



NetScaler VPX 13.0

Contents

Support matrix and usage guidelines	5
Optimize Citrix ADC VPX performance on VMware ESX, Linux KVM, and Citrix Hypervisors	12
Apply Citrix ADC VPX configurations at the first boot of the Citrix ADC appliance in cloud	25
Install a Citrix ADC VPX instance on a bare metal server	60
Install a Citrix ADC VPX instance on Citrix Hypervisor/XenServer	60
Configure VPX instances to use single root I/O virtualization (SR-IOV) network interfaces	64
Install a Citrix ADC VPX instance on VMware ESX	69
Configure a Citrix ADC VPX instance to use VMXNET3 network interface	74
Configure a Citrix ADC VPX instance to use SR-IOV network interface	86
Migrating the Citrix ADC VPX from E1000 to SR-IOV or VMXNET3 Network Interfaces	104
Configure a Citrix ADC VPX instance to use PCI passthrough network interface	105
Apply Citrix ADC VPX configurations at the first boot of the Citrix ADC appliance on VMware ESX hypervisor	108
Install a Citrix ADC VPX instance on VMware cloud on AWS	114
Install a Citrix ADC VPX instance on Microsoft Hyper-V server	117
Install a Citrix ADC VPX instance on Linux-KVM platform	122
Prerequisites for installing a Citrix ADC VPX instance on Linux-KVM platform	123
Provision the Citrix ADC VPX instance by using OpenStack	128
Provision the Citrix ADC VPX instance by using the Virtual Machine Manager	137
Configure a Citrix ADC VPX instance to use SR-IOV network interfaces	151
Configure a Citrix ADC VPX instance to use PCI passthrough network interfaces	162
Provision the Citrix ADC VPX instance by using the virsh program	166
Manage the Citrix ADC VPX guest VMs	170

Provision the Citrix ADC VPX instance with SR-IOV, on OpenStack	173
Configure a Citrix ADC VPX instance on KVM to use OVS DPDK-based host interfaces	179
Citrix ADC VPX on AWS	189
AWS terminology	192
VPX-AWS support matrix	194
Limitations and usage guidelines	197
Prerequisites	199
How a Citrix ADC VPX instance on AWS works	201
Deploy a Citrix ADC VPX standalone instance on AWS	202
Scenario: standalone instance	208
Download a Citrix ADC VPX license	217
Load balancing servers in different availability zones	222
How high availability on AWS works	223
Deploy a VPX HA pair in the same AWS availability zone	226
High availability across different AWS availability zones	237
Deploy a VPX high-availability pair with elastic IP addresses across different AWS zones	238
Deploy a VPX high-availability pair with private IP addresses across different AWS zones	243
Deploy a Citrix ADC VPX instance on AWS Outposts	251
Add back-end AWS Autoscaling service	253
Configure a Citrix ADC VPX instance to use SR-IOV network interface	260
Configure a Citrix ADC VPX instance to use Enhanced Networking with AWS ENA	263
Upgrade a Citrix ADC VPX instance on AWS	263
Troubleshoot a VPX instance on AWS	268
AWS FAQs	269

Deploy a Citrix ADC VPX instance on Microsoft Azure	272
Azure terminology	277
Network architecture for Citrix ADC VPX instances on Microsoft Azure	281
Configure a Citrix ADC VPX standalone instance	284
Configure multiple IP addresses for a Citrix ADC VPX standalone instance	298
Configure a high-availability setup with multiple IP addresses and NICs	304
Configure a high-availability setup with multiple IP addresses and NICs by using PowerShell commands	314
Configure a Citrix ADC VPX instance to use Azure accelerated networking	326
Configure HA-INC nodes by using the Citrix high availability template with Azure ILB	341
Configure HA-INC nodes by using the Citrix high availability template for internet-facing applications	354
Configure a high-availability setup with Azure external and internal load balancers simultaneously	365
Install a Citrix ADC VPX instance on Azure VMware Solution	369
Add Azure Autoscale settings	386
Azure tags for Citrix ADC VPX deployment	392
Configure GSLB on Citrix ADC VPX instances	398
Configure GSLB on an active-standby high-availability setup	406
Configure address pools intranet IP for a Citrix Gateway appliance	410
Configure multiple IP addresses for a Citrix ADC VPX standalone instance by using PowerShell commands	412
Additional PowerShell scripts for Azure deployment	420
Azure FAQs	436
Deploy a Citrix ADC VPX instance on the Google Cloud Platform	436

