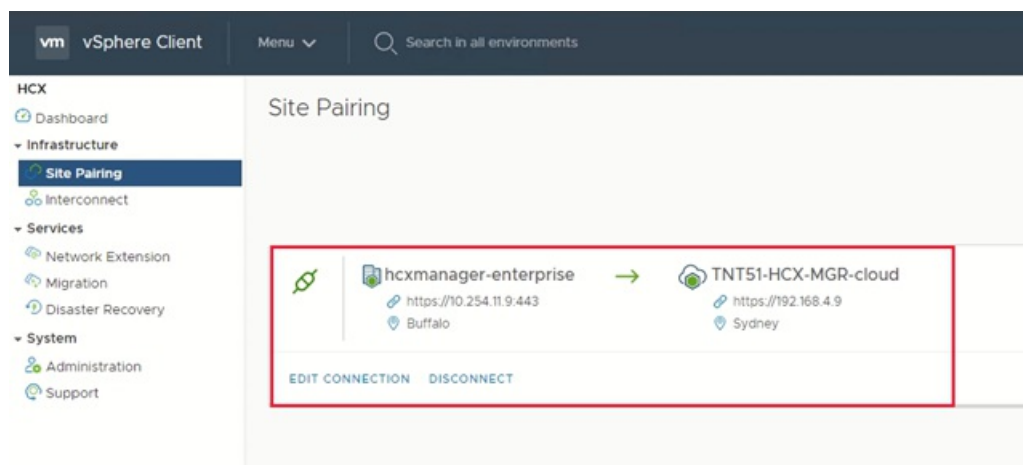
1. Sign in to your source's vCenter Server, and under **Home**, select **HCX**.
2. Under **Infrastructure**, select **Site Pairing** and select the **Connect To Remote Site** option (in the middle of the screen).
3. Enter the Azure VMware Solution HCX Cloud Manager URL or IP address you noted earlier `https://x.x.x.9`, the Azure VMware Solution cloudadmin@vsphere.local username, and the password. Then select **Connect**.

NOTE

To successfully establish a site pair:

- Your VMware HCX Connector must be able to route to your HCX Cloud Manager IP over port 443.
- Use the same password that you used to sign in to vCenter Server. You defined this password on the initial deployment screen.

You'll see a screen showing that your VMware HCX Cloud Manager in Azure VMware Solution and your on-premises VMware HCX Connector are connected (paired).



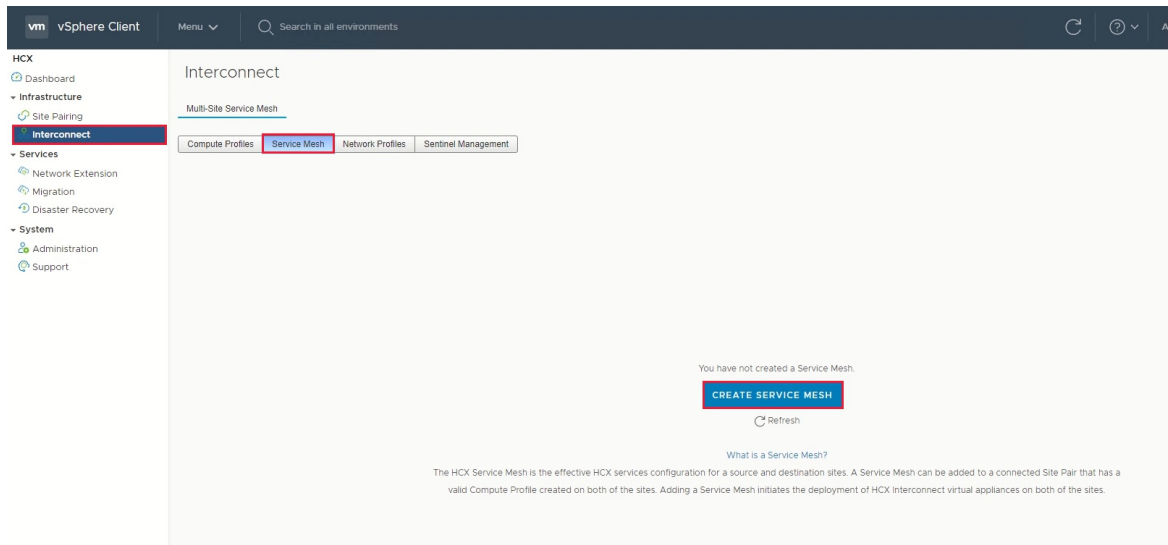
Create a service mesh between private clouds

NOTE

To successfully establish a service mesh with Azure VMware Solution:

- Ports UDP 500/4500 are open between your on-premises VMware HCX Connector 'uplink' network profile addresses and the Azure VMware Solution HCX Cloud 'uplink' network profile addresses.
- Be sure to review the [VMware HCX required ports](#).

1. Under Infrastructure, select Interconnect > Service Mesh > Create Service Mesh.



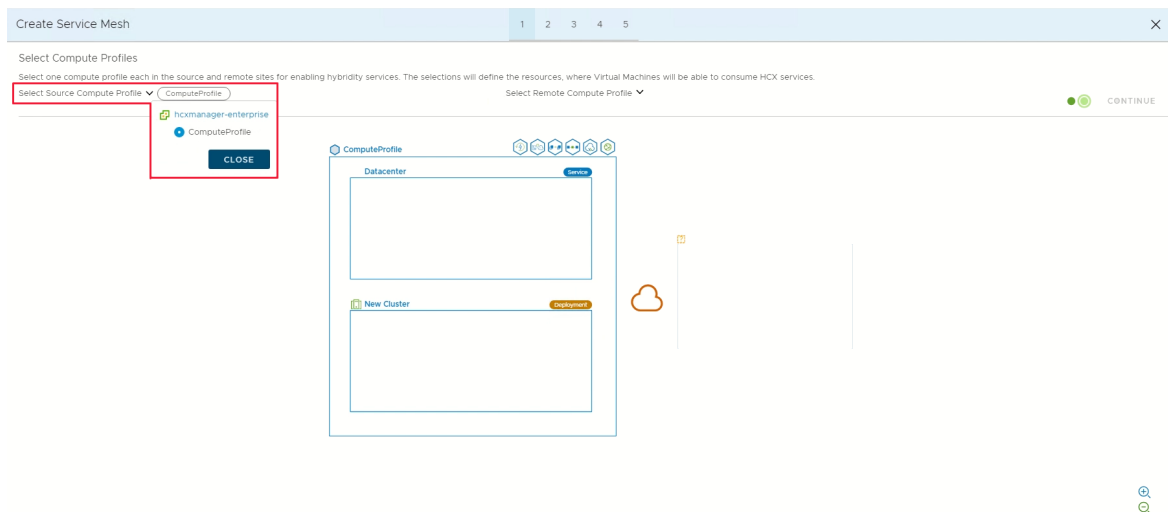
2. Review the pre-populated sites, and then select Continue.

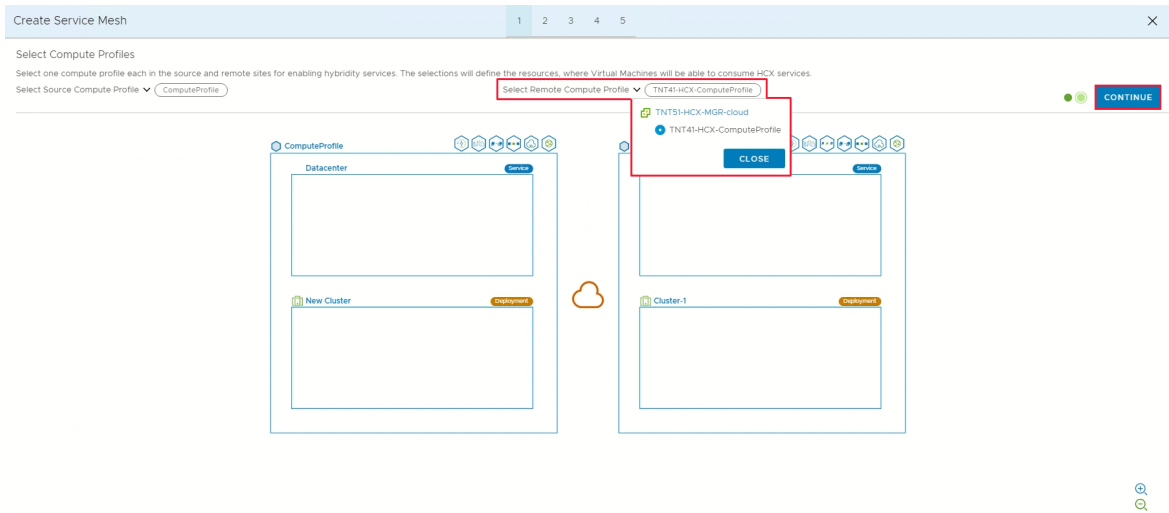
NOTE

If this is your first service mesh configuration, you won't need to modify this screen.

3. Select the source and remote compute profiles from the drop-down lists, and then select Continue.

The selections define the resources where VMs can consume VMware HCX services.






4. Review services that you want to be enabled, and then select **Continue**.

5. In **Advanced Configuration - Override Uplink Network profiles**, select **Continue**.

Uplink network profiles connect to the network through which the remote site's interconnect appliances can be reached.

6. In **Advanced Configuration - Network Extension Appliance Scale Out**, review and select **Continue**.

You can have up to eight VLANs per appliance, but you can deploy another appliance to add another eight VLANs. You must also have IP space to account for the more appliances, and it's one IP per appliance. For more information, see [VMware HCX Configuration Limits](#).

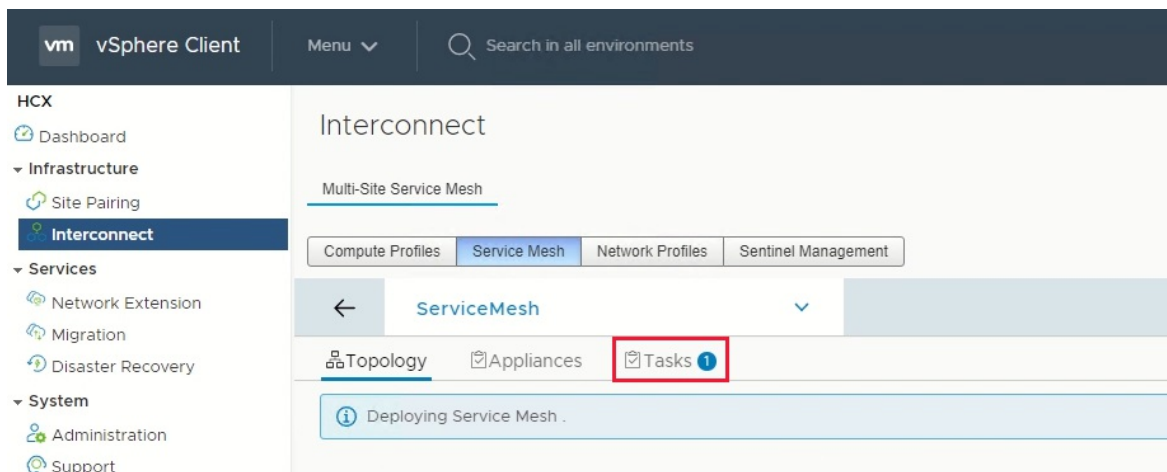


7. In **Advanced Configuration - Traffic Engineering**, review and make any modifications that you feel are necessary, and then select **Continue**.

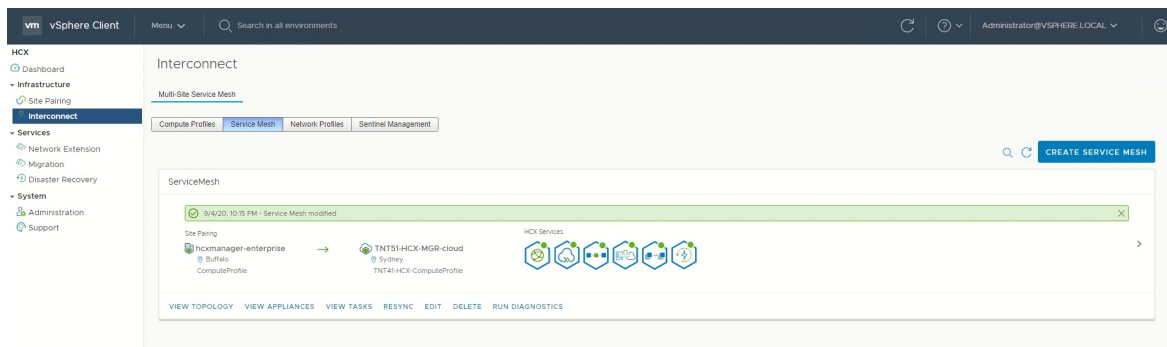
8. Review the topology preview and select **Continue**.

9. Enter a user-friendly name for this service mesh and select **Finish** to complete.

10. Select **View Tasks** to monitor the deployment.

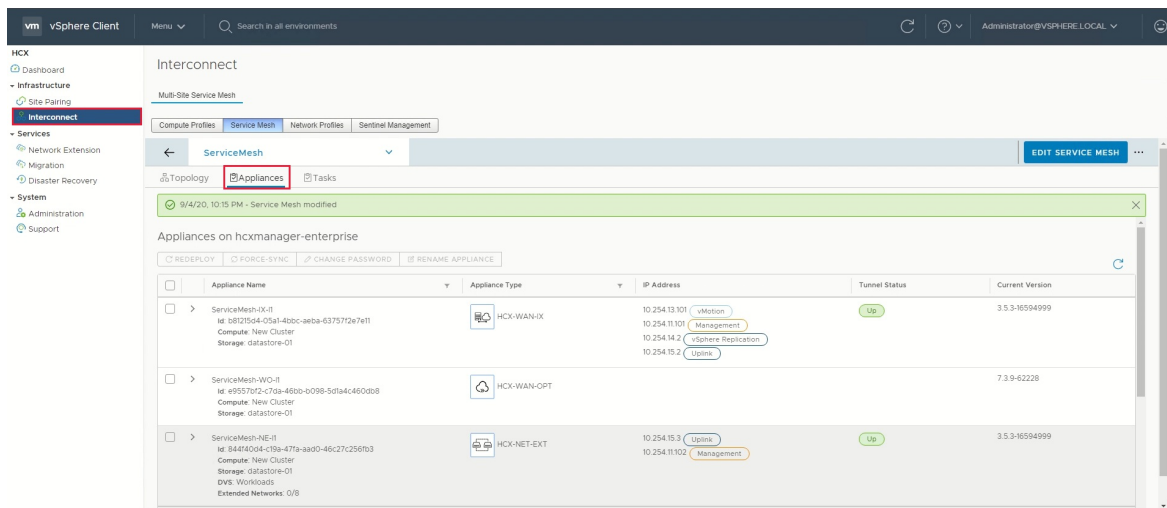


When the service mesh deployment finishes successfully, you'll see the services as green.



11. Verify the service mesh's health by checking the appliance status.

12. Select Interconnect > Appliances.



Move

The following steps show how to move your Azure VMware Solution private cloud resources to another Azure VMware Solution private cloud in a different region.

In this section, you'll migrate the:

- Resource pool configuration and folder creation
- VM templates and the associated tags
- Logical segments deployment based on the source's port groups and associated VLANs
- Network security services and groups

- Gateway firewall policy and rules based on the source's firewall policies

Migrate the source vSphere configuration

In this step, you'll copy the source's vSphere configuration and move it to the target environment.

1. From the source's vCenter Server, use the same resource pool configuration and [create the same resource pool configuration](#) on the target's vCenter Server.
2. From the source's vCenter Server, use the same VM folder name and [create the same VM folder](#) on the target's vCenter Server under **Folders**.
3. Use VMware HCX to migrate all VM templates from the source's vCenter Server to the target's vCenter.
 - a. From the source, convert the existing templates to VMs and then migrate them to the target.
 - b. From the target, convert the VMs to VM templates.
4. From the source environment, use the same VM Tags name and [create them on the target's vCenter](#).
5. From the source's vCenter Server Content Library, use the subscribed library option to copy the ISO, OVF, OVA, and VM Templates to the target content library:
 - a. If the content library isn't already published, select the **Enable publishing** option.
 - b. From the source's Content Library, copy the URL of the published library.
 - c. From the target, [create a subscribed content library](#) with the URL from the source's library.
 - d. Select **Sync Now**.

Configure the target NSX-T Data Center environment

In this step, you'll use the source NSX-T Data Center configuration to configure the target NSX-T environment.

NOTE

You'll have multiple features configured on the source NSX-T Data Center, so you must copy or read from the source NSX-T Data Center and recreate it in the target private cloud. Use L2 Extension to keep same IP address and Mac Address of the VM while migrating Source to target AVS Private Cloud to avoid downtime due to IP change and related configuration.

1. [Configure NSX-T Data Center network components](#) required in the target environment under default Tier-1 gateway.
2. [Create the security group configuration](#).
3. [Create the distributed firewall policy and rules](#).
4. [Create the gateway firewall policy and rules](#).
5. [Create the DHCP server or DHCP relay service](#).
6. [Configure port mirroring](#).
7. [Configure DNS forwarder](#).
8. [Configure a new Tier-1 gateway \(other than default\)](#). This configuration is based on the NSX-T Data Center configured on the source.

Migrate the VMs from the source

In this step, you'll use VMware HCX to migrate the VMs from the source to the target. You'll have the option to do a Layer-2 extension from the source and use HCX to vMotion the VMs from the source to the target with

minimal interruption.

Besides vMotion, other methods, like Bulk and Cold vMotion, are also recommended. Learn more about:

- [Plan an HCX Migration](#)
- [Migrate Virtual Machines with HCX](#)

Cutover extended networks

In this step, you'll do a final gateway cutover to terminate the extended networks. You'll also move (migrate) the gateways from the source Azure VMware Solution environment to the target environment.

IMPORTANT

You must do the gateway cutover post VLAN workload migration to the target Azure VMware Solution environment. Also, there shouldn't be any VM dependency on the source and target environments.

Before the gateway cutover, verify all migrated workload services and performance. Once application and web service owners accept the performance (except for any latency issues), you can continue with the gateway cutover. Once you've completed the cutover, you'll need to modify the public DNS A and PTR records.

For VMware recommendations, see [Cutover of extended networks](#).

Public IP DNAT for migrated DMZ VMs

To this point, you've migrated the workloads to the target environment. These application workloads must be reachable from the public internet. The target environment provides two ways of hosting any application. Applications can be:

- Hosted and published under the application gateway load balancer.
- Published through the public IP feature in vWAN.

Public IP is typically the destination NAT translated into the Azure firewall. With DNAT rules, firewall policy would translate the public IP address request to a private address (webserver) with a port. For more information, see [How to use the public IP functionality in Azure Virtual WAN](#).

NOTE

SNAT is by default configured in Azure VMware Solution, so you must enable SNAT from Azure VMware Solution private cloud connectivity settings under the Manage tab.

Decommission

For this last step, you'll verify that all the VM workloads were migrated successfully, including the network configuration. If there's no dependency, you can disconnect the HCX service mesh, site pairing, and network connectivity from the source environment.

NOTE

Once you decommission the private cloud, you cannot undo it as the configuration and data will be lost.