Deploy a VPX high-availability pair on Google Cloud Platform	453
Deploy a VPX high-availability pair with external static IP address on the Google Cloud Platform	454
Deploy a VPX high-availability pair with private IP address on Google Cloud Platform	464
Add back-end GCP Autoscaling service	473
VIP scaling support for Citrix ADC VPX instance on GCP	478
Troubleshoot a VPX instance on GCP	484
Jumbo frames on Citrix ADC VPX instances	484
Automate deployment and configurations of Citrix ADC	486
FAQs	489

Support matrix and usage guidelines

September 13, 2024

This document lists the different hypervisors, and features supported on a Citrix ADC VPX instance. The document also describes their usage guidelines, and known limitations.

Table 1. VPX instance on Citrix Hypervisor

Citrix Hypervisor version	SysID	VPX models
8.2 supported 13.0 64.x onwards, 8.0, 7.6, 7.1	450000	VPX 10, VPX 25, VPX 200, VPX 1000, VPX 3000, VPX 5000, VPX 8000, VPX 10G, VPX 15G, VPX 25G, VPX 40G

Table 2. VPX instance on VMware ESXi hypervisor

The following VPX models with 450010 (Sys ID) supports the VMware ESX versions listed in the table.

VPX models: VPX 10, VPX 25, VPX 200, VPX 1000, VPX 3000, VPX 5000, VPX 8000, VPX 10G, VPX 15G, VPX 25G, VPX 40G, and VPX 100G.

ESXi version	ESXi release date in MM/DD/YYYY format	ESXi build number	Citrix ADC VPX version
ESXi 8.0 update 2	09/21/2023	22380479	13.0-92.x and higher builds
ESXi 8.0 update 1	04/18/2023	21495797	13.0-90.x and higher builds
ESXi 8.0c	03/30/2023	21493926	13.0-90.x and higher builds
ESXi 8.0	10/11/2022	20513097	13.0-90.x and higher builds
ESXi 7.0 update 3n	07/06/2023	21930508	13.0-91.x and higher builds
ESXi 7.0 update 3m	05/03/2023	21686933	13.0-91.x and higher builds
ESXi 7.0 update 3i	12/08/2022	20842708	13.0-90.x and higher builds
ESXi 7.0 update 3f	07/12/2022	20036589	13.0-86.x and higher builds

ESXi version	ESXi release date in MM/DD/YYYY format	ESXi build number	Citrix ADC VPX version
ESXi 7.0 update 3d	03/29/2022	19482537	13.0-86.x and higher builds
ESXi 7.0 update 3c	01/27/2022	19193900	13.0-85.x and higher builds
ESXi 7.0 update 2d	09/14/2021	18538813	13.0-83.x and higher builds
ESXi 7.0 update 2a	12/17/2020	17867351	13.0-82.x and higher builds
ESXi 6.7 P04	11/19/2020	17167734	13.0-67.x and higher builds
ESXi 6.7 P03	08/20/2020	16713306	13.0-67.x and higher builds
ESXi 6.7 P02	04/28/2020	16075168	13.0-67.x and higher builds
ESXi 6.7 update 3	08/20/2019	14320388	13.0-58.x and higher builds
ESXi 6.5 U1g	3/20/2018	7967591	13.0 47.x and higher builds

Note:

Each ESXi patch support is validated on the Citrix ADC VPX version specified in the preceding table and is applicable for all the higher builds of Citrix ADC 13.0 version.

Table 3. VPX on Microsoft Hyper-V

Hyper-V version	SysID	VPX models
2012, 2012R2, 2016, 2019	450020	VPX 10, VPX 25, VPX 200, VPX 1000, VPX 3000

VPX instance on Nutanix AHV

NetScaler VPX is supported on Nutanix AHV through the [Citrix Ready partnership](#). Citrix Ready is a technology partner program that helps software and hardware vendors develop and integrate their products with NetScaler technology for digital workspace, networking, and analytics.

For more information on a step-by-step method to deploy a NetScaler VPX instance on Nutanix AHV, see [Deploying a NetScaler VPX on Nutanix AHV](#).

Third-party support:

If you experience any issues with a particular third-party (Nutanix AHV) integration on a NetScaler environment, open a support incident directly with the third-party partner (Nutanix).

If the partner determines that the issue appears to be with NetScaler, the partner can approach NetScaler support for further assistance. A dedicated technical resource from partners works with the NetScaler support until the issue is resolved.

Table 4. VPX instance on generic KVM

Generic KVM version	SysID	VPX models
RHEL 7.6, RHEL 8.0, RHEL 9.3 Ubuntu 16.04, Ubuntu 18.04, RHV 4.2	450070	VPX 10, VPX 25, VPX 200, VPX 1000, VPX 3000, VPX 5000, VPX 8000, VPX 10G, VPX 15G, VPX 25G, VPX 40G, VPX 100G

Points to note:

Consider the following points while using KVM hypervisors.

- The VPX instance is qualified for hypervisor release versions mentioned in table 1–4, and not for patch releases within a version. However, the VPX instance is expected to work seamlessly with patch releases of a supported version. If it does not, log a support case for troubleshooting and debugging.
- Use the `ip link` commands to configure RHEL 8.2 network bridges.
- Before using RHEL 7.6, complete the following steps on the KVM host:
 1. Edit `/etc/default/grub` and append `"kvm_intel.preemption_timer=0"` to `GRUB_CMDLINE_LINUX` variable.
 2. Regenerate `grub.cfg` with the command `"# grub2-mkconfig -o /boot/grub2/grub.cfg"`.
 3. Restart the host machine.
- Before using Ubuntu 18.04, complete the following steps on the KVM host:
 1. Edit `/etc/default/grub` and append `"kvm_intel.preemption_timer=0"` to `GRUB_CMDLINE_LINUX` variable.
 2. Regenerate `grub.cfg` with the command `"# grub-mkconfig -o /boot/grub/grub.cfg"`.
 3. Restart the host machine.

Table 5. VPX instance on AWS

AWS version	SysID	VPX models
N/A	450040	VPX 10, VPX 200, VPX 1000, VPX 3000, VPX 5000, VPX BYOL, VPX 8000, VPX 10G, VPX 15G, and VPX 25G are available only with BYOL with EC2 instance types (C5, M5, and C5n)

Note:

The VPX 25G offering doesn't give the desired 25G throughput in AWS but can give higher SSL transactions rate compared to VPX 15G offering.

Table 6. VPX instance on Azure

Azure version	SysID	VPX models
N/A	450020	VPX 10, VPX 200, VPX 1000, VPX 3000, VPX 5000, VPX BYOL

Table 7. VPX feature matrix

Features	VPX on XenServer		VPX on VMware ESX				VPX on Microsoft Hyper-V	VPX on generic KVM			VPX on AWS	VPX on Azure	VPX on GCP
	PV	SR-IOV	PV	SR-IOV	Emulated	PCI Passthrough	PV	PV	SR-IOV	PCI Passthrough			
Multi-PE Support	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustering Support	Yes	Yes ¹	Yes	Yes ¹	Yes	Yes	Yes	Yes	Yes ¹	Yes	No	No	No
VLAN Tagging	Yes	Yes	Yes	Yes	Yes	Yes	Yes (only on 2012R2)	Yes	Yes	Yes	No	No	No
Detecting Link Events	No ²	Yes ³	No ²	Yes ³	No ²	Yes ³	No ²	No ²	Yes ³	Yes ³	No ²	No ²	No ²
Interface Parameter Configuration	No	No	No	No	No	Yes	No	No	No	Yes	No	No	No
Static LA	Yes ²	Yes ³	Yes ²	No	Yes ²	Yes ³	Yes ²	Yes ²	Yes ³	Yes ²	No	No	No
LACP	No	Yes ³	Yes ²	No	Yes ²	Yes ³	No	Yes ²	Yes ³	Yes ³	No	No	No
Static CLAG	No	No	No	No	No	No	No	No	No	No	No	No	No
LACP CLAG	No	No	Yes ²	No	Yes ²	Yes ³	No	Yes ²	Yes ³	Yes ³	No	No	No
Hot-plug	No	No	No	No	No	No	No	No	No	No	Yes	No	No

The superscript numbers (1, 2, 3) used in the preceding table refers to the following points with respective numbering:

1. Clustering support is available on SRIOV for client-facing and server-facing interfaces and not for the backplane.
2. Interface DOWN events are not recorded in Citrix ADC VPX instances.
3. For Static LA, traffic might still be sent on the interface whose physical status is DOWN.
4. For LACP, the peer device knows the interface DOWN event based on the LACP timeout mechanism.
 - Short timeout: 3 seconds
 - Long timeout: 90 seconds
5. For LACP, do not share interfaces across VMs.
6. For Dynamic routing, convergence time depends on the Routing Protocol since link events are not detected.
7. Monitored static Route functionality fails if you do not bind monitors to static routes because the Route state depends on the VLAN status. The VLAN status depends on the link status.
8. Partial failure detection does not happen in high availability if there's link failure. High availability-split brain condition might happen if there's link failure.
 - When any link event (disable/enable, reset) is generated from a VPX instance, the physical status of the link does not change. For static LA, any traffic initiated by the peer gets dropped on the instance.
 - For the VLAN tagging feature to work, do the following:

On the VMware ESX, set the port group's VLAN ID to 1–4095 on the vSwitch of the VMware ESX server.

Table 8. Supported browsers

Operating system	Browser and versions
Windows 7	Internet Explorer- 8, 9, 10, and 11; Mozilla Firefox 3.6.25 and above; Google Chrome- 15 and above
Windows 64 bit	Internet Explorer - 8, 9; Google Chrome - 15 and above
MAC	Mozilla Firefox - 12 and above; Safari - 5.1.3; Google Chrome - 15 and above

Usage guidelines

Follow these usage guidelines:

- We recommend you to deploy a VPX instance on local disks of the server or SAN-based storage volumes.

See the **VMware ESXi CPU Considerations** section in the document [Performance Best Practices for VMware vSphere 6.5](#). Here's an extract:

- It is not recommended that virtual machines with high CPU/Memory demand sit on a Host/Cluster that is overcommitted.
- In most environments, ESXi allows significant levels of CPU overcommitment without impacting virtual machine performance. On a host, you can run more vCPUs than the total number of physical processor cores in that host.
- If an ESXi host becomes CPU saturated, that is, the virtual machines and other loads on the host demand all the CPU resources the host has, latency-sensitive workloads might not perform well. In this case you might want to reduce the CPU load, for example, by powering off some virtual machines or migrating them to a different host, or allowing DRS to migrate them automatically.
- Citrix recommends the latest hardware compatibility version to avail the latest feature sets of the ESXi hypervisor for the virtual machine. For more information about the hardware and ESXi version compatibility, see [VMware documentation](#).
- The Citrix ADC VPX is a latency-sensitive, high-performance virtual appliance. To deliver its expected performance, the appliance requires vCPU reservation, memory reservation, vCPU pinning on the host. Also, hyper threading must be disabled on the host. If the host does not meet these requirements, issues such as high-availability failover, CPU spike within the VPX instance, sluggishness in accessing the VPX CLI, pit boss daemon crash, packet drops, and low throughput occur.

A hypervisor is considered over-provisioned if one of the following two conditions is met:

- The total number of virtual cores (vCPU) provisioned on the host is greater than the total number of physical cores (pCPUs).
- The total number of provisioned VMs consume more vCPUs than the total number of pCPUs.

If an instance is over-provisioned, the hypervisor might not guarantee the resources reserved (such as CPU, memory, and others) for the instance due to hypervisor scheduling over-heads, bugs, or limitations with the hypervisor. This behavior can cause lack of CPU resource for Citrix ADC and might lead to the issues mentioned in the first point under **Usage guidelines**. As administrators, you're recommended to reduce the tenancy on the host so that the total number of vCPUs provisioned on the host is lesser or equal to the total number of pCPUs.

Example

For ESX hypervisor, if the `%RDY%` parameter of a VPX vCPU is greater than 0 in the `esx top` command output, the ESX host is said to be having scheduling overheads, which can cause latency

related issues for the VPX instance.

In such a situation, reduce the tenancy on the host so that %RDY% returns to 0 always. Alternatively, contact the hypervisor vendor to triage the reason for not honoring the resource reservation done.

- Hot adding is supported only for PV and SRIOV interfaces with Citrix ADC on AWS. VPX instances with ENA interfaces do not support hot-plug, and the behavior of the instances can be unpredictable if hot-plugging is attempted.
- Hot removing either through the AWS Web console or AWS CLI interface is not supported with PV, SRIOV, and ENA interfaces for Citrix ADC. The behavior of the instances can be unpredictable if hot-removal is attempted.