Next steps

Learn more about:

- [Move operation support for Microsoft.AVS](#)
- [Move guidance for networking resources](#)
- [Move guidance for virtual machines](#)
- [Move guidance for App Service resources](#)

Request host quota for Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

In this how-to, you'll request host quota/capacity for [Azure VMware Solution](#). You'll submit a support ticket to have your hosts allocated whether it's for a new deployment or an existing one.

If you have an existing Azure VMware Solution private cloud and want more hosts allocated, you'll follow the same process.

IMPORTANT

It can take up to five business days to allocate the hosts, depending on the number requested. So request what is needed for provisioning, so you don't need to request a quota increase as often.

Eligibility criteria

You'll need an Azure account in an Azure subscription that adheres to one of the following criteria:

- A subscription under an [Azure Enterprise Agreement \(EA\)](#) with Microsoft.
- A Cloud Solution Provider (CSP) managed subscription under an existing CSP Azure offers contract or an Azure plan.
- A [Microsoft Customer Agreement \(MCA\)](#) with Microsoft.

Request host quota for EA and MCA customers

1. In your Azure portal, under **Help + Support**, create a [New support request](#) and provide the following information:

- **Issue type:** Technical
- **Subscription:** Select your subscription
- **Service:** All services > Azure VMware Solution
- **Resource:** General question
- **Summary:** Need capacity
- **Problem type:** Capacity Management Issues
- **Problem subtype:** Customer Request for Additional Host Quota/Capacity

2. In the **Description** of the support ticket, on the **Details** tab, provide information for:

- Region Name
- Number of hosts
- Any other details

NOTE

Azure VMware Solution requires a minimum of three hosts and recommends redundancy of N+1 hosts.

3. Select **Review + Create** to submit the request.

Request host quota for CSP customers

CSPs must use [Microsoft Partner Center](#) to enable Azure VMware Solution for their customers. This article uses [CSP Azure plan](#) as an example to illustrate the purchase procedure for partners.

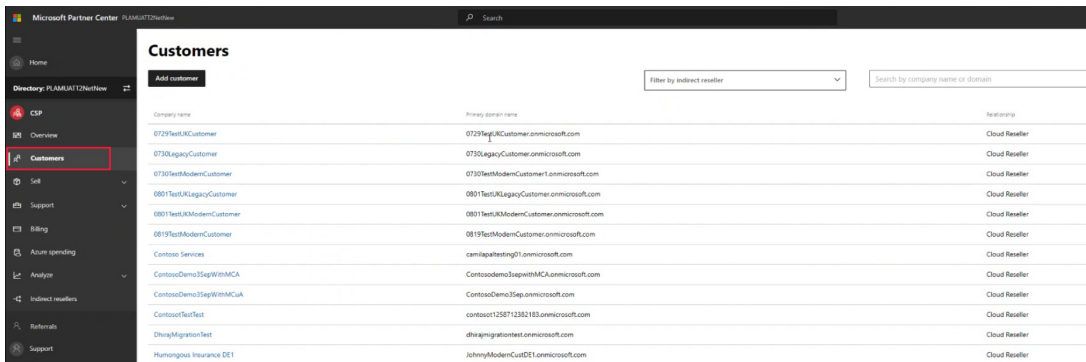
Access the Azure portal using the **Admin On Behalf Of (AOBO)** procedure from Partner Center.

IMPORTANT

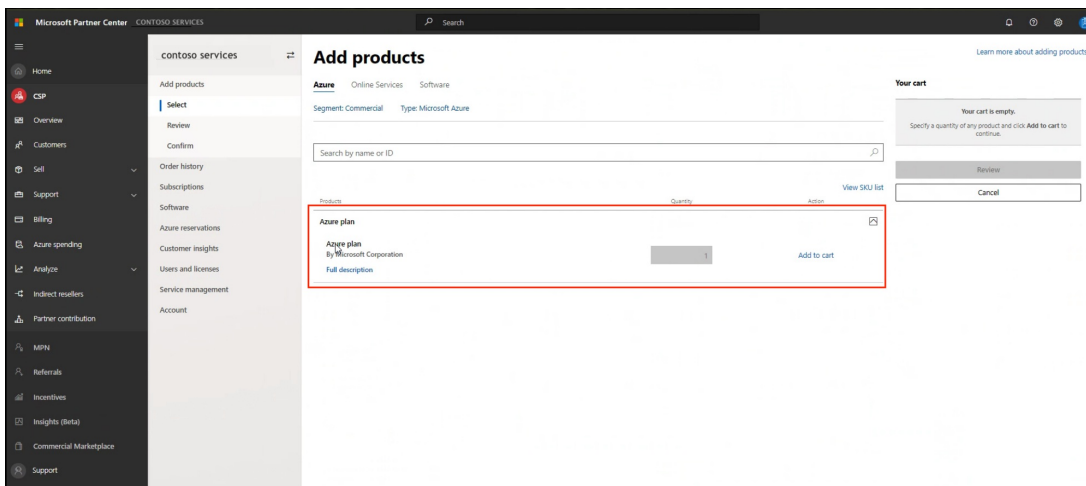
Azure VMware Solution service does not provide multi-tenancy support. Hosting partners requiring it are not supported.

1. Configure the CSP Azure plan:

a. In **Partner Center**, select **CSP** to access the **Customers** area.



b. Select your customer and then select **Add products**.



c. Select **Azure plan** and then select **Add to cart**.

d. Review and finish the general setup of the Azure plan subscription for your customer. For more information, see [Microsoft Partner Center documentation](#).

2. After you configure the Azure plan and you have the needed [Azure RBAC permissions](#) in place for the subscription, you'll request the quota for your Azure plan subscription.

a. Access Azure portal from [Microsoft Partner Center](#) using the **Admin On Behalf Of (AOBO)** procedure.

b. Select **CSP** to access the **Customers** area.

c. Expand customer details and select **Microsoft Azure Management Portal**.

d. In the Azure portal, under **Help + Support**, create a **New support request** and provide the following information:

- **Issue type:** Technical

- **Subscription:** Select your subscription
 - **Service:** All services > Azure VMware Solution
 - **Resource:** General question
 - **Summary:** Need capacity
 - **Problem type:** Capacity Management Issues
 - **Problem subtype:** Customer Request for Additional Host Quota/Capacity
- e. In the **Description** of the support ticket, on the **Details** tab, provide information for:
- Region Name
 - Number of hosts
 - Any other details
 - Is intended to host multiple customers?

NOTE

Azure VMware Solution requires a minimum of three hosts and recommends redundancy of N+1 hosts.

- f. Select **Review + Create** to submit the request.

Next steps

Before deploying Azure VMware Solution, you must first [register the resource provider](#) with your subscription to enable the service.

Rotate the cloudadmin credentials for Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

IMPORTANT

Currently, rotating your NSX-T Manager *admin* credentials isn't supported. To rotate your NSX-T Manager password, submit a [support request](#). This process might impact running HCX services.

In this article, you'll rotate the cloudadmin credentials (vCenter Server *CloudAdmin* credentials) for your Azure VMware Solution private cloud. Although the password for this account doesn't expire, you can generate a new one at any time.

Caution

If you use your cloudadmin credentials to connect services to vCenter Server in your private cloud, those connections will stop working once you rotate your password. Those connections will also lock out the cloudadmin account unless you stop those services before rotating the password.

Prerequisites

Consider and determine which services connect to vCenter Server as *cloudadmin@vsphere.local* before you rotate the password. These services may include VMware services such as HCX, vRealize Orchestrator, vRealize Operations Manager, VMware Horizon, or other third-party tools used for monitoring or provisioning.

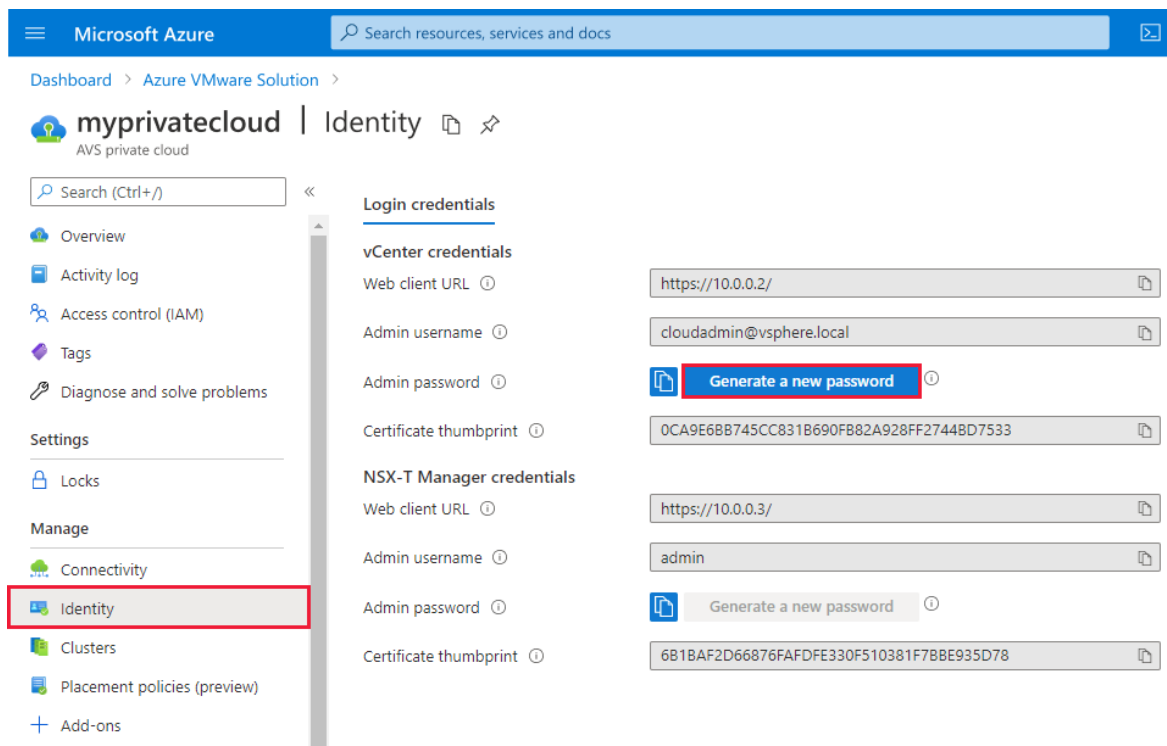
One way to determine which services authenticate to vCenter Server with the cloudadmin user is to inspect vSphere events using the vSphere Client for your private cloud. After you identify such services, and before rotating the password, you must stop these services. Otherwise, the services won't work after you rotate the password. You'll also experience temporary locks on your vCenter Server CloudAdmin account, as these services continuously attempt to authenticate using a cached version of the old credentials.

Instead of using the cloudadmin user to connect services to vCenter, we recommend individual accounts for each service. For more information about setting up separate accounts for connected services, see [Access and Identity Concepts](#).

Reset your vCenter Server credentials

- [Portal](#)
- [Azure CLI](#)

1. In your Azure VMware Solution private cloud, select **Identity**.
2. Select **Generate new password**.




3. Select the confirmation checkbox and then select **Generate password**.

Update HCX Connector

1. Go to the on-premises HCX Connector at [https://\[ip of the HCX connector appliance\]:443](https://[ip of the HCX connector appliance]:443) and sign in using the new credentials.

Be sure to use port 443.

2. On the VMware HCX Dashboard, select **Site Pairing**.



3. Select the correct connection to Azure VMware Solution and select **Edit Connection**.

4. Provide the new vCenter Server user credentials and select **Edit**, which saves the credentials. Save should show successful.

Next steps

Now that you've covered resetting your vCenter Server credentials for Azure VMware Solution, you may want to learn about:

- [Integrating Azure native services in Azure VMware Solution](#)
- [Deploying disaster recovery for Azure VMware Solution workloads using VMware HCX](#)

Save costs with Azure VMware Solution

12/16/2022 • 6 minutes to read • [Edit Online](#)

When you commit to a reserved instance of [Azure VMware Solution](#), you save money. The reservation discount automatically applies to the running Azure VMware Solution hosts that match the reservation scope and attributes. In addition, a reserved instance purchase covers only the compute part of your usage and includes software licensing costs.

Purchase restriction considerations

Reserved instances are available with some exceptions.

- **Clouds** - Reservations are available only in the regions listed on the [Products available by region](#) page.
- **Insufficient quota** - A reservation scoped to a single/shared subscription must have hosts quota available in the subscription for the new reserved instance. You can [create quota increase request](#) to resolve this issue.
- **Offer eligibility** - You'll need an [Azure Enterprise Agreement \(EA\)](#) with Microsoft.
- **Capacity restrictions** - In rare circumstances, Azure limits the purchase of new reservations for Azure VMware Solution host SKUs because of low capacity in a region.

Buy a reservation

You can buy a reserved instance of an Azure VMware Solution host instance in the [Azure portal](#).

You can pay for the reservation [up front or with monthly payments](#).

These requirements apply to buying a reserved dedicated host instance:

- You must be in an *Owner* role for at least one EA subscription or a subscription with a pay-as-you-go rate.
- For EA subscriptions, you must enable the **Add Reserved Instances** option in the [EA portal](#). If disabled, you must be an EA Admin for the subscription to enable it.
- For subscription under a Cloud Solution Provider (CSP) Azure Plan, the partner must purchase the customer's reserved instances in the Azure portal.

Buy reserved instances for an EA subscription

1. Sign in to the [Azure portal](#).
2. Select **All services > Reservations**.
3. Select **Purchase Now**, then select **Azure VMware Solution**.
4. Enter the required fields. The selected attributes that match running Azure VMware Solution hosts qualify for the reservation discount. Attributes include the SKU, regions (where applicable), and scope. Reservation scope selects where the reservation savings apply.

If you have an EA agreement, you can use the **Add more option** to add instances quickly. The option isn't available for other subscription types.

FIELD	DESCRIPTION
Subscription	The subscription used to pay for the reservation. The payment method on the subscription is charged the costs for the reservation. The subscription type must be an enterprise agreement (offer numbers: MS-AZR-0017P or MS-AZR-0148P), Microsoft Customer Agreement, or an individual subscription with pay-as-you-go rates (offer numbers: MS-AZR-0003P or MS-AZR-0023P). The charges are deducted from the Azure Prepayment (previously called monetary commitment) balance, if available, or charged as overage. For a subscription with pay-as-you-go rates, the charges are billed to the subscription's credit card or an invoice payment method.
Scope	The reservation's scope can cover one subscription or multiple subscriptions (shared scope). If you select: <ul style="list-style-type: none"> • Single resource group scope - Applies the reservation discount to the matching resources in the selected resource group only. • Single subscription scope - Applies the reservation discount to the matching resources in the selected subscription. • Shared scope - Applies the reservation discount to matching resources in eligible subscriptions that are in the billing context. For EA customers, the billing context is the enrollment. Therefore, the billing scope is all eligible subscriptions created by the account administrator for individual subscriptions with pay-as-you-go rates. • Management group - Applies the reservation discount to the matching resource in the list of subscriptions that are a part of both the management group and billing scope.

Region The Azure region that's covered by the reservation. Host Size AV36 Term One year or three years.

Quantity The number of instances to purchase within the reservation. The quantity is the number of running Azure VMware Solution hosts that can get the billing discount.

Buy reserved instances for a CSP subscription

CSPs that want to purchase reserved instances for their customers must use the **Admin On Behalf Of (AOBO)** procedure from the [Partner Center documentation](#). For more information, view the Admin on behalf of (AOBO) video.

1. Sign in to [Partner Center](#).
2. Select **CSP** to access the **Customers** area.
3. Expand customer details and select **Microsoft Azure Management Portal**.

Company name	Primary domain name	Relationship
0729TestUKCustomer	0729TestUKCustomer.onmicrosoft.com	Cloud Reseller
0730LegacyCustomer	0730LegacyCustomer.onmicrosoft.com	Cloud Reseller
0730TestModernCustomer	0730TestModernCustomer1.onmicrosoft.com	Cloud Reseller
0801TestUKLegacyCustomer	0801TestUKLegacyCustomer.onmicrosoft.com	Cloud Reseller
0801TestUKModernCustomer	0801TestUKModernCustomer.onmicrosoft.com	Cloud Reseller
0819TestModernCustomer	0819TestModernCustomer.onmicrosoft.com	Cloud Reseller
Contoso Services	camilapaltesting01.onmicrosoft.com	Cloud Reseller
Microsoft ID:	4316479f-9cb2-4485-a3ad-b47844c1d4d2	
Products:	Add products View orders	
Subscriptions:	Add subscriptions View subscriptions	
Licenses:	Users and licenses	
Administer services:	Azure Active Directory Microsoft 365 Microsoft Azure Management Portal Visual Studio Marketplace Manage Visual Studio subscriptions View services	
Indirect reseller(s):	--	
Service costs:	View service costs	
Service alerts:		
ContosoDemo3SepWithMCA	Contosodemo3sepwithMCA.onmicrosoft.com	Cloud Reseller
ContosoDemo3SepWithMCAuA	Contosodemo3sep.onmicrosoft.com	Cloud Reseller
ContosotTestTest	contosot1258712382183.onmicrosoft.com	Cloud Reseller
DhrinjMigrationTest	dhrinjmigrationtest.onmicrosoft.com	Cloud Reseller

4. In the Azure portal, select **All services > Reservations**.

5. Select **Purchase Now** and then select **Azure VMware Solution**.

Microsoft Azure

Home > Purchase reservations

Products Review + buy

Receive discounts on your Azure services by purchasing reservations. See FAQs

Filter by name...

- Virtual machine
- Azure Blob Storage
- Azure Database for PostgreSQL
- Azure Data Explorer
- Azure Red Hat OpenShift
- SQL Database
- Azure Dedicated Host
- Azure Managed Disks
- SUSE Linux
- App Services
- Azure Synapse Analytics (formerly SQL Data Warehouse)
- Azure Database for MySQL
- Azure Databricks
- Red Hat Plans
- Azure VMware Solution**
- Azure Cosmos DB
- Azure Database for MariaDB
- Azure Cache for Redis
- Azure VMware Solution by CloudSimple

Next: Review + buy

6. Enter the required fields. The selected attributes that match running Azure VMware Solution hosts qualify for the reservation discount. Attributes include the SKU, regions (where applicable), and scope. Reservation scope selects where the reservation savings apply.

FIELD	DESCRIPTION
Subscription	The subscription that funds the reservation. The payment method on the subscription is charged the costs for the reservation. The subscription type must be an eligible one, which in this case is a CSP subscription

FIELD	DESCRIPTION
Scope	<p>The reservation's scope can cover one subscription or multiple subscriptions (shared scope). If you select:</p> <ul style="list-style-type: none"> • Single resource group scope - Applies the reservation discount to the matching resources in the selected resource group only. • Single subscription scope - Applies the reservation discount to the matching resources in the selected subscription. • Shared scope - Applies the reservation discount to matching resources in eligible subscriptions that are in the billing context. For EA customers, the billing context is the enrollment. Therefore, the billing scope is all eligible subscriptions created by the account administrator for individual subscriptions with pay-as-you-go rates. • Management group - Applies the reservation discount to the matching resource in the list of subscriptions that are a part of both the management group and billing scope.

Region The Azure region that's covered by the reservation. Host Size AV36 Term One year or three years.

Quantity The number of instances to purchase within the reservation. The quantity is the number of running Azure VMware Solution hosts that can get the billing discount.

To learn more about viewing the purchased reservations for your customer, see [View Azure reservations as a Cloud Solution Provider \(CSP\)](#) article.

Usage data and reservation usage

Your usage that gets a reservation discount has an effective price of zero. You can see which Azure VMware Solution instance received the reservation discount for each reservation.

For more information about how reservation discounts appear in usage data:

- For EA customers, see [Understand Azure reservation usage for your Enterprise enrollment](#)
- For individual subscriptions, see [Understand Azure reservation usage for your Pay-As-You-Go subscription](#)

Change a reservation after purchase

You can make these changes to a reservation after purchase:

- Update reservation scope
- Instance size flexibility (if applicable)
- Ownership

You can also split a reservation into smaller chunks or merge reservations. None of the changes cause a new commercial transaction or change the end date of the reservation.

For details about CSP-managed reservations, see [Sell Microsoft Azure reservations to customers using Partner Center, the Azure portal, or APIs](#).

NOTE

Once you've purchased your reservation, you won't be able to make these types of changes directly:

- An existing reservation's region
- SKU
- Quantity
- Duration

However, you can *exchange* a reservation if you want to make changes.

Cancel, exchange, or refund reservations

You can cancel, exchange, or refund reservations with certain limitations. For more information, see [Self-service exchanges and refunds for Azure Reservations](#).

CSPs can cancel, exchange, or refund reservations, with certain limitations, purchased for their customer. For more information, see [Manage, cancel, exchange, or refund Microsoft Azure reservations for customers](#).

Next steps

Now that you've covered reserved instance of Azure VMware Solution, you may want to learn about:

- [Creating an Azure VMware Solution assessment](#).
- [Configure DHCP for Azure VMware Solution](#).
- [Integrating Azure native services in Azure VMware Solution](#).

Operating system support for Azure VMware Solution virtual machines

12/16/2022 • 2 minutes to read • [Edit Online](#)

Azure VMware Solution supports a wide range of operating systems to be used in the guest virtual machines. Being based on VMware vSphere, currently 7.0 version, all operating systems currently supported by vSphere can be used by any Azure VMware Solution customer for their workloads.

Check the list of operating systems and configurations supported in the [VMware Compatibility Guide](#), create a query for ESXi 7.0 Update 3 and select all operating systems and vendors.

Additionally to the supported operating systems by VMware for vSphere, we have worked with Red Hat, SUSE and Canonical to extend the support model currently in place for Azure Virtual Machines to the workloads running on Azure VMware Solution, given that it is a first-party Azure service. You can check the following sites of vendors for more information about the benefits of running their operating system on Azure.

- [Red Hat Enterprise Linux](#)
- [Ubuntu Server](#)
- [SUSE Enterprise Linux Server](#)

Backup solutions for Azure VMware Solution virtual machines (VMs)

12/16/2022 • 2 minutes to read • [Edit Online](#)

A key principle of Azure VMware Solution is to enable you to continue to use your investments and your favorite VMware solutions running on Azure. Independent software vendor (ISV) technology support, validated with Azure VMware Solution, is an important part of this strategy.

Our backup partners have industry-leading backup and restore solutions in VMware-based environments. Customers have widely adopted these solutions for their on-premises deployments. Now these partners have extended their solutions to Azure VMware Solution, using Azure to provide a backup repository and a storage target for long-term retention and archival.

Back up network traffic between Azure VMware Solution VMs and the backup repository in Azure travels over a high-bandwidth, low-latency link. Replication traffic across regions travels over the internal Azure backplane network, which lowers bandwidth costs for users.

NOTE

For common questions, see [our third-party backup solution FAQ](#).

You can find more information on these backup solutions here:

- [Cohesity](#)
- [Commvault](#)
- [Dell Technologies](#)
- [Rubrik](#)
- [Veeam](#)
- [Veritas](#)

Disaster recovery solutions for Azure VMware Solution virtual machines (VMs)

12/16/2022 • 2 minutes to read • [Edit Online](#)

One of the most important aspects of any Azure VMware Solution deployment is disaster recovery, which can be achieved by creating disaster recovery plans between different Azure VMware Solution regions or between Azure and an on-premises vSphere environment.

We currently offer customers the possibility to implement their disaster recovery plans using state-of-the-art VMware solution like [SRM](#) or [HCX](#).

Following our principle of giving customers the choice to apply their investments in skills and technology we've collaborated with some of the leading partners in the industry.

You can find more information about their solutions in the links below:

- [Jetstream](#)
- [Zerto](#)
- [RiverMeadow](#)

Migration solutions for Azure VMware Solution virtual machines (VMs)

12/16/2022 • 2 minutes to read • [Edit Online](#)

One of the most common use cases for using Azure VMware Solution is data center evacuation. It allows you to continue to maximize your VMware investments, because Azure VMware Solution will always be up to date. Additionally, you can enhance your workloads with the full range of native Azure services. An initial key step in this process is the migration of your legacy VMware-based environment onto Azure VMware Solution.

Our migration partners have industry-leading migration solutions in VMware-based environments. Customers around the world have used these solutions for their migrations to both Azure and Azure VMware Solution.

You aren't required to use VMware HCX as a migration tool, which means you can also migrate physical workloads into Azure VMware Solution. Additionally, migrations to your Azure VMware Solution environment don't need an ExpressRoute connection if it's not available within your source environment. Migrations can be done to multiple locations if you decide to host those workloads in multiple Azure regions.

You can find more information on these migration solutions here:

- [RiverMeadow](#).

Security solutions for Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

A fundamental part of Azure VMware Solution is security. It allows customers to run their VMware-based workloads in a safe and trustable environment.

Our security partners have industry-leading solutions in VMware-based environments that cover many aspects of the security ecosystem like threat protection and security scanning. Our customers have adopted many of these solutions integrated with VMware NSX-T Data Center for their on-premises deployments. As one of our key principles, we want to enable them to continue to use their investments and VMware solutions running on Azure. Many of these Independent Software Vendors (ISV) have validated their solutions with Azure VMware Solution.

You can find more information about these solutions here:

- [Bitdefender](#)
- [Trend Micro Deep Security](#)
- [Check Point](#)

Application performance monitoring and troubleshooting solutions for Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

A key objective of Azure VMware Solution is to maintain the performance and security of applications and services across VMware on Azure and on-premises. Getting there requires visibility into complex infrastructures and quickly pinpointing the root cause of service disruptions across the hybrid cloud.

Microsoft solutions

Microsoft recommends [Application Insights](#), a feature of [Azure Monitor](#), to maximize the availability and performance of your applications and services.

Learn how modern monitoring with Azure Monitor can transform your business by reviewing the [product overview, features, getting started guide and more](#).

Third-party solutions

Our application performance monitoring and troubleshooting partners have industry-leading solutions in VMware-based environments that assure the availability, reliability, and responsiveness of applications and services. Our customers have adopted many of these solutions integrated with VMware NSX-T Data Center for their on-premises deployments. As one of our key principles, we want to enable them to continue to use their investments and VMware solutions running on Azure. Many of these Independent Software Vendors (ISV) have validated their solutions with Azure VMware Solution.

You can find more information about these solutions here:

- [NETSCOUT](#)
- [Turbonomic](#)

Bitnami appliance deployment

12/16/2022 • 3 minutes to read • [Edit Online](#)

Bitnami by VMware provides a rich catalog of turnkey virtual appliances. You can deploy any vSphere compatible appliance by Bitnami available in the [VMware Marketplace](#), including many of the most common open-source software projects.

In this article, you'll learn how to install and configure the following virtual appliances packaged by Bitnami on your Azure VMware Solution private cloud:

- LAMP
- Jenkins
- PostgreSQL
- NGINX
- RabbitMQ

Prerequisites

- Azure VMware Solution private cloud [deployed with a minimum of three nodes](#).
- Networking configured as described in [Network planning checklist](#).

Step 1. Download the Bitnami virtual appliance OVA/OVF file

1. Go to the [VMware Marketplace](#) and download the virtual appliance you want to install on your Azure VMware Solution private cloud:
 - [LAMP virtual appliance packaged by Bitnami](#)
 - [Jenkins](#)
 - [PostgreSQL](#)
 - [NGINX](#)
 - [RabbitMQ](#)
2. Select the version, select **Download**, and then accept the EULA license.

NOTE

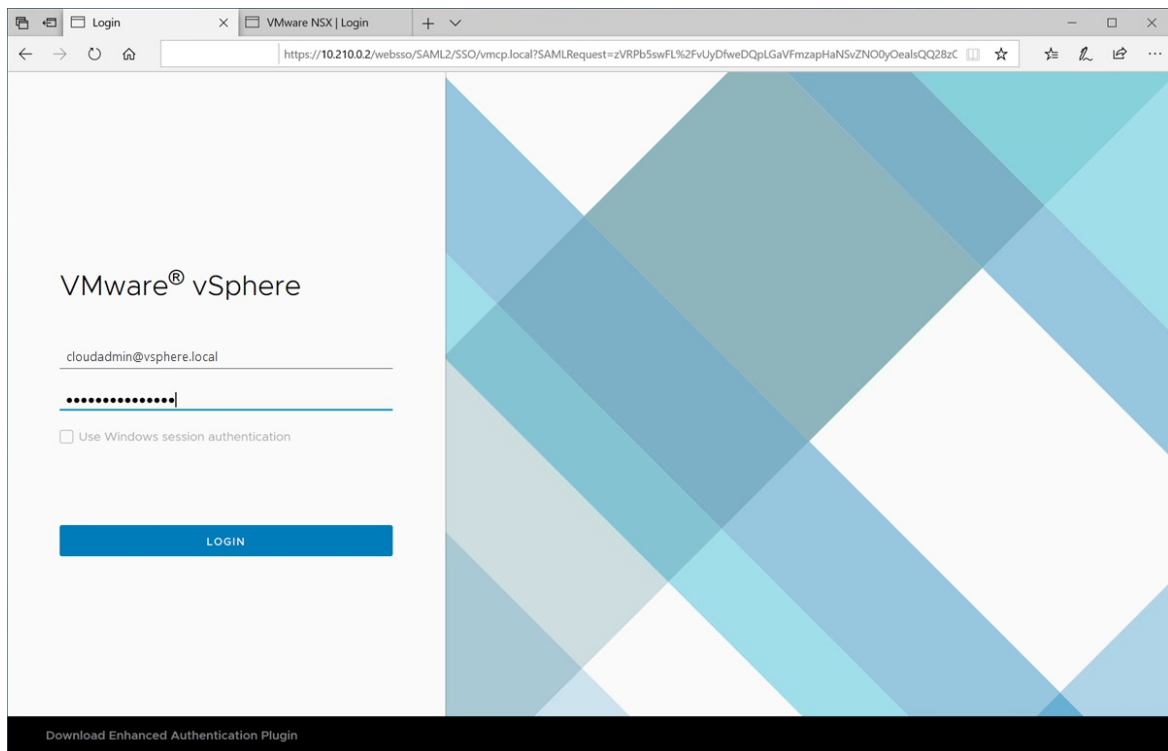
Make sure the file is accessible from the virtual machine.

Step 2. Access the local vCenter Server of your private cloud

1. Sign in to the [Azure portal](#), select your private cloud, and then **Manage > Identity**.
2. Copy the vCenter Server URL, username, and password. You'll use them to access your virtual machine (VM).
3. Select **Overview**, select the VM, and then connect to it through RDP. If you need help with connecting,

see [connect to a virtual machine](#) for details.

4. In the VM, open a browser and navigate to the vCenter URL.
5. Sign in with the `cloudadmin@vsphere.local` user credentials you copied earlier.



Step 3. Install the Bitnami OVA/OVF file in vCenter Server

1. Right-click the cluster that you want to install the LAMP virtual appliance and select **Deploy OVF Template**.
2. Select **Local file** and navigate to the OVF file you downloaded earlier. Then select **Next**.
3. Select your data center and provide a name for your virtual appliance VM, for example, **bitnami-lampstack**. Then select **Next**.
4. Select the ESXi host as the compute resource to run your VM and then select **Next**.
5. Review the details and select **Next**.
6. Accept the license agreement and select **Next**.
7. Select the storage for your VM and select **Next**.
8. Select the destination network for your VM and select **Next**.
9. Provide the required information to customize the template, such as the VM and networking properties. Then select **Next**.
10. Review the configuration settings and then select **Finish**.
11. From the **Task Console**, verify that the status of the OVF template deployment has completed successfully.
12. After the installation finishes, under **Actions**, select **Power on** to turn on the appliance.
13. From the vCenter Server console, select **Launch Web Console** and sign in to the Bitnami virtual appliance. Check the [Bitnami virtual appliance support documentation](#) for the default username and password.

NOTE

You can change the default password to a more secure one. For more information, see ...

Step 4. Assign a static IP to the virtual appliance

In this step, you'll modify the *bootproto* and *onboot* parameters and assign a static IP address to the Bitnami virtual appliance.

1. Search for the network configuration file.

```
sudo find /etc -name \*ens160\*
```

2. Edit the */etc/sysconfig/network-scripts/ifcfg-ens160* file and modify the boot parameters. Then add the static IP, netmask, and gateway addresses.

- `bootproto=static`
- `onboot=yes`

3. View and confirm the changes to the *ifcfg-ens160* file.

```
cat ifcfg-ens160
```

4. Restart the networking service. This stops the networking services first and then applies the IP configuration.

```
sudo systemctl restart network
```

5. Ping the gateway IP address to verify the configuration and VM connectivity to the network.
6. Confirm that the default route 0.0.0.0 is listed.

```
sudo route -n
```

Step 5. Enable SSH access to the virtual appliance

In this step, you'll enable SSH on your virtual appliance for remote access control. The SSH service is disabled by default. You'll also use PuTTY to connect to the host console.

1. Enable and start the SSH service.

```
sudo rm /etc/ssh/sshd_not_to_be_run
sudo systemctl enable sshd
sudo systemctl start sshd
```

2. Edit the */etc/ssh/sshd_config* file to change the password authentication.

```
PasswordAuthentication yes
```

3. View and confirm the changes to the *sshd_config* file.

```
sudo cat sshd_config
```

4. Reload the changes made to the file.

```
sudo /etc/init.d/ssh force-reload
```


5. Open PuTTY, select the SSH option and provide the host name and *22 for the port. Then select **Open**.
6. At the virtual appliance console prompt, enter the Bitnami username and password to connect to the host.

Create a HCX network extension

12/16/2022 • 2 minutes to read • [Edit Online](#)


This is an optional step to extend any networks from your on-premises environment to Azure VMware Solution.

1. Under **Services**, select **Network Extension** > **Create a Network Extension**.




2. Select each of the networks you want to extend to Azure VMware Solution, and then select **Next**.

Extend Networks Select source networks for extension to remote site



3. Enter the on-premises gateway IP for each of the networks you're extending, and then select **Submit**.

Extend Networks Select source networks for extension to remote site



It takes a few minutes for the network extension to finish. When it does, you see the status change to **Extension complete**.

Extensions: 1

+EXTEND NETWORKS

1

Transport Zones / DVS

Extension Appliance	Status
ServiceMesh-NE-11	✓ Extension complete

1 extension

Next steps

Now that you've configured the HCX Network Extension, you can also learn about:

- [VMware HCX Mobility Optimized Networking \(MON\) guidance](#)

HCX Network extension high availability (HA)

12/16/2022 • 3 minutes to read • [Edit Online](#)

VMware HCX is an application mobility platform that's designed to simplify application migration, workload rebalancing, and business continuity across data centers and clouds.

The HCX Network Extension service provides layer 2 connectivity between sites. Network Extension HA protects extended networks from a Network Extension appliance failure at either the source or remote site.

HCX 4.3.0 or later allows network extension high availability. Network Extension HA operates in Active/Standby mode. In this article, you'll learn how to configure HCX network extension High Availability on Azure private cloud.

Prerequisites

The Network Extension High Availability (HA) setup requires four Network Extension appliances, with two appliances at the source site and two appliances at the remote site. Together, these two pairs form the HA Group, which is the mechanism for managing Network Extension High Availability. Appliances on the same site require a similar configuration and must have access to the same set of resources.

- Network Extension HA requires an HCX Enterprise license.
- In the HCX Compute Profile, the Network Extension Appliance Limit is set to allow for the number of Network Extension appliances. The Azure VMware Solutions Limit is automatically set to unlimited.
- In the HCX Service Mesh, the Network Extension Appliance Scale Out Appliance Count is set to provide enough appliances to support network extension objectives, including any Network Extension HA groups.

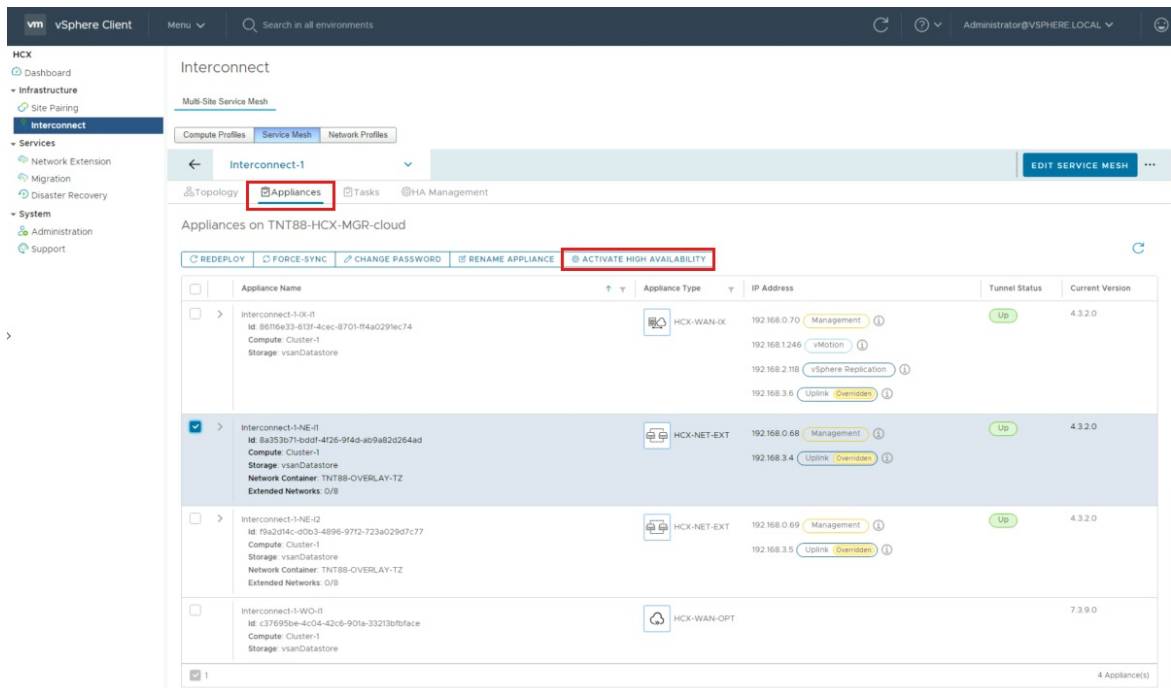
When you create a service mesh, set the appliance count to a minimum of two. For an existing service mesh, you can edit and adjust the appliance count to provide the required appliance count.