Commands to control the packet engine CPU usage

You can use two commands (`set ns vpxparam` and `show ns vpxparam`) to control the packet engine (non-management) CPU usage behavior of VPX instances in hypervisor and cloud environments:

- `set ns vpxparam [-cpuyield (YES | NO | DEFAULT)] [-masterclockcpu1 (YES | NO)]`

Allow each VM to use CPU resources that have been allocated to another VM but are not being used.

Set `ns vpxparam` parameters:

-cpuyield: Release or do not release of allocated but unused CPU resources.

- **YES:** Allow allocated but unused CPU resources to be used by another VM.
- **NO:** Reserve all CPU resources for the VM to which they have been allocated. This option shows higher percentage in hypervisor and cloud environments for VPX CPU usage.
- **DEFAULT:** No.

Note:

On all the Citrix ADC VPX platforms, the vCPU usage on the host system is 100 percent. Type the `set ns vpxparam -cpuyield YES` command to override this usage.

If you want to set the cluster nodes to “yield”, you must perform the following extra configurations on CCO:

- If a cluster is formed, all the nodes come up with “yield=DEFAULT”.
- If a cluster is formed using the nodes that are already set to “yield=YES”, then the nodes are added to cluster using the “DEFAULT” yield.

Note:

If you want to set the cluster nodes to “yield=YES”, you can configure only after forming the cluster but not before the cluster is formed.

-masterclockcpu1: You can move the main clock source from CPU0 (management CPU) to CPU1. This parameter has the following options:

- **YES:** Allow the VM to move the main clock source from CPU0 to CPU1.
- **NO:** VM uses CPU0 for the main clock source. By default, CPU0 is the main clock source.

- `show ns vpxparam`

Display the current `vpxparam` settings.

Other References

- For Citrix Ready products, visit [Citrix Ready Marketplace](#).
- For Citrix Ready product support, see the [FAQ page](#).
- For VMware ESX hardware versions, see [Upgrading VMware Tools](#).

Optimize Citrix ADC VPX performance on VMware ESX, Linux KVM, and Citrix Hypervisors

December 17, 2024

The Citrix ADC VPX performance greatly varies depending on the hypervisor, allocated system resources, and the host configurations. To achieve the desired performance, first follow the recommendations in the VPX data sheet, and then further optimize it using the best practices provided in this document.

Citrix ADC VPX instance on VMware ESX hypervisors

This section contains details of configurable options and settings, and other suggestions that help you achieve optimal performance of Citrix ADC VPX instance on VMware ESX hypervisors.

- [Recommended configuration on ESX hosts](#)
- [Citrix ADC VPX with E1000 network interfaces](#)
- [Citrix ADC VPX with VMXNET3 network interfaces](#)
- [Citrix ADC VPX with SR-IOV and PCI passthrough network interfaces](#)

Recommended configuration on ESX hosts

To achieve high performance for VPX with E1000, VMXNET3, SR-IOV, and PCI passthrough network interfaces, follow these recommendations:

- The total number of virtual CPUs (vCPUs) provisioned on the ESX host must be less than or equal to the total number of physical CPUs (pCPUs) on the ESX host.
- Non-uniform Memory Access (NUMA) affinity and CPU affinity must be set for the ESX host to achieve good results.

–To find the NUMA affinity of a Vmnic, log in to the host locally or remotely, and type:

```
1 #vsish -e get /net/pNics/vmnic7/properties | grep NUMA
2 Device NUMA Node: 0
```

- To set NUMA and vCPU affinity for a VM, see [VMware documentation](#).

Citrix ADC VPX with E1000 network interfaces

Perform the following settings on the VMware ESX host:

- On the VMware ESX host, create two vNICs from one pNIC vSwitch. Multiple vNICs create multiple Rx threads in the ESX host. This increases the Rx throughput of the pNIC interface.
- Enable VLANs on the vSwitch port group level for each vNIC that you have created.
- To increase vNIC transmit (Tx) throughput, use a separate Tx thread in the ESX host per vNIC. Use the following ESX command:

- For ESX version 5.5:

```
1 esxcli system settings advanced set -o /Net/NetTxWorldlet -i
```

- For ESX version 6.0 onwards:

```
1 esxcli system settings advanced set -o /Net/NetVMTxType -i 1
```

- To further increase the vNIC Tx throughput, use a separate Tx completion thread and Rx threads per device (NIC) queue. Use the following ESX command:


```
1 esxcli system settings advanced set -o /Net/NetNetqRxQueueFeatPairEnable -i 0
```

Note:

Make sure that you reboot the VMware ESX host to apply the updated settings.