- The Network Extension appliances selected for HA activation must have no networks extended over them.
- Only Network Extension appliances upgraded to HCX 4.3.0 or later can be added to HA Groups.
- Learn more about the [Network Extension High Availability](#) feature, prerequisites, considerations and limitations.

Activate high availability (HA)

Use the following steps to activate HA, create HA groups, and view the HA roles and options available.

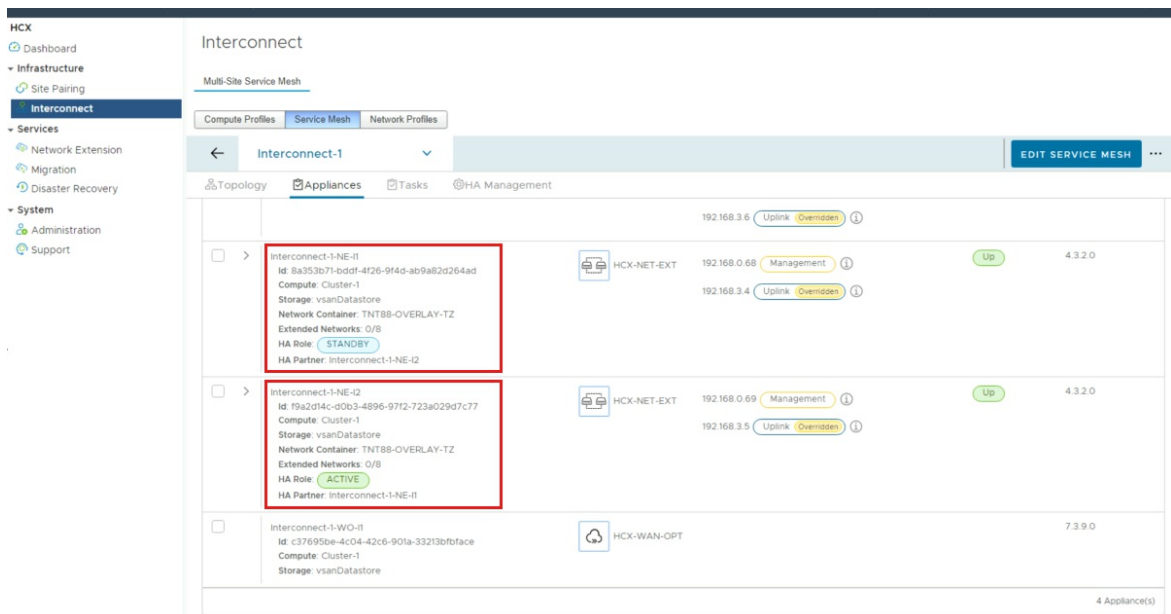
1. Sign in to HCX Manager UI in one of two ways:
 - a. cloudadmin@vsphere.local.
 - b. HCX UI through vCenter HCX Plugin.
2. Navigate to **Infrastructure**, then **Interconnect**.
3. Select **Service Mesh**, then select **View Appliances**.
4. Select **Appliances** from the **Interconnect** tab options.
 - a. Check the network appliance that you want to make highly available and select **Activate High Availability**.



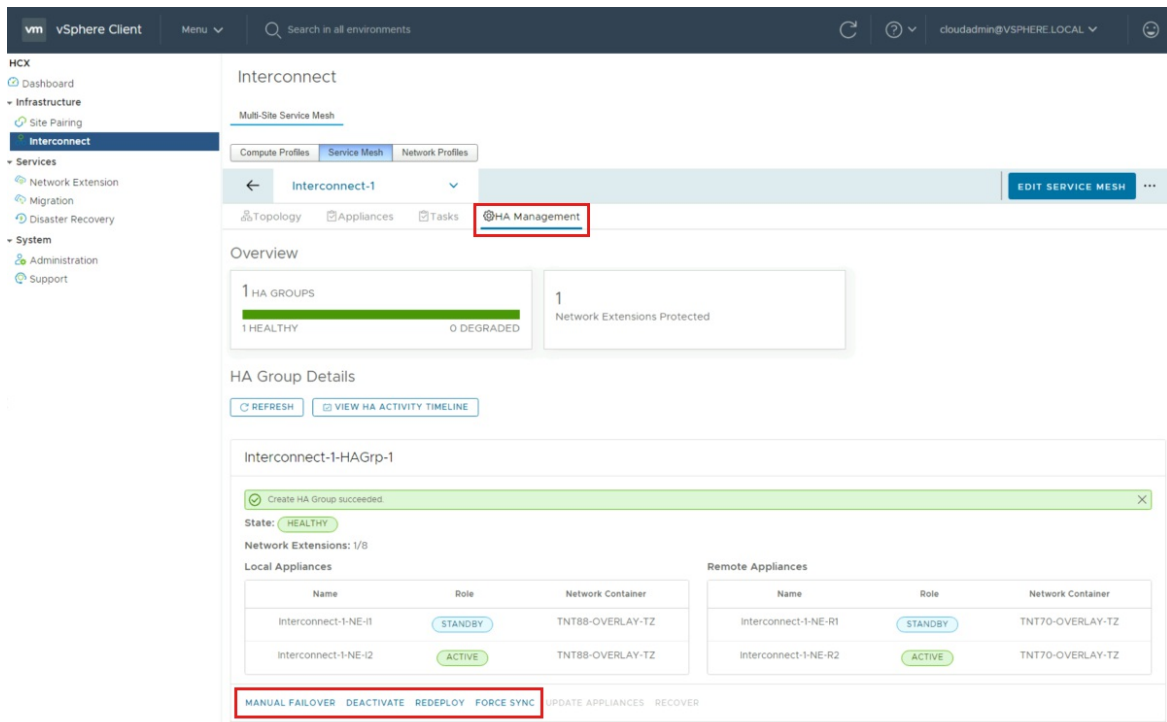
5. Confirm by selecting **Activate HA**.

a. Activating HA initiates the process to create an HA group. The process automatically selects an HA partner from the available NE Appliances.

6. After the HA group is created, the HA Roles for the local and remote appliances display **Active** and **Standby**.

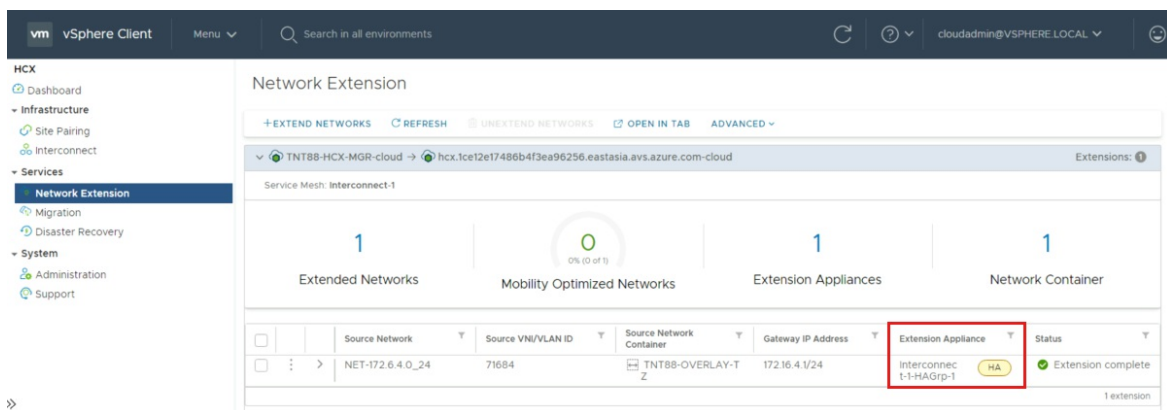


7. Select **HA Management** from the **Interconnect** tab options to view the HA group details and the available options: **Manual failover**, **Deactivate**, **Redeploy**, and **Force Sync**.



Extend network using network HA group

1. Locate **Services** in the left navigation and select **Network Extension**.
2. Select **Create a Network Extension**.
3. Choose the Network you want and select **Next**.
4. In **mandatory fields**, provide the gateway IP address in CIDR format, select the HA group under **Extension Appliances** (this was created in the previous step), and select **Submit** to extend the Network.
5. After the network is extended, under **Extension Appliance**, you can see the extension details and HA group.



6. To migrate virtual machines (VMs), navigate to **Services** and select **Migration**.
 - a. Select **Migrate** from the **Migration** window to start the workload mobility wizard.
7. In **Workload Mobility**, add and replace details as needed, then select **Validate**.
8. After validation completes, select **Go** to start the migration using Extended Network.

Workload Mobility

Remote Site Connection: Reverse Migration

Source: TNT88-HCX-MGR-cloud / VC: 192.168.0.2 → Destination: hcx.1ce12e17486b4f3ea96256.eastasia.azure.com-cloud / VC: 192.168.192.2 Reload Connections

Group Name: _____ Batch size: 1 VM / 40 GB / 4 GB / 2 vCPU Select VMs for Migration

Transfer and Placement:

Cluster-1 vsanDatastore (33.8 TB / 41.9 TB) Bulk Migration

Discovered virtual machine Same format as source (Optional: Switchover Schedule)

Switchover:

Force Power-off VM Remove Snapshots

Force unmount ISO Images

Extended Options:

Edit Extended Options Retain MAC

VM for Migration	Disk / Memory / vCPU	Migration Info
1 windows2019	40 GB / 4 GB / 2 vCPU	Bulk Migration

GO VALIDATE SAVE CLOSE

Next steps

Now that you've learned how to configure and extend HCX network extension high availability (HA), use the following resource to learn more about how to manage HCX network extension HA.

[Managing Network Extension High Availability](#)

Configure vRealize Operations for Azure VMware Solution

12/16/2022 • 3 minutes to read • [Edit Online](#)

vRealize Operations is an operations management platform that allows VMware infrastructure administrators to monitor system resources. These system resources could be application-level or infrastructure level (both physical and virtual) objects. Most VMware administrators have used vRealize Operations to monitor and manage the VMware private cloud components – vCenter Server, ESXi, NSX-T Data Center, vSAN, and VMware HCX. Each provisioned Azure VMware Solution private cloud includes a dedicated vCenter Server, NSX-T Data Center, vSAN, and HCX deployment.

Thoroughly review [Before you begin](#) and [Prerequisites](#) first. Then, we'll walk you through the three typical deployment topologies:

- [On-premises vRealize Operations managing Azure VMware Solution deployment](#)
- [vRealize Operations Cloud managing Azure VMware Solution deployment](#)
- [vRealize Operations running on Azure VMware Solution deployment](#)

Before you begin

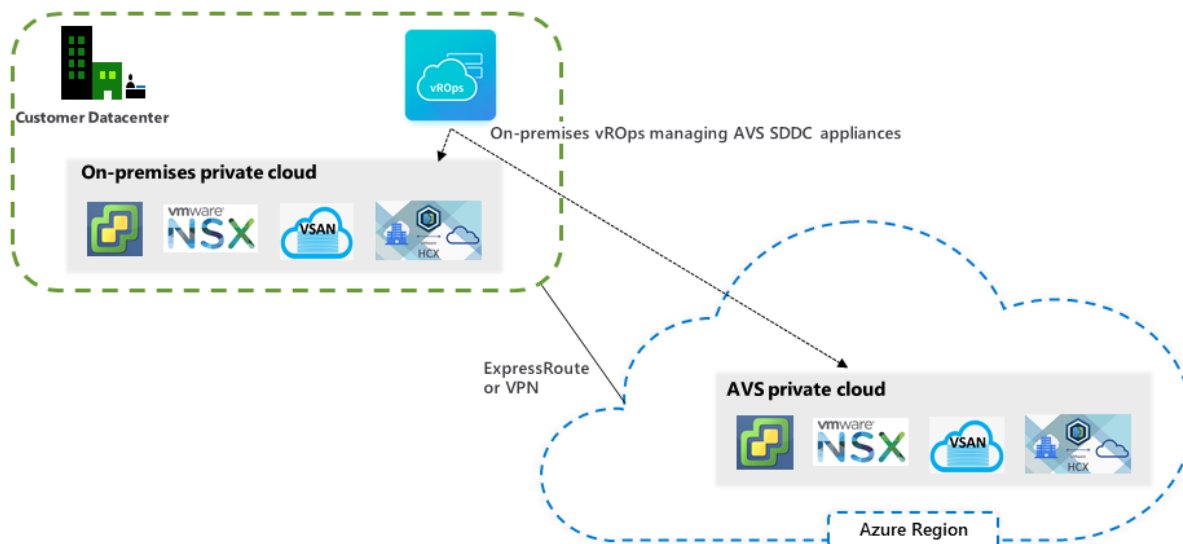
- Review the [vRealize Operations Manager product documentation](#) to learn more about deploying vRealize Operations.
- Review the basic Azure VMware Solution Software-Defined Datacenter (SDDC) [tutorial series](#).
- Optionally, review the [vRealize Operations Remote Controller](#) product documentation for the on-premises vRealize Operations managing Azure VMware Solution deployment option.

Prerequisites

- [vRealize Operations Manager](#) installed.
- A VPN or an Azure ExpressRoute configured between on-premises and Azure VMware Solution SDDC.
- An Azure VMware Solution private cloud has been deployed in Azure.

On-premises vRealize Operations managing Azure VMware Solution deployment

Most customers have an existing on-premises deployment of vRealize Operations to manage one or more on-premises vCenter Server domains. When they provision an Azure VMware Solution private cloud, they connect their on-premises environment with their private cloud using an Azure ExpressRoute or a Layer 3 VPN solution.



To extend the vRealize Operations capabilities to the Azure VMware Solution private cloud, you create an adapter [instance for the private cloud resources](#). It collects data from the Azure VMware Solution private cloud and brings it into on-premises vRealize Operations. The on-premises vRealize Operations Manager instance can directly connect to the vCenter Server and NSX-T Manager on Azure VMware Solution. Optionally, you can deploy a vRealize Operations Remote Collector on the Azure VMware Solution private cloud. The collector compresses and encrypts the data collected from the private cloud before it's sent over the ExpressRoute or VPN network to the vRealize Operations Manager running on-premises.

TIP

Refer to the [VMware documentation](#) for step-by-step guide for installing vRealize Operations Manager.

vRealize Operations Cloud managing Azure VMware Solution deployment

VMware vRealize Operations Cloud supports the Azure VMware Solution, including the vCenter Server, vSAN and NSX-T Data Center adapters.

IMPORTANT

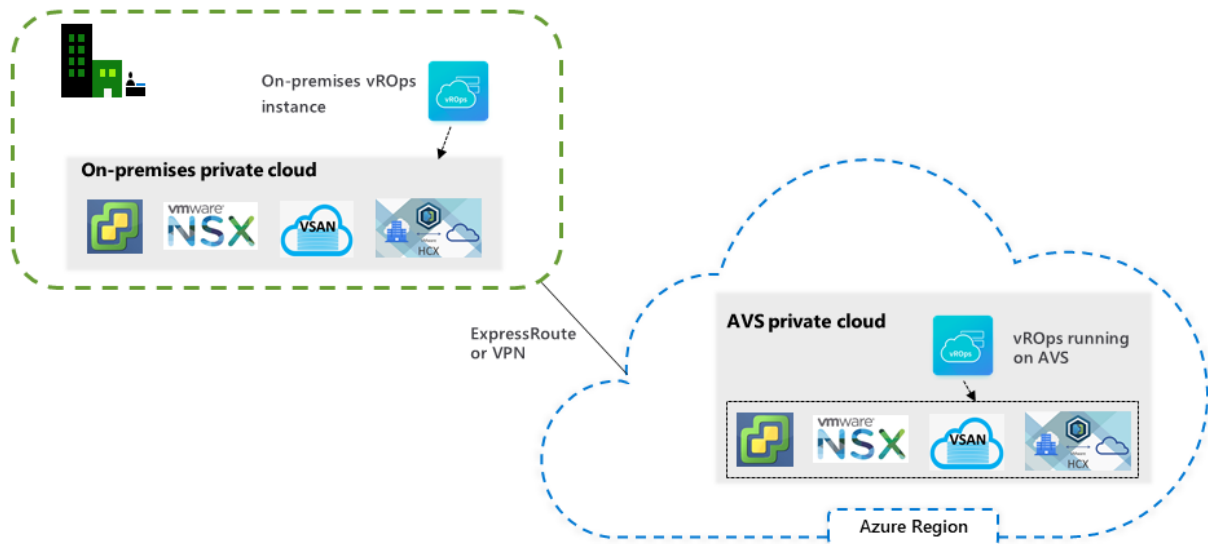
Refer to the [VMware documentation](#) for step-by-step guide for connecting vRealize Operations Cloud to Azure VMware Solution.

vRealize Operations running on Azure VMware Solution deployment

Another option is to deploy an instance of vRealize Operations Manager on a vSphere cluster in the private cloud.

IMPORTANT

This option isn't currently supported by VMware.



Once the instance has been deployed, you can configure vRealize Operations to collect data from vCenter Server, ESXi, NSX-T Data Center, vSAN, and HCX.

Known limitations

- The `cloudadmin@vsphere.local` user in Azure VMware Solution has [limited privileges](#). Virtual machines (VMs) on Azure VMware Solution doesn't support in-guest memory collection using VMware tools. Active and consumed memory utilization continues to work in this case.
- Workload optimization for host-based business intent doesn't work because Azure VMware Solutions manage cluster configurations, including DRS settings.
- Workload optimization for the cross-cluster placement within the SDDC using the cluster-based business intent is fully supported with vRealize Operations Manager 8.0 and onwards. However, workload optimization isn't aware of resource pools and places the VMs at the cluster level. A user can manually correct it in the Azure VMware Solution vCenter Server interface.
- You can't sign in to vRealize Operations Manager using your Azure VMware Solution vCenter Server credentials.
- Azure VMware Solution doesn't support the vRealize Operations Manager plugin.

When you connect the Azure VMware Solution vCenter Server to vRealize Operations Manager using a vCenter Server Cloud Account, you'll see a warning:

Warning



Adapter instance creation succeeded, but vCenter Server registration failed.

OK

The warning occurs because the `cloudadmin@vsphere.local` user in Azure VMware Solution doesn't have sufficient privileges to do all vCenter Server actions required for registration. However, the privileges are sufficient for the adapter instance to do data collection, as seen below:

Name	Adapter Type	Object Type	Collection State	Collection Status
vc-avs	vCenter Adapter	vCenter Server		

For more information, see [Privileges Required for Configuring a vCenter Server Adapter Instance](#).

NOTE

VMware vRealize Automation(vRA) integration with the NSX-T Data Center component of the Azure VMware Solution requires the "auditor" role to be added to the user with the NSX-T Manager cloudadmin role.

Deploy Horizon on Azure VMware Solution

12/16/2022 • 11 minutes to read • [Edit Online](#)

NOTE

This document focuses on the VMware Horizon product, formerly known as Horizon 7. Horizon is a different solution than Horizon Cloud on Azure, although there are some shared components. Key advantages of the Azure VMware Solution include both a more straightforward sizing method and the integration of VMware Cloud Foundation management into the Azure portal.

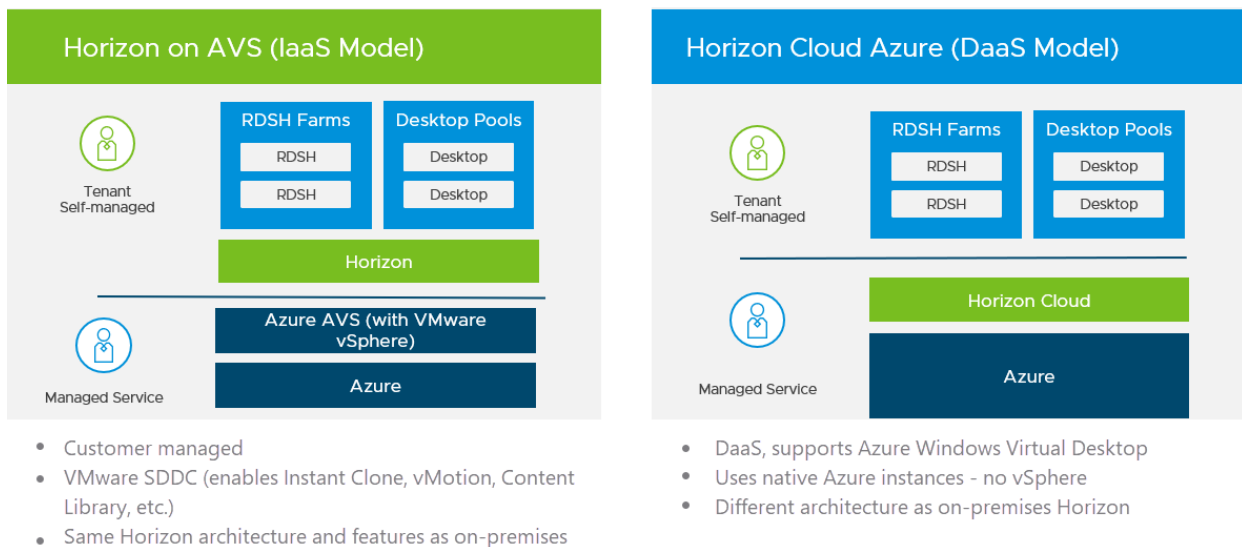
VMware Horizon®, a virtual desktop and applications platform, runs in the data center and provides simple and centralized management. It delivers virtual desktops and applications on any device, anywhere. Horizon lets you create, and broker connections to Windows and Linux virtual desktops, Remote Desktop Server (RDS) hosted applications, desktops, and physical machines.

Here, we focus specifically on deploying Horizon on Azure VMware Solution. For general information on VMware Horizon, refer to the Horizon production documentation:

- [What is VMware Horizon?](#)
- [Learn more about VMware Horizon](#)
- [Horizon Reference Architecture](#)

With Horizon's introduction on Azure VMware Solution, there are now two Virtual Desktop Infrastructure (VDI) solutions on the Azure platform. The following diagram summarizes the key differences at a high level.

Differences between VMware Horizon on Azure VMware Solution and VMware Horizon Cloud on Azure



Horizon 2006 and later versions on the Horizon 8 release line supports both on-premises and Azure VMware Solution deployment. There are a few Horizon features that are supported on-premises but not on Azure VMware Solution. Other products in the Horizon ecosystem are also supported. For more information, see [feature parity and interoperability](#).

Deploy Horizon in a hybrid cloud

You can deploy Horizon in a hybrid cloud environment by using Horizon Cloud Pod Architecture (CPA) to interconnect on-premises and Azure data centers. CPA scales up your deployment, builds a hybrid cloud, and provides redundancy for Business Continuity and Disaster Recovery. For more information, see [Expanding Existing Horizon 7 Environments](#).

IMPORTANT

CPA is not a stretched deployment; each Horizon pod is distinct, and all Connection Servers that belong to each of the individual pods are required to be located in a single location and run on the same broadcast domain from a network perspective.

Like on-premises or private data centers, you can deploy Horizon in an Azure VMware Solution private cloud. We'll discuss key differences in deploying Horizon on-premises and Azure VMware Solution in the following sections.

The *Azure private cloud* is conceptually the same as the *VMware SDDC*, a term typically used in Horizon documentation. The rest of this document uses both terms interchangeably.

The Horizon Cloud Connector is required for Horizon on Azure VMware Solution to manage subscription licenses. You can deploy Cloud Connector in Azure Virtual Network alongside Horizon Connection Servers.

IMPORTANT

Horizon Control Plane support for Horizon on Azure VMware Solution is not yet available. Be sure to download the VHD version of Horizon Cloud Connector.

vCenter Server Cloud Admin role

Since Azure VMware Solution is an SDDC service and Azure manages the lifecycle of the SDDC on Azure VMware Solution, the vCenter Server permission model on Azure VMware Solution is limited by design.

Customers are required to use the Cloud Admin role, which has a limited set of vCenter Server permissions. The Horizon product was modified to work with the Cloud Admin role on Azure VMware Solution, specifically:

- Instant clone provisioning was modified to run on Azure VMware Solution.
- A specific vSAN policy (VMware_Horizon) was created on Azure VMware Solution to work with Horizon, which must be available and used in the SDDCs deployed for Horizon.
- vSphere Content-Based Read Cache (CBRC), also known as View Storage Accelerator, is disabled when running on the Azure VMware Solution.

IMPORTANT

CBRC must not be turned back on.

NOTE

Azure VMware Solution automatically configures specific Horizon settings as long as you deploy Horizon 2006 (aka Horizon 8) and above on the Horizon 8 branch and select the **Azure** option in the Horizon Connection Server installer.

Horizon on Azure VMware Solution deployment architecture

A typical Horizon architecture design uses a pod and block strategy. A block is a single vCenter Server, while multiple blocks combined make a pod. A Horizon pod is a unit of organization determined by Horizon scalability limits. Each Horizon pod has a separate management portal, and so a standard design practice is to minimize the number of pods.

Every cloud has its own network connectivity scheme. Combined with VMware SDDC networking / NSX-T Data Center, the Azure VMware Solution network connectivity presents unique requirements for deploying Horizon that is different from on-premises.

Each Azure private cloud and SDDC can handle 4,000 desktop or application sessions, assuming:


- The workload traffic aligns with the LoginVSI task worker profile.
- Only protocol traffic is considered, no user data.
- NSX Edge is configured to be large.

NOTE

Your workload profile and needs may be different, and therefore results may vary based on your use case. User Data volumes may lower scale limits in the context of your workload. Size and plan your deployment accordingly. For more information, see the sizing guidelines in the [Size Azure VMware Solution hosts for Horizon deployments](#) section.

Given the Azure private cloud and SDDC max limit, we recommend a deployment architecture where the Horizon Connection Servers and VMware Unified Access Gateways (UAGs) are running inside the Azure Virtual Network. It effectively turns each Azure Desktop private cloud and SDDC into a block. In turn, maximizing the scalability of Horizon running on Azure VMware Solution.


The connection from Azure Virtual Network to the Azure private clouds / SDDCs should be configured with ExpressRoute FastPath. The following diagram shows a basic Horizon pod deployment.



Network connectivity to scale Horizon on Azure VMware Solution

This section lays out the network architecture at a high level with some common deployment examples to help you scale Horizon on Azure VMware Solution. The focus is specifically on critical networking elements.

Single Horizon pod on Azure VMware Solution



A single Horizon pod is the most straight forward deployment scenario because you deploy just one Horizon pod in the US East region. Since each private cloud and SDDC is estimated to handle 4,000 desktop sessions, you deploy the maximum Horizon pod size. You can plan the deployment of up to three private clouds/SDDCs.

With the Horizon infrastructure virtual machines (VMs) deployed in Azure Virtual Network, you can reach the 12,000 sessions per Horizon pod. The connection between each private cloud and SDDC to the Azure Virtual Network is ExpressRoute Fast Path. No east-west traffic between private clouds is needed.

Key assumptions for this basic deployment example include that:


- You don't have an on-premises Horizon pod that you want to connect to this new pod using Cloud Pod Architecture (CPA).
- End users connect to their virtual desktops through the internet (vs. connecting via an on-premises datacenter).

You connect your AD domain controller in Azure Virtual Network with your on-premises AD through VPN or ExpressRoute circuit.

A variation on the basic example might be to support connectivity for on-premises resources. For example, users access desktops and generate virtual desktop application traffic or connect to an on-premises Horizon pod using CPA.

The diagram shows how to support connectivity for on-premises resources. To connect to your corporate network to the Azure Virtual Network, you'll need an ExpressRoute circuit. You'll also need to connect your corporate network with each of the private cloud and SDDCs using ExpressRoute Global Reach. It allows the connectivity from the SDDC to the ExpressRoute circuit and on-premises resources.

Region 1- US East Horizon POD 1

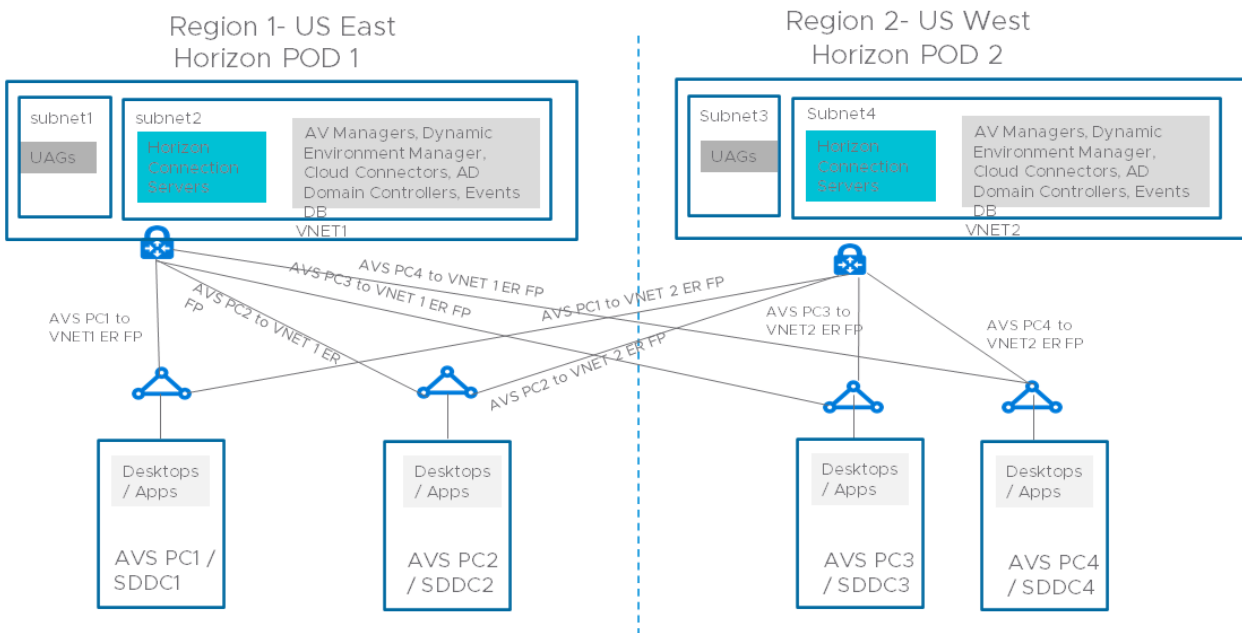


Multiple Horizon pods on Azure VMware Solution across multiple regions

Another scenario is scaling Horizon across multiple pods. In this scenario, you deploy two Horizon pods in two different regions and federate them using CPA. It's similar to the network configuration in the previous example, but with some additional cross-regional links.

You'll connect the Azure Virtual Network in each region to the private clouds/SDDCs in the other region. It allows Horizon connection servers part of the CPA federation to connect to all desktops under management. Adding extra private clouds/SDDCs to this configuration would allow you to scale to 24,000 sessions overall.

The same principles apply if you deploy two Horizon pods in the same region. Make sure to deploy the second Horizon pod in a *separate Azure Virtual Network*. Just like the single pod example, you can connect your corporate network and on-premises pod to this multi-pod/region example using ExpressRoute and Global Reach.



Size Azure VMware Solution hosts for Horizon deployments

Horizon's sizing methodology on a host running in Azure VMware Solution is simpler than Horizon on-premises. That's because the Azure VMware Solution host is standardized. Exact host sizing helps determine the number of hosts needed to support your VDI requirements. It's central to determining the cost-per-desktop.

Sizing tables

Specific vCPU/vRAM requirements for Horizon virtual desktops depend on the customer's specific workload profile. Work with your MSFT and VMware sales team to help determine your vCPU/vRAM requirements for your virtual desktops.

V C P U P E R F O R M A N C E	V R A M P E R M E N T (G B)	IN S T A N C E	10	20	30	40	50	60	70	80	90	100	200	300	400	500	600	6400
			V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S	V M S
2	3.5	A V S	3	3	4	4	5	6	6	7	8	9	17	25	33	41	49	53
2	4	A V S	3	3	4	5	6	6	7	8	9	9	18	26	34	42	50	54
2	6	A V S	3	4	5	6	7	9	10	11	12	13	22	33	51	62	73	79
2	8	A V S	3	5	6	8	9	11	12	14	16	18	34	51	68	84	100	106
2	12	A V S	4	6	9	11	13	16	19	22	26	30	55	77	100	124	148	158
2	16	A V S	5	8	11	14	18	22	27	32	38	44	66	90	115	140	165	176
4	3.5	A V S	3	3	4	5	6	7	8	9	10	11	22	33	44	55	66	70
4	4	A V S	3	3	4	5	6	7	8	9	10	11	22	33	44	55	66	70
4	6	A V S	3	4	5	6	7	9	10	11	12	13	26	38	51	64	77	82

V C P U P E R V M	V R A M P E R V M (G B)	I N S T A N C E	10	20	30	40	50	60	70	80	90	10	20	30	40	50	60	64
			0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	0 V M S	00 V M S	00 V M S	00 V M S	00 V M S	00 V M S
4	8	A V S	3	5	6	8	9	1 1	1 2	1 4	1 6	1 8	3 4	5 1	6 7	8 4	1 0 0	1 0 6
4	1 2	A V S	4	6	9	1 1	1 3	1 6	1 9	2 1	2 3	2 6	5 1	7 5	1 0 0	1 2 4	1 4 9	1 5 8
4	1 6	A V S	5	8	1 1	1 4	1 8	2 1	2 4	2 7	3 0	3 4	6 7	1 0 0	1 3 3	1 6 5	1 9 8	2 1 1
6	3. 5	A V S	3	4	5	6	7	9	1 0	1 1	1 3	1 4	2 7	4 1	5 4	6 8	8 1	8 6
6	4	A V S	3	4	5	6	7	9	1 0	1 1	1 3	1 4	2 7	4 1	5 4	6 8	8 1	8 6
6	6	A V S	3	4	5	6	7	9	1 0	1 1	1 3	1 4	2 7	4 1	5 4	6 8	8 1	8 6
6	8	A V S	3	5	6	8	9	1 1	1 2	1 4	1 6	1 8	3 4	5 1	6 7	8 4	1 0 0	1 0 6
6	1 2	A V S	4	6	9	1 1	1 3	1 6	1 9	2 1	2 3	2 6	5 1	7 5	1 0 0	1 2 4	1 4 9	1 5 8
6	1 6	A V S	5	8	1 1	1 4	1 8	2 1	2 4	2 7	3 0	3 4	6 7	1 0 0	1 3 3	1 6 5	1 9 8	2 1 1
8	3. 5	A V S	3	4	6	7	9	1 0	1 2	1 4	1 5	1 7	3 3	4 9	6 6	8 2	9 8	1 0 5
8	4	A V S	3	4	6	7	9	1 0	1 2	1 4	1 5	1 7	3 3	4 9	6 6	8 2	9 8	1 0 5
8	6	A V S	3	4	6	7	9	1 0	1 2	1 4	1 5	1 7	3 3	4 9	6 6	8 2	9 8	1 0 5

V C P U P E R V M	R A M P E S (G B)	IN S T R U C T I O N S	10	20	30	40	50	60	70	80	90	10	20	30	40	50	60	64
			0	0	0	0	0	0	0	0	0	0	00	00	00	00	00	00
8	8	A V S	3	5	6	8	9	1	1	1	1	1	3	5	6	8	1	1
8	1 2	A V S	4	6	9	1	1	1	1	2	2	2	5	7	1	1	1	1
8	1 6	A V S	5	8	1	1	1	2	2	2	3	3	6	1	1	1	1	2

Horizon sizing inputs

Here's what you'll need to gather for your planned workload:

- Number of concurrent desktops
- Required vCPU per desktop
- Required vRAM per desktop
- Required storage per desktop

In general, VDI deployments are either CPU or RAM constrained, which determines the host size. Let's take the following example for a LoginVSI Knowledge Worker type of workload, validated with performance testing:

- 2,000 concurrent desktop deployment
- 2vCPU per desktop.
- 4-GB vRAM per desktop.
- 50 GB of storage per desktop

For this example, the total number of hosts factors out to 18, yielding a VM-per-host density of 111.

IMPORTANT

Customer workloads will vary from this example of a LoginVSI Knowledge Worker. As a part of planning your deployment, work with your VMware EUC SEs for your specific sizing and performance needs. Be sure to run your own performance testing using the actual, planned workload before finalizing host sizing and adjust accordingly.

Horizon on Azure VMware Solution licensing

There are four components to the overall costs of running Horizon on Azure VMware Solution.

Azure VMware Solution Capacity Cost

For information on the pricing, see the [Azure VMware Solution pricing](#) page

Horizon Licensing Cost

There are two available licenses for use with the Azure VMware Solution, which can be either Concurrent User (CCU) or Named User (NU):

- Horizon Subscription License
- Horizon Universal Subscription License

If only deploying Horizon on Azure VMware Solution for the foreseeable future, then use the Horizon Subscription License as it is a lower cost.

If deployed on Azure VMware Solution and on-premises, choose the Horizon Universal Subscription License as a disaster recovery use case. However, it includes a vSphere license for on-premises deployment, so it has a higher cost.

Work with your VMware EUC sales team to determine the Horizon licensing cost based on your needs.

Azure Instance Types

To understand the Azure virtual machine sizes that are required for the Horizon Infrastructure, see [Horizon Installation on Azure VMware Solution](#).

References

[System Requirements For Horizon Agent for Linux](#)

Next steps

To learn more about VMware Horizon on Azure VMware Solution, read the [VMware Horizon FAQ](#).

Create a content library to deploy VMs in Azure VMware Solution

12/16/2022 • 2 minutes to read • [Edit Online](#)

A content library stores and manages content in the form of library items. A single library item consists of files you use to deploy virtual machines (VMs).

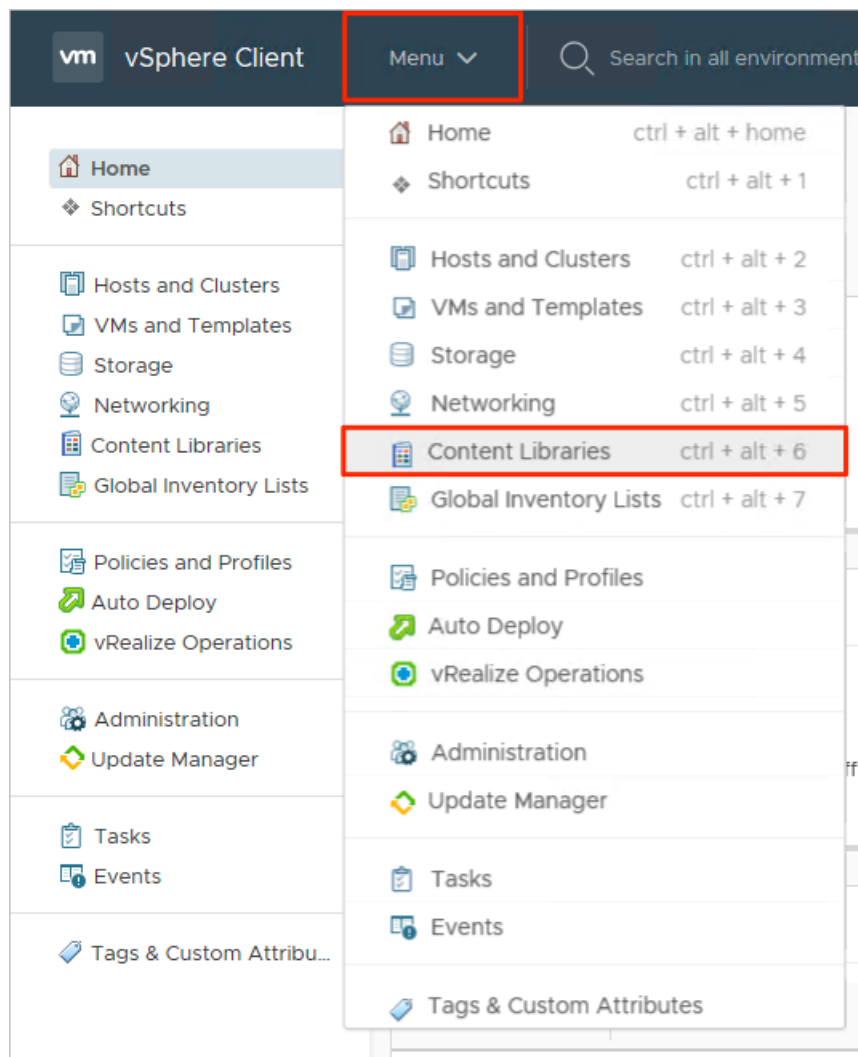
In this article, you'll create a content library in the vSphere Client and then deploy a VM using an ISO image from the content library.

Prerequisites

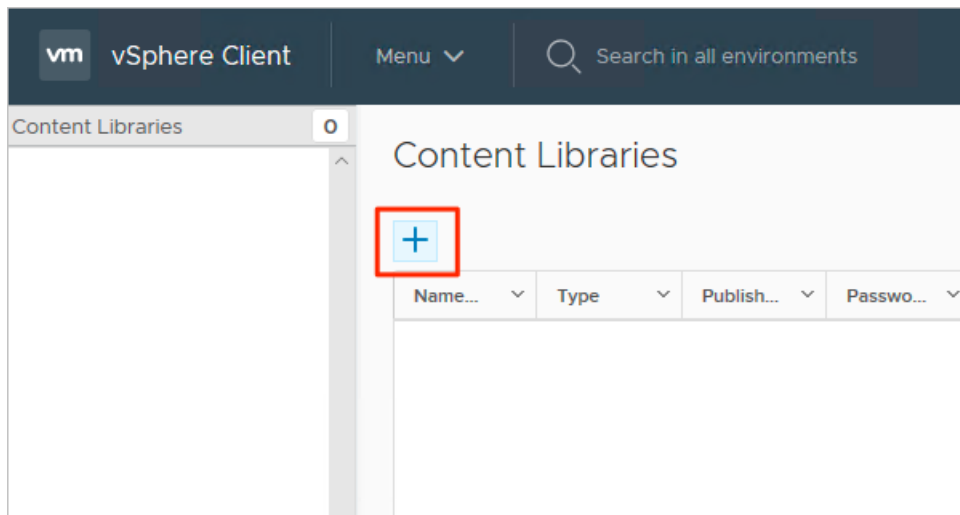
An NSX-T Data Center segment and a managed DHCP service are required to complete this tutorial. For more information, see [Configure DHCP for Azure VMware Solution](#).

Create a content library

1. From the on-premises vSphere Client, select **Menu** > **Content Libraries**.



2. Select **Add** to create a new content library.



3. Provide a name and confirm the IP address of the vCenter Server and select **Next**.

New Content Library

1 Name and location
2 Configure content library
3 Add storage
4 Ready to complete

Name and location
Specify content library name and location.

Name:

Notes:

vCenter Server:

[CANCEL](#) [BACK](#) [NEXT](#)

4. Select the **Local** content library and select **Next**.

New Content Library

- ✓ 1 Name and location
- 2 Configure content library**
- 3 Add storage
- 4 Ready to complete

Configure content library
Local libraries can be published externally and optimized for syncing over HTTP. Subscribed libraries originate from other published libraries.

Local content library

Publish externally

Optimize for syncing over HTTP
Once published, it cannot be reverted back to a local library and cannot be used to deploy virtual machines.

Enable authentication

Subscribed content library

Subscription URL:

Enable authentication

Download content immediately when needed

[CANCEL](#)
[BACK](#)
[NEXT](#)


5. Select the datastore for storing your content library, and then select **Next**.

New Content Library

- ✓ 1 Name and location
- ✓ 2 Configure content library
- 3 Add storage**
- 4 Ready to complete

Add storage
Select a storage location for the library contents. Use a file system backing for published content libraries to store the uploaded OVF packages. Use a datastore backing for local and subscribed content libraries to store content optimized for cloning.

Filter

Name ↑	Status	Type	Datastore...
 vsanDatastore	✓ Normal	vSAN	

1 items

[CANCEL](#)
[BACK](#)
[NEXT](#)

6. Review the content library settings and select **Finish**.

New Content Library

- ✓ 1 Name and location
- ✓ 2 Configure content library
- ✓ 3 Add storage
- 4 Ready to complete

Ready to complete
Review content library settings.

Name:	AVSV-Library
Notes:	
vCenter Server:	10.61.0.2
Type:	Local Content Library
Published:	No
Storage:	vsanDatastore

CANCEL
BACK
FINISH

Upload an ISO image to the content library

Now that you've created the content library, you can add an ISO image to deploy a VM to a private cloud cluster.

1. From the vSphere Client, select **Menu > Content Libraries**.
2. Right-click the content library you want to use for the new ISO and select **Import Item**.
3. Import a library item for the Source by doing one of the following, and then select **Import**:
 - a. Select **URL** and provide a URL to download an ISO.
 - b. Select **Local File** to upload from your local system.

TIP

Optional, you can define a custom item name and notes for the Destination.

4. Open the library and select the **Other Types** tab to verify that your ISO was uploaded successfully.

Deploy a VM to a private cloud cluster

1. From the vSphere Client, select **Menu > Hosts and Clusters**.
2. In the left panel, expand the tree and select a cluster.
3. Select **Actions > New Virtual Machine**.
4. Go through the wizard and modify the settings you want.
5. Select **New CD/DVD Drive > Client Device > Content Library ISO File**.
6. Select the ISO uploaded in the previous section and then select **OK**.
7. Select the **Connect** check box so the ISO is mounted at power-on time.
8. Select **New Network > Select dropdown > Browse**.
9. Select the **logical switch (segment)** and select **OK**.

10. Modify any other hardware settings and select **Next**.

11. Verify the settings and select **Finish**.

Next steps

Now that you've created a content library to deploy VMs in Azure VMware Solution, you may want to learn about:

- [Migrating VM workloads to your private cloud](#)
- [Integrating Azure native services in Azure VMware Solution](#)

Enable HCX access over the internet

12/16/2022 • 6 minutes to read • [Edit Online](#)

In this article, you'll learn how to perform HCX migration over a public IP address using Azure VMware Solution.

IMPORTANT

Before configuring a public IP on your Azure VMware Solution private cloud, consult your network administrator to understand the implications and the impact to your environment.

You'll also learn how to pair HCX sites and create service mesh from on-premises to an Azure VMware Solution private cloud using Public IP. The service mesh allows you to migrate a workload from an on-premises datacenter to an Azure VMware Solution private cloud over the public internet. This solution is useful when the customer isn't using ExpressRoute or VPN connectivity with the Azure cloud.

IMPORTANT

The on-premises HCX appliance should be reachable from the internet to establish HCX communication from on-premises to the Azure VMware Solution private cloud.

Configure public IP block

For HCX manager to be available over the public IP address, you'll need one public IP address for DNAT rule.

To perform HCX migration over the public internet, you'll need other IP addresses. You can have a /29 subnet to create minimum configuration when defining HCX network profile (usable IPs in subnet will be assigned to IX, NE appliances). You can choose a bigger subnet based on the requirements. You'll create an NSX-T segment using this public subnet. This segment can be used for creating HCX network profile.

NOTE

After assigning a subnet to NSX-T segment, you can't use an IP from that subnet to create a DNAT rule. Both subnets should be different.

Configure a Public IP block through portal by using the [Public IP feature of the Azure VMware Solution](#) private cloud.

Use public IP address for Cloud HCX Manager public access

Cloud HCX manager can be available over a public IP address by using a DNAT rule. However, since Cloud HCX manager is in the provider space, the null route is necessary to allow HCX Manager to route back to the client by way of the DNAT rule. It forces the NAT traffic through NSX-T Tier-0 router.

Add static null route to the Tier1 router

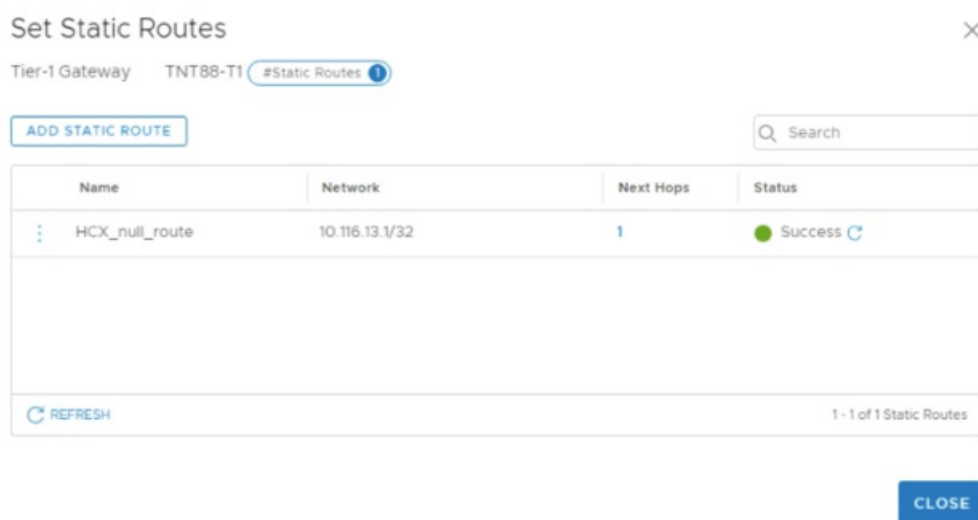
The static null route is used to allow HCX private IP to route through the NSX Tier-1 for public endpoints. This static route can be the default Tier-1 router created in your private cloud or you can create a new tier-1 router.

1. Sign in to NSX-T manager, and select **Networking**.

2. Under the **Connectivity** section, select **Tier-1 Gateways**.
3. Edit the existing Tier-1 gateway.
4. Expand **STATIC ROUTES**.
5. Select the number next to **Static Routes**.
6. Select **ADD STATIC ROUTE**.
A pop-up window is displayed.
7. Under **Name**, enter the name of the route.
8. Under **Network**, enter a non-overlapping /32 IP address under Network.

NOTE

This address should not overlap with any other IP addresses on the private cloud network and the customer network.



9. Under **Next hops**, select **Set**.
10. Select **NULL** as IP Address.
Leave defaults for Admin distance and scope.
11. Select **ADD**, then select **APPLY**.
12. Select **SAVE**, then select **CLOSE**.

Next Hops

Tier-1 Gateway TNT88-T1 | Static Route HCX_null_ro... #Next Hops 1

Search

IP Address	Admin Distance	Scope
NULL	1	None

CLOSE

13. Select **CLOSE EDITING**.

Add NAT rule to Tier-1 gateway

1. Sign in to NSX-T Manager, and select **Networking**.
2. Select **NAT**.
3. Select the Tier-1 Gateway. Use same Tier-1 router to create NAT rule that you used to create null route in previous steps.
4. Select **ADD NAT RULE**.
5. Add one SNAT rule and one DNAT rule for HCX Manager.
 - a. The DNAT Rule Destination is the Public IP for HCX Manager. The Translated IP is the HCX Manager IP in the cloud.
 - b. The SNAT Rule Destination is the HCX Manager IP in the cloud. The Translated IP is the non-overlapping /32 IP from the Static Route.
 - c. Make sure to set the Firewall option on DNAT rule to **Match External Address**.

NAT

Gateway TNT88-T1 | #Total NAT Rules 2 | View NAT

ADD NAT RULE | COLLAPSE ALL | Filter by Name, Path and more

	Name	Action	Match		Translated	Apply To	Enabled	Status
			Source	Destination				
⋮	HCX_dnat_rule	DNAT	Any	20.95.78.20	192.168.128.9	0	● Enabled	● Success
	Service	Any			Description	Not Set		
	Logging	● No			Translated Port	Any		
	Firewall	● Match External Address			Priority	0		
⋮	HCX_snat_rule	SNAT	Any	192.168.128.9	10.116.13.1/32	0	● Enabled	● Success
	Service	Any			Description	Not Set		
	Logging	● No			Translated Port	Any		
	Firewall	● Match Internal Address			Priority	0		

6. Create Tier-1 Gateway Firewall rules to allow only expected traffic to the Public IP for HCX Manager and drop everything else.
 - a. Create a Gateway Firewall rule on the T1 that allows your on-premises as the **Source IP** and the Azure VMware Solution reserved Public as the **Destination IP**. This rule should be the highest priority.
 - b. Create a Gateway Firewall rule on the Tier-1 that denies all other traffic where the **Source IP** is **Any**

and **Destination IP** is the Azure VMware Solution reserved Public IP.

For more information, see [HCX ports](#)

NOTE

HCX manager can now be accessed over the internet using public IP.

Pair sites using HCX Cloud manager's public IP address

Site pairing is required before you create service mesh between source and destination sites.

1. Sign in to the **Source** site HCX Manager.
2. Select **Site Pairing** and select **ADD SITE PAIRING**.
3. Enter the **Cloud HCX Manager Public URL** as remote site and sign in credentials, then select **Connect**.

After pairing is done, it will appear under site pairing.

Create public IP segment on NSX-T

Before you create a Public IP segment, get your credentials for NSX-T Manager from Azure VMware Solution portal.

1. Under the **Networking** section select **Connectivity, Segments**, and then select **ADD SEGMENT**.
2. Provide Segment name, select **Tier-1 router** as connected gateway, and provide the reserved public IP under subnets.
3. Select **Save**.

Create network profile for HCX at destination site

1. Sign in to Destination HCX Manager (cloud manager in this case).
2. Select **Interconnect** and then select the **Network Profiles** tab.
3. Select **Create Network Profile**.
4. Select **NSX Networks** as network type under **Network**.
5. Select the **Public-IP-Segment** created on NSX-T.
6. Enter **Name**.
7. Under IP pools, enter the **IP Ranges** for HCX uplink, **Prefix Length**, and **Gateway** of public IP segment.
8. Scroll down and select the **HCX Uplink** checkbox under **HCX Traffic Type** as this profile will be used for HCX uplink.
9. Select **Create** to create the network profile.

Create service mesh

Service Mesh will deploy HCX WAN Optimizer, HCX Network Extension and HCX-IX appliances.

1. Sign in to **Source** site HCX Manager.
2. Select **Interconnect** and then select the **Service Mesh** tab.
3. Select **CREATE SERVICE MESH**.
4. Select the **destination** site to create service mesh with and then select **Continue**.
5. Select the compute profiles for both sites and select **Continue**.
6. Select the HCX services to be activated and select **Continue**.

NOTE

Premium services require an additional HCX Enterprise license.

7. Select the network profile of source site.
8. Select the network profile of destination that you created in the **Network Profile** section.
9. Select **Continue**.
10. Review the **Transport Zone** information, and then select **Continue**.
11. Review the **Topological view**, and select **Continue**.
12. Enter the **Service Mesh name** and select **FINISH**.
13. Add the public IP addresses in firewall to allow required ports only.

Extend network

The HCX Network Extension service provides layer 2 connectivity between sites. The extension service also allows you to keep the same IP and MAC addresses during virtual machine migrations.

1. Sign in to **source** HCX Manager.
2. Under the **Network Extension** section, select the site for which you want to extend the network, and then select **EXTEND NETWORKS**.
3. Select the network that you want to extend to destination site, and select **Next**.
4. Enter the subnet details of network that you're extending.
5. Select the destination first hop route (Tier-1), and select **Submit**.
6. Sign in to the **destination** NSX, you'll see Network 10.14.27.1/24 has been extended.

After the network is extended to destination site, VMs can be migrated over Layer 2 extension.

Next steps

[Enable Public IP to the NSX Edge for Azure VMware Solution](#)

For detailed information on HCX network underlay minimum requirements, see [Network Underlay Minimum Requirements](#).

Enable SQL Azure hybrid benefit for Azure VMware Solution (Preview)

12/16/2022 • 2 minutes to read • [Edit Online](#)

In this article, you'll learn how to configure SQL Azure hybrid benefits to an Azure VMware Solution private cloud by configuring a placement policy. The placement policy defines the hosts that are running SQL as well as the virtual machines on that host.

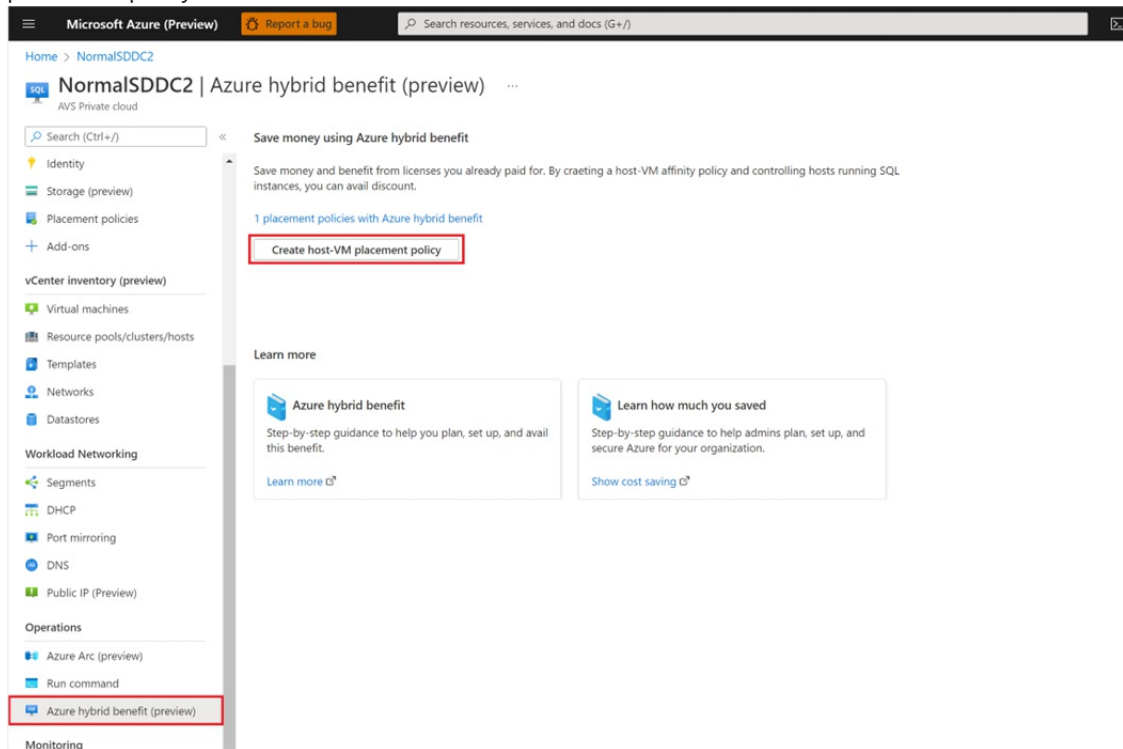
IMPORTANT

It is important to note that SQL benefits are applied at the host level.

For example, if each host in Azure VMware Solution has 36 cores and you signal that two hosts run SQL, then SQL Azure hybrid benefit will apply to 72 cores irrespective of the number of SQL or other virtual machines on that host.

Configure host-VM placement policy

1. From your Azure VMware Solution private cloud, select Azure hybrid benefit, then Create host-VM placement policy.



2. Fill in the required fields for creating the placement policy.
 - a. **Name** – Select the name that identifies this policy.
 - b. **Type** – Select the type of policy. This type must be a VM-Host affinity rule only.
 - c. **Azure hybrid benefit** – Select the checkbox to apply the SQL Azure hybrid benefit.
 - d. **Cluster** – Select the correct cluster. The policy is scoped to host in this cluster only.
 - e. **Enabled** – Select enabled to apply the policy immediately once created.

Microsoft Azure (Preview) [Report a bug](#)

Home > NormalSDDC2 > **Create placement policy** ...

Basics Review and Create

Placement policies are a way of specifying constraints for running virtual machines in an Azure VMware Solution private cloud. [Learn more](#)

Details

Name: *

Type: *
 This policy decides which VMs should (but are not required) to run on any of the selected hosts

Azure hybrid benefit: Use this policy to ensure host to VM affinity. I confirm I have eligible SQL Server licenses with software assurance or SQL server subscription to apply this Azure hybrid benefit. [Learn more](#)

Cluster: *

Enabled: * Enabled Disabled

3. Select the hosts and VMs that will be applied to the VM-Host affinity policy.
 - a. **Add Hosts** – Select the hosts that will be running SQL. When hosts are replaced, policies are re-created on the new hosts automatically.
 - b. **Add VMs** – Select the VMs that should run on the selected hosts.
 - c. **Review and Create** the policy.

Select hosts

[+ Edit hosts](#) [Unassign](#)

<input type="checkbox"/>	Name ↑	Associated policies	Virtual machines
<input type="checkbox"/>	esx13-r16.p01.eastus.avslab.azure.com	0	0

Select virtual machines

[+ Edit virtual machines](#) [Unassign](#)

<input type="checkbox"/>	Display name ↑
<input type="checkbox"/>	vsan-healthcheck-disposable-04-15-1502-44-esx22-r07.p01.7c8287cb0d474c449d58df.e

[Next: Review and Create](#)

Manage placement policies

After creating the placement policy, you can review, manage, or edit the policy by way of the Placement policies menu in the Azure VMware Solution private cloud.

By checking the Azure hybrid benefit checkbox in the configuration setting, you can enable existing host-VM affinity policies with the SQL Azure hybrid benefit.

Microsoft Azure (Preview) Report a bug Search resources, services, and docs (G+)

Home > NormalSDDC2

NormalSDDC2 | Placement policies

AVS Private cloud

Search (Ctrl+F)

Settings

- Locks
- Manage
 - Connectivity
 - Clusters
 - Encryption
 - VMware credentials
 - Identity
 - Storage (preview)
 - Placement policies**
 - + Add-ons

+ Create policy Refresh Restrict VM movement Feedback

Total policies VMs with restricted movement

Search policies Cluster: Cluster-3 Enable: All Policy type: All Azure hybrid benefit: All

Name ↑	Type	Azure hybrid benefit	Hosts	VMs	Provisioning state	State
sql_ahb_noe_to_sql_test	VM-Host affinity	-	1	2	Succeeded	Enabled
sqlahbenabled	VM-Host affinity	Yes	1	1	Succeeded	Enabled

Next steps

[Azure Hybrid Benefit](#)

[Attach Azure NetApp Files datastores to Azure VMware Solution hosts \(Preview\)](#)

Enable VMware Cloud Director service with Azure VMware Solution (Preview)

12/16/2022 • 8 minutes to read • [Edit Online](#)

VMware Cloud Director service (CDs) with Azure VMware Solution enables enterprise customers to use APIs or the Cloud Director services portal to self-service provision and manage virtual datacenters through multi-tenancy with reduced time and complexity.

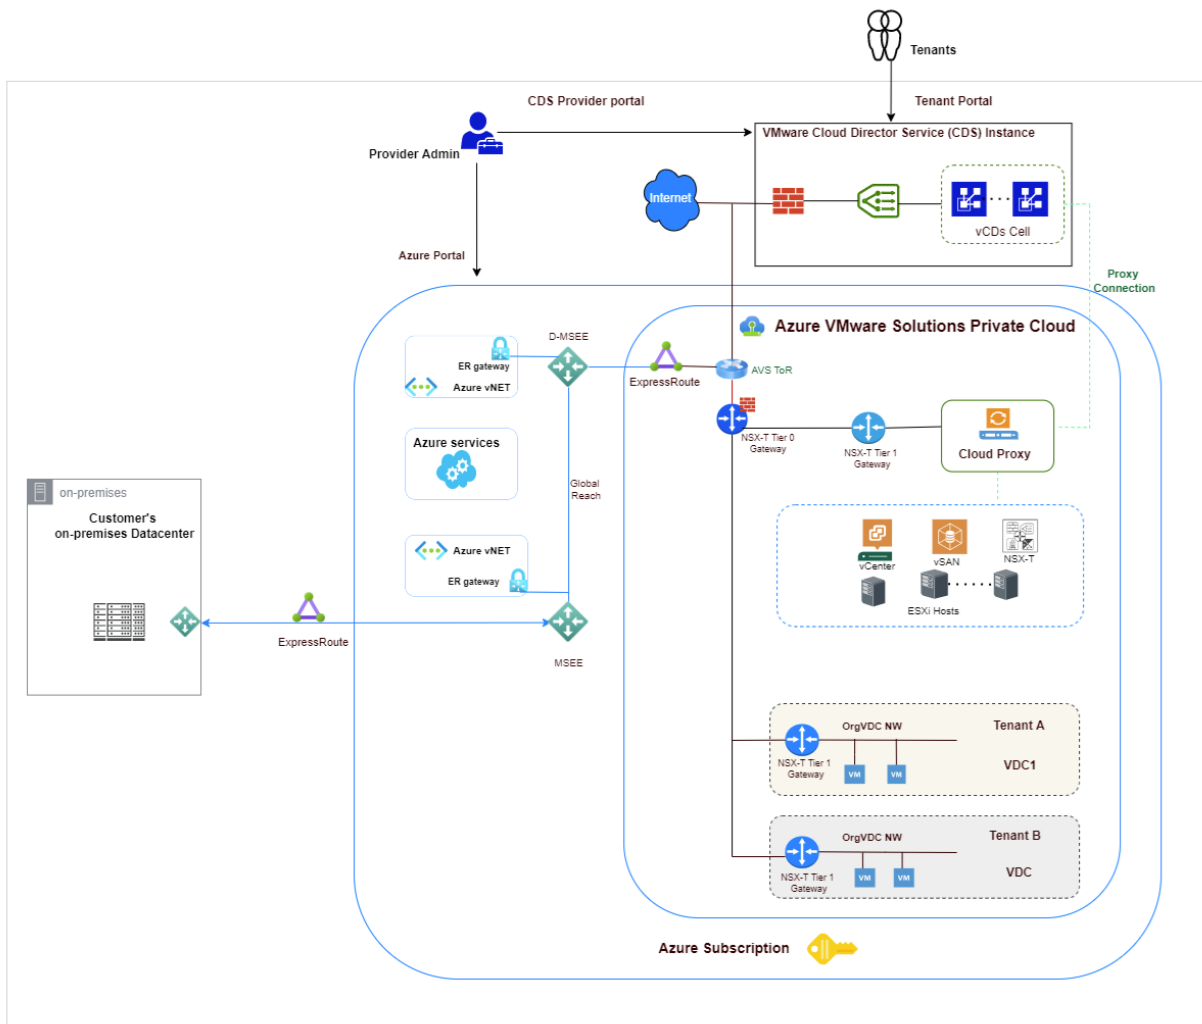
In this article, you'll learn how to enable VMware Cloud Director service with Azure VMware Solution for enterprise customers to use Azure VMware Solution resources and Azure VMware Solution private clouds with underlying resources for virtual datacenters.

IMPORTANT

VMware Cloud Director service is now available to use with Azure VMware Solution under the Enterprise Agreement (EA) model only. It's not suitable for MSP / Hosters to resell Azure VMware Solution capacity to customers at this point. For more information, see [Azure Service terms](#).

Reference architecture

The following diagram shows typical architecture for Cloud Director services with Azure VMware Solution and how they're connected. Communications to Azure VMware Solution endpoints from Cloud Director service are supported by an SSL reverse proxy.

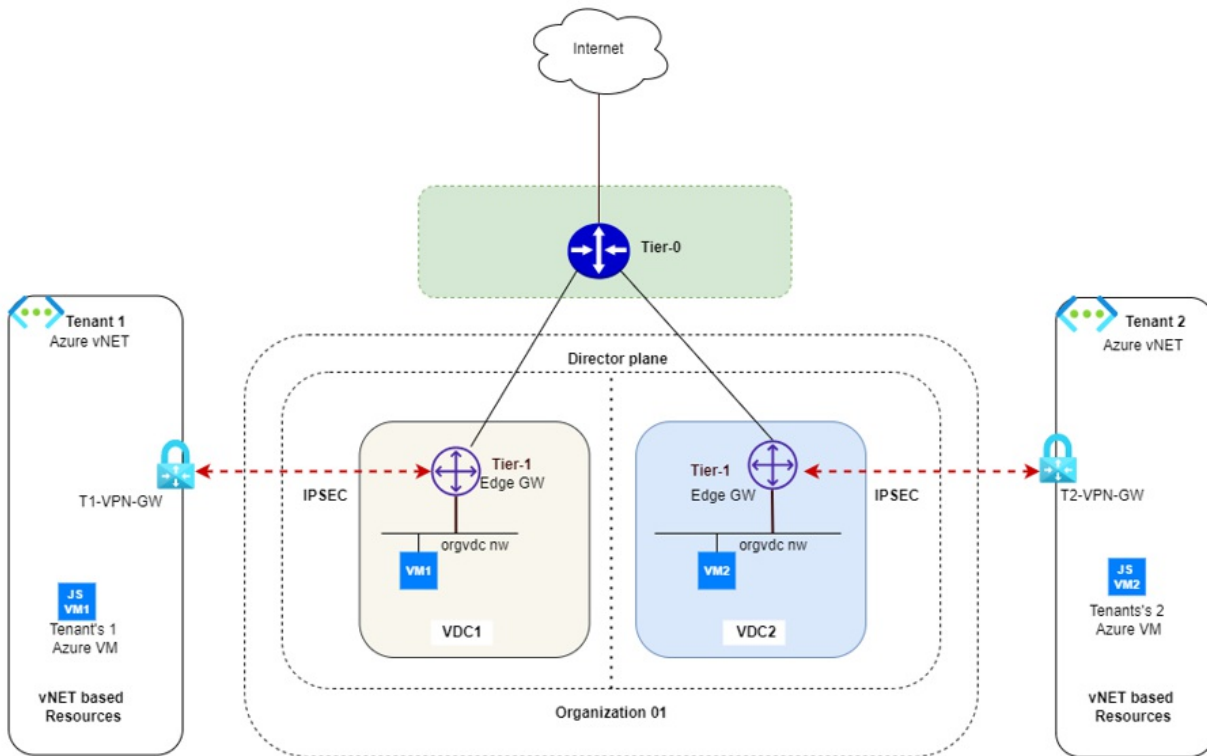


VMware Cloud Director supports multi-tenancy by using organizations. A single organization can have multiple organization virtual data centers (VDC). Each Organization's VDC can have their own dedicated Tier-1 router (Edge Gateway) which is further connected with the provider's managed shared Tier-0 router.

[Learn more about CDs on Azure VMware Solutions reference architecture](#)

Connect tenants and their organization virtual datacenters to Azure vNet based resources

To provide access to vNET based Azure resources, each tenant can have their own dedicated Azure vNET with Azure VPN gateway. A site-to-site VPN between customer organization VDC and Azure vNET is established. To achieve this connectivity, the provider will provide public IP to the organization VDC. Organization VDC's administrator can configure IPSEC VPN connectivity from the Cloud Director service portal.



As shown in the diagram above, organization 01 has two organization virtual datacenters: VDC1 and VDC2. The virtual datacenter of each organization has its own Azure vNETs connected with their respective organization VDC Edge gateway through IPSEC VPN. Providers provide public IP addresses to the organization VDC Edge gateway for IPSEC VPN configuration. An ORG VDC Edge gateway firewall blocks all traffic by default, specific allow rules needs to be added on organization Edge gateway firewall.

Organization VDCs can be part of a single organization and still provide isolation between them. For example, VM1 hosted in organization VDC1 cannot ping Azure VM JSVM2 for tenant2.

Prerequisites

- Organization VDC is configured with an Edge gateway and has Public IPs assigned to it to establish IPSEC VPN by provider.
- Tenants have created a routed Organization VDC network in tenant's virtual datacenter.
- Test VM1 and VM2 are created in the Organization VDC1 and VDC2 respectively. Both VMs are connected to the routed orgVDC network in their respective VDCs.
- Have a dedicated [Azure vNET](#) configured for each tenant. For this example, we created Tenant1-vNet and Tenant2-vNet for tenant1 and tenant2 respectively.
- Create an [Azure Virtual network gateway](#) for vNETs created earlier.
- Deploy Azure VMs JSVM1 and JSVM2 for tenant1 and tenant2 for test purposes.

NOTE

VMware Cloud Director service supports a policy-based VPN. Azure VPN gateway configures route-based VPN by default and to configure policy-based VPN policy-based selector needs to be enabled.

Configure Azure vNet

Create the following components in tenant's dedicated Azure vNet to establish IPSEC tunnel connection with the tenant's ORG VDC edge gateway.

- Azure Virtual network gateway

- Local network gateway.
- Add IPSEC connection on VPN gateway.
- Edit connection configuration to enable policy-based VPN.

Create Azure virtual network gateway

To create an Azure virtual network gateway, see the [create-a-virtual-network-gateway tutorial](#).

Create local network gateway

1. Log in to the Azure portal and select **Local network gateway** from marketplace and then select **Create**.
2. Local Network Gateway represents remote end site details. Therefore provide tenant1 OrgVDC public IP address and orgVDC Network details to create local end point for tenant1.
3. Under **Instance details**, select **Endpoint** as IP address
4. Add IP address (add Public IP address from tenant's OrgVDC Edge gateway).
5. Under **Address space** add **Tenants Org VDC Network**.
6. Repeat steps 1-5 to create a local network gateway for tenant 2.

Create IPSEC connection on VPN gateway

1. Select tenant1 VPN Gateway (created earlier) and then select **Connection** (in left pane) to add new IPSEC connection with tenant1 orgVDC Edge gateway.
2. Enter the following details.

NAME	CONNECTION
Connection Type	Site to Site
VPN Gateway	Tenant's VPN Gateway
Local Network Gateway	Tenant's Local Gateway
PSK	Shared Key (provide a password)
IKE Protocol	IKEV2 (ORG-VDC is using IKEv2)

3. Select **Ok** to deploy local network gateway.

Configure IPsec Connection

VMware Cloud Director service supports a policy-based VPN. Azure VPN gateway configures route-based VPN by default and to configure policy-based VPN policy-based selector needs to be enabled.

1. Select the connection you created earlier and then select **configuration** to view the default settings.
2. **IPSEC/IKE Policy**
3. **Enable policy base traffic selector**
4. Modify all other parameters to match what you have in OrgVDC.

NOTE

Both source and destination of the tunnel should have identical settings for IKE,SA, DPD etc.

5. Select **Save**.

Configure VPN on organization VDC Edge router

1. Log in to Organization VMware Cloud Director service tenant portal and select tenant's edge gateway.

2. Select **IPSEC VPN** option under **Services** and then select **New**.
3. Under general setting, provide **Name** and select desired security profile. Ensure that security profile settings (IKE, Tunnel, and DPD configuration) are same on both sides of the IPsec tunnel.
4. Modify Azure VPN gateway to match the Security profile, if necessary. You can also do security profile customization from CDS tenant portal.

NOTE

VPN tunnel won't establish if these settings were mismatched.

5. Under **Peer Authentication Mode**, provide the same pre-shared key that is used at the Azure VPN gateway.
6. Under **Endpoint configuration**, add the Organization's public IP and network details in local endpoint and Azure VNet details in remote endpoint configuration.
7. Under **Ready to complete**, review applied configuration.
8. Select **Finish** to apply configuration.

Apply firewall configuration

Organization VDC Edge router firewall denies traffic by default. You'll need to apply specific rules to enable connectivity. Use the following steps to apply firewall rules.

1. Add IP set in VMware Cloud Director service portal
 - a. Log in to Edge router then select **IP SETS** under the **Security** tab in left plane.
 - b. Select **New** to create IP sets.
 - c. Enter **Name** and **IP address** of test VM deployed in orgVDC.
 - d. Create another IP set for Azure vNET for this tenant.
2. Apply firewall rules on ORG VDC Edge router.
 - a. Under **Edge gateway**, select **Edge gateway** and then select **firewall** under **services**.
 - b. Select **Edit rules**.
 - c. Select **NEW ON TOP** and enter rule name.
 - d. Add **source** and **destination** details. Use created IPSET in source and destination.
 - e. Under **Action**, select **Allow**.
 - f. Select **Save** to apply configuration.
3. Verify tunnel status
 - a. Under **Edge gateway** select **Service**, then select **IPSEC VPN**,
 - b. Select **View statistics**.
Status of tunnel should show **UP**.
4. Verify IPsec connection
 - a. Log in to Azure VM deployed in tenants vNET and ping tenant's test VM IP address in tenant's OrgVDC.
For example, ping VM1 from JSVM1. Similarly, you should be able to ping VM2 from JSVM2. You can verify isolation between tenants Azure vNETs. Tenant1's VM1 won't be able to ping Tenant2's Azure VM JSVM2 in tenant2 Azure vNETs.

Connect Tenant workload to public Internet

- Tenants can use public IP to do SNAT configuration to enable Internet access for VM hosted in organization VDC. To achieve this connectivity, the provider can provide public IP to the organization VDC.

- Each organization VDC can be created with dedicated T1 router (created by provider) with reserved Public & Private IP for NAT configuration. Tenants can use public IP SNAT configuration to enable Internet access for VM hosted in organization VDC.
- OrgVDC administrator can create a routed OrgVDC network connected to their OrgVDC Edge gateway. To provide Internet access.
- OrgVDC administrator can configure SNAT to provide a specific VM or use network CIDR to provide public connectivity.
- OrgVDC Edge has default DENY ALL firewall rule. Organization administrators will need to open appropriate ports to allow access through the firewall by adding a new firewall rule. Virtual machines configured on such OrgVDC network used in SNAT configuration should be able to access the Internet.

Prerequisites

1. Public IP is assigned to the organization VDC Edge router. To verify, log in to the organization's VDC. Under **Networking > Edges**, select **Edge Gateway**, then select **IP allocations** under **IP management**. You should see a range of assigned IP address there.
2. Create a routed Organization VDC network. (Connect OrgVDC network to the edge gateway with public IP address assigned)

Apply SNAT configuration

1. Log in to Organization VDC. Navigate to your Edge gateway and then select **NAT** under **Services**.
2. Select **New** to add new SNAT rule.
3. Provide **Name** and select **Interface type** as SNAT.
4. Under **External IP**, enter public IP address from public IP pool assigned to your orgVDC Edge router.
5. Under **Internal IP**, enter IP address for your test VM. This IP address is one of the orgVDC network IP assigned to the VM.
6. **State** should be enabled.
7. Under **Priority**, select a higher number. For example, 4096.
8. Select **Save** to save the configuration.

Apply firewall rule

1. Log in to Organization VDC and navigate to **Edge Gateway**, then select **IP set** under security.
2. Create an IPset. Provide IP address of your VM (you can use CIDR also). Select **Save**.
3. Under **services**, select **Firewall**, then select **Edit rules**.
4. Select **New ON TOP** and create a firewall rule to allow desired port and destination.
5. Select the **IPset** your created earlier as source. Under **Action**, select **Allow**.
6. Select **Keep** to save the configuration.
7. Log in to your test VM and ping your destination address to verify outbound connectivity.

Migrate workloads to VMware Cloud Director service on Azure VMware Solution

VMware Cloud Director Availability can be used to migrate VMware Cloud Director workload into the VMware Cloud Director service on Azure VMware Solution. Enterprise customers can drive self-serve one-way warm migration from the on-premises Cloud Director Availability vSphere plugin, or they can run the Cloud Director Availability plugin from the provider-managed Cloud Director instance and move workloads into Azure VMware Solution.

For more information about VMware Cloud Director Availability, see [VMware Cloud Director Availability | Disaster Recovery & Migration](#)

FAQs

Question: What are the supported Azure regions for the VMware Cloud Director service?

Answer: This offering is supported in all Azure regions where Azure VMware Solution is available except for Brazil South and South Africa. Ensure that the region you wish to connect to VMware Cloud Director service is within a 150-milliseconds round trip time for latency with VMware Cloud Director service.

Question: How do I configure VMware Cloud Director service on Microsoft Azure VMware Solutions?

Answer [Learn about how to configure CDs on Azure VMware Solutions](#)

Next steps

[VMware Cloud Director Service Documentation](#)

[Migration to Azure VMware Solutions with Cloud Director service](#)

Upgrade HCX on Azure VMware Solution

12/16/2022 • 3 minutes to read • [Edit Online](#)

In this article, you'll learn how to upgrade Azure VMware Solution for HCX service updates that may include new features, software fixes, or security patches.

You can update HCX Connector and HCX Cloud systems during separate maintenance windows, but for optimal compatibility, it's recommended you update both systems together. Apply service updates during a maintenance window where no new HCX operations are queued up.

IMPORTANT

Starting with HCX 4.4.0, HCX appliances install the VMware Photon Operating System. When upgrading to HCX 4.4.x or later from an HCX version prior to version 4.4.0, you must also upgrade all Service Mesh appliances.

System requirements

- For systems requirements, compatibility, and upgrade prerequisites, see the [VMware HCX release notes](#).
- For more information about the upgrade path, see the [Product Interoperability Matrix](#).
- Ensure HCX manager and site pair configurations are healthy.
- As part of HCX update planning, and to ensure that HCX components are updated successfully, review the service update considerations and requirements. For planning HCX upgrade, see [Planning for HCX Updates](#).
- Ensure that you have a backup and snapshot of HCX connector in the on-premises environment, if applicable.

Backup HCX

- Azure VMware Solution backs up HCX Cloud Manager configuration daily.
- Use the appliance management interface to create backup of HCX in on-premises, see [Backing Up HCX Manager](#). You can use the configuration backup to restore the appliance to its state before the backup. The contents of the backup file supersede configuration changes made before restoring the appliance.
- HCX cloud manager snapshots are taken automatically during upgrades to HCX 4.4 or later. HCX retains automatic snapshots for 24 hours before deleting them. To take a manual snapshot on HCX Cloud Manager or help with reverting from a snapshot, [create a support ticket](#).

Upgrade HCX

The upgrade process is in two steps:

1. Upgrade HCX Manager
 - a. HCX cloud manager
 - b. HCX connector (You can update site-paired HCX Managers simultaneously)
2. Upgrade HCX Service Mesh appliances

Upgrade HCX manager

The HCX update is first applied to the HCX Manager systems.

What to expect

- HCX manager is rebooted as part of the upgrade process.
- HCX vCenter Plugins will be updated.
- There's no data-plane outage during this procedure.

Prerequisites

- Verify the HCX Manager system reports healthy connections to the connected (vCenter Server, NSX Manager (if applicable)).
- Verify the HCX Manager system reports healthy connections to the HCX Interconnect service components. (Ensure HCX isn't in an out of sync state)
- Verify that Site Pair configurations are healthy.
- No VM migrations should be in progress during this upgrade.

Procedure

To follow the HCX Manager upgrade process, see [Upgrading the HCX Manager](#)

Upgrade HCX Service Mesh appliances

While Service Mesh appliances are upgraded independently to the HCX Manager, they must be upgraded. These appliances are flagged for new available updates anytime the HCX Manager has newer software available.

What to expect

- Service VMs will be rebooted as part of the upgrade.
- There is a small data plane outage during this procedure.
- In-service upgrade of Network-extension can be considered to reduce downtime during HCX Network extension upgrades.

Prerequisites

- All paired HCX Managers on both the source and the target site are updated and all services have returned to a fully converged state.
- Service Mesh appliances must be initiated using the HCX plug-in of vCenter or the 443 console at the source site
- No VM migrations should be in progress during this upgrade.

Procedure

To follow the Service Mesh appliances upgrade process, see [Upgrading the HCX Service Mesh Appliances](#)

FAQ

What is the impact of an HCX upgrade?

Apply service updates during a maintenance window where no new HCX operations and migration are queued up. The upgrade window accounts for a brief disruption to the Network Extension service, while the appliances are redeployed with the updated code.

For individual HCX component upgrade impact, see [Planning for HCX Updates](#).

Do I need to upgrade the service mesh appliances?

The HCX Service Mesh can be upgraded once all paired HCX Manager systems are updated, and all services have returned to a fully converged state. Check HCX release notes for upgrade requirements. Starting with HCX 4.4.0, HCX appliances installed the VMware Photon Operating System. When upgrading to HCX 4.4.x or later from an HCX version prior to 4.4.0 version, you must upgrade all Service Mesh appliances.

How do I roll back HCX upgrade using a snapshot?

See [Rolling Back an Upgrade Using Snapshots](#). On the cloud side, open a [support ticket](#) to roll back the upgrade.

Next steps

[Software Versioning, Skew and Legacy Support Policies](#)

[Updating VMware HCX](#)

Open a support request for an Azure VMware Solution deployment or provisioning failure

12/16/2022 • 3 minutes to read • [Edit Online](#)

This article shows you how to open a [support request](#) and provide key information for an Azure VMware Solution deployment or provisioning failure.

When you have a failure on your private cloud, you need to open a support request in the Azure portal. To open a support request, first get some key information in the Azure portal:

- Correlation ID
- Error messages
- Azure ExpressRoute circuit ID

Get the correlation ID

When you create a private cloud or any resource in Azure, a correlation ID for the resource is automatically generated for the resource. Include the private cloud correlation ID in your support request to more quickly open and resolve the request.

In the Azure portal, you can get the correlation ID for a resource in two ways:

- **Overview** pane
- Deployment logs

Get the correlation ID from the resource overview

Here's an example of the operation details of a failed private cloud deployment, with the correlation ID selected:

The screenshot shows the Azure portal interface for a failed deployment. At the top, there are action buttons: Delete, Cancel, Redeploy, and Refresh. Below these is a red error message: "The resource operation completed with terminal provisioning state 'Failed'. Click here for details →". The main heading is "Your deployment failed". Below this, there are details for the deployment: Deployment name: VMCP-20200528091210, Subscription: [redacted], Resource group: contoso-a01. The start time is 5/28/2020, 9:12:16 AM. The correlation ID is cc2ffcdd-ca98-2020-91dc-1e3c20362020, which is highlighted with a red box. Below the details is a link for "Deployment details (Download)". At the bottom, there is a table with columns: Resource, Type, Status, and Operation details. The table contains one row: Resource: pc03, Type: Microsoft.AVS/privateClouds, Status: Conflict, Operation details: [Operation details](#).

Resource	Type	Status	Operation details
pc03	Microsoft.AVS/privateClouds	Conflict	Operation details

To access deployment results in a private cloud **Overview** pane:

1. In the Azure portal, select your private cloud.
2. In the left menu, select **Overview**.

After a deployment is initiated, the results of the deployment are shown in the private cloud **Overview** pane.

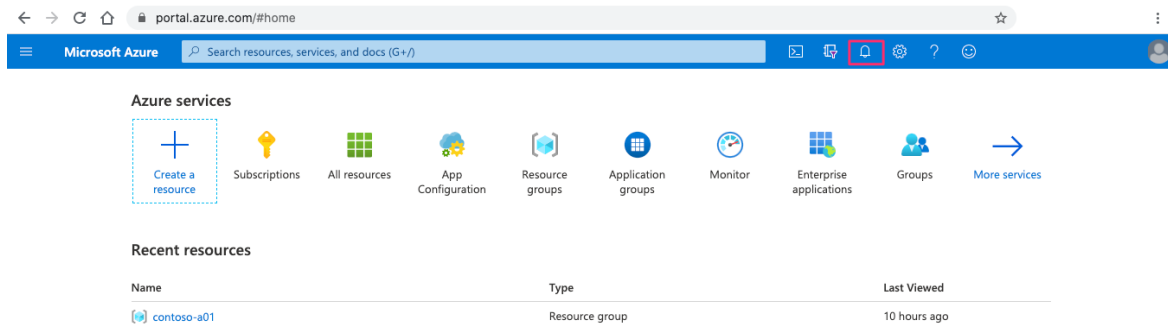
Copy and save the private cloud deployment correlation ID to include in the service request.

Get the correlation ID from the deployment log

You can get the correlation ID for a failed deployment by searching the deployment activity log located in the Azure portal.

To access the deployment log:

1. In the Azure portal, select your private cloud, and then select the notifications icon.

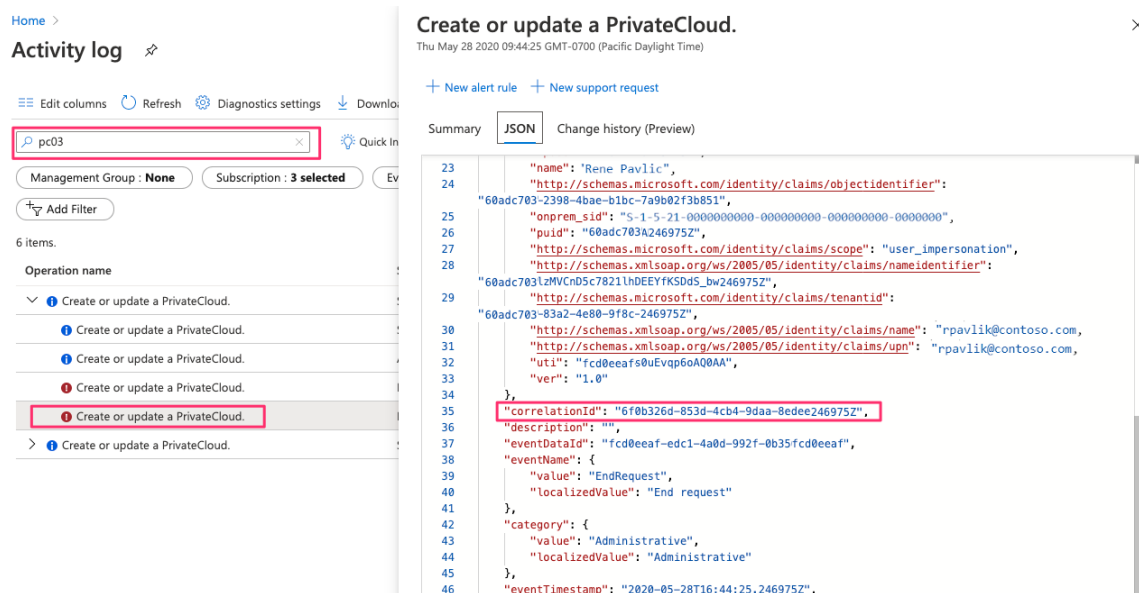


2. In the Notifications pane, select **More events in the activity log**:



3. To find the failed deployment and its correlation ID, search for the name of the resource or other information that you used to create the resource.

The following example shows search results for a private cloud resource named pc03.



4. In the search results in the **Activity log** pane, select the operation name of the failed deployment.
5. In the **Create or update a PrivateCloud** pane, select the **JSON** tab, and then look for `correlationId` in the log that is shown. Copy the `correlationId` value to include it in your support request.

Copy error messages

To help resolve your deployment issue, include any error messages that are shown in the Azure portal. Select a warning message to see a summary of errors:

Errors



Summary

Raw Error

ERROR DETAILS



- ✓ The resource operation completed with terminal provisioning state 'Failed'. (Code: ResourceDeploymentFailure)
 - The -mgmt-d deployment in the 02-amst01 resource group has failed. The status is: Failed at 05/28/2020 16:17:45. Details: Code: 'AnotherOperationInProgress' Message: 'Another operation on this or dependent resource is in progress. To retrieve status of the operation use uri: https://management.azure.com/subscriptions/providers/Microsoft.Network/locations/westeurope/operations/providers/Microsoft.Network/?api-version=2019-02-01.' Code: 'AnotherOperationInProgress' Message: 'The access token for this request was issued by the tenant.'

WAS THIS HELPFUL?

Troubleshooting Options

- [Common Azure deployment errors](#)
- [Check Usage + Quota](#)
- [New Support Request](#)

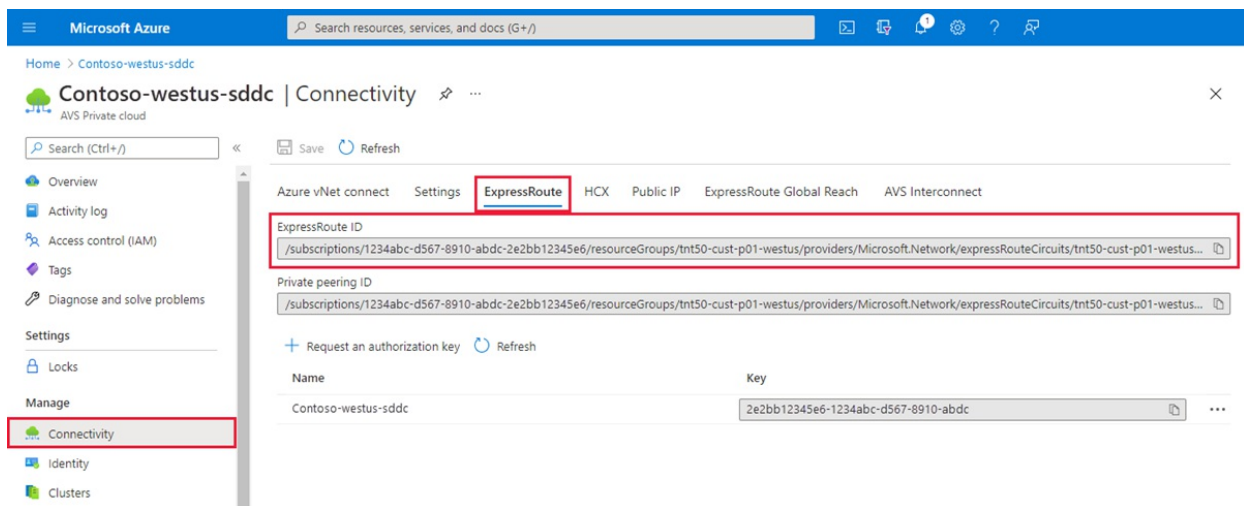
To copy the error message, select the copy icon. Save the copied message to include in your support request.

Get the ExpressRoute ID (URI)

Perhaps you're trying to scale or peer an existing private cloud with the private cloud ExpressRoute circuit, and it fails. In that scenario, you need the ExpressRoute ID to include in your support request.

To copy the ExpressRoute ID:

1. In the Azure portal, select your private cloud.
2. In the left menu, under **Manage**, select **Connectivity**.
3. In the right pane, select the **ExpressRoute** tab.
4. Select the copy icon for **ExpressRoute ID** and save the value to use in your support request.



Pre-validation failures

If your private cloud pre-validations check failed (before deployment), a correlation ID won't have been

generated. In this scenario, you can provide the following information in your support request:

- Error and failure messages. These messages can be helpful in many failures, for example, for quota-related issues. It's important to copy these messages and include them in the support request, as described in this article.
- Information you used to create the Azure VMware Solution private cloud, including:
 - Location
 - Resource group
 - Resource name

Create your support request

For general information about creating a support request, see [How to create an Azure support request](#).

To create a support request for an Azure VMware Solution deployment or provisioning failure:

1. In the Azure portal, select the **Help** icon, and then select **New support request**.

The screenshot shows the Microsoft Azure portal interface. At the top, the navigation bar includes the Microsoft Azure (Preview) logo, a search bar, and several utility icons. A red box highlights the question mark icon, which represents the Help section. Below the navigation bar, the page title is 'Help + support | New support request'. A search bar is present, followed by tabs for 'Basics', 'Solutions', 'Details', and 'Review + create'. The 'Basics' tab is active. On the left, there is a sidebar with 'Support' options, including 'New support request' (highlighted), 'All support requests', 'Support Plans', 'Service Health', and 'Advisor'. The main content area contains instructions and a form. The instructions state: 'Create a new support request to get assistance with billing, subscription, technical (including advisory) or quota management issues. Complete the Basics tab by selecting the options that best describe your problem. Providing detailed, accurate information can help to solve your issues faster.' The form fields are: Issue type (Technical), Subscription (9e93-d505c-4e34-d505c-4e34-426c-9e93-7e5e8ff744...), Service (My services selected), Summary (My AVS private cloud failed to deploy), Problem type (Configuration and Setup Issues), and Problem subtype (Provision a Private Cloud). A 'Next: Solutions >>' button is at the bottom.

2. Enter or select the required information:

- a. On the **Basics** tab:

- a. For **Problem type**, select **Configuration and Setup Issues**.

- b. For **Problem subtype**, select **Provision a private cloud**.

- b. On the **Details** tab:

- a. Enter or select the required information.

- b. Paste your Correlation ID or ExpressRoute ID where this information is requested. If you don't see a specific text box for these values, paste them in the **Provide details about the issue** text box.

- c. Paste any error details, including the error or failure messages you copied, in the **Provide details about the issue** text box.

3. Review your entries, and then select **Create** to create your support request.