Two vNICs per pNIC deployment

The following is a sample topology and configuration commands for the **Two vNICs per pNIC** model of deployment that delivers better network performance.



Citrix ADC VPX sample configuration:

To achieve the deployment shown in the preceding sample topology, perform the following configuration on the Citrix ADC VPX instance:

- On the client side, bind the SNIP (1.1.1.2) to network interface 1/1 and enable the VLAN tag mode.

```
1 bind vlan 2 -ifnum 1/1 -tagged
2 bind vlan 2 -IPAddress 1.1.1.2 255.255.255.0
```

- On the server side, bind the SNIP (2.2.2.2) to network interface 1/1 and enable the VLAN tag mode.

```
1 bind vlan 3 -ifnum 1/2 -tagged
2 bind vlan 3 -IPAddress 2.2.2.2 255.255.255.0
```

- Add an HTTP virtual server (1.1.1.100) and bind it to a service (2.2.2.100).

```

1  add lb vserver v1 HTTP 1.1.1.100 80 -persistenceType NONE -
    Listenpolicy None -cltTimeout 180
2  add service s1 2.2.2.100 HTTP 80 -gslb NONE -maxClient 0 -maxReq
    0 -cip DISABLED -usip NO -useproxyport YES -sp ON -cltTimeout
    180 -svrTimeout 360 -CKA NO -TCPB NO -CMP NO
3  bind lb vserver v1 s1

```

Note:

Make sure that you include the following two entries in the route table:

- 1.1.1.0/24 subnet with gateway pointing to SNIP 1.1.1.2
- 2.2.2.0/24 subnet with gateway pointing to SNIP 2.2.2.2

Citrix ADC VPX with VMXNET3 network interfaces

To achieve high performance for VPX with VMXNET3 network interfaces, do the following settings on the VMware ESX host:

- Create two vNICs from one pNIC vSwitch. Multiple vNICs create multiple Rx threads in the ESX host. This increases the Rx throughput of the pNIC interface.
- Enable VLANs on the vSwitch port group level for each vNIC that you have created.
- To increase vNIC transmit (Tx) throughput, use a separate Tx thread in the ESX host per vNIC. Use the following ESX commands:

- For ESX version 5.5:

```
1  esxcli system settings advanced set -o /Net/NetTxWorldlet -i
```

- For ESX version 6.0 onwards:

```
1  esxcli system settings advanced set -o /Net/NetVMTxType -i 1
```

On the VMware ESX host, perform the following configuration:

- On the VMware ESX host, create two vNICs from 1 pNIC vSwitch. Multiple vNICs create multiple Tx and Rx threads in the ESX host. This increases the Tx and Rx throughput of the pNIC interface.
- Enable VLANs on the vSwitch port group level for each vNIC that you have created.
- To increase Tx throughput of a vNIC, use a separate Tx completion thread and Rx threads per device (NIC) queue. Use the following command:

```
1  esxcli system settings advanced set -o /Net/
    NetNetqRxQueueFeatPairEnable -i 0
```

- Configure a VM to use one transmit thread per vNIC, by adding the following setting to the VM's configuration:

```
1 ethernetX.ctxPerDev = "1"
```

Note:

Make sure that you reboot the VMware ESX host to apply the updated settings.

You can configure VMXNET3 as a **Two vNICs per pNIC** deployment. For more information, see [Two vNICs per pNIC deployment](#).

Citrix ADC VPX with SR-IOV and PCI passthrough network interfaces

To achieve high performance for VPX with SR-IOV and PCI passthrough network interfaces, see [Recommended configuration on ESX hosts](#).

Citrix ADC VPX instance on Linux-KVM platform

This section contains details of configurable options and settings, and other suggestions that help you achieve optimal performance of Citrix ADC VPX instance on Linux-KVM platform.

- [Performance settings for KVM](#)
- [Citrix ADC VPX with PV network interfaces](#)
- [Citrix ADC VPX with SR-IOV and Fortville PCIe passthrough network interfaces](#)

Performance settings for KVM

Perform the following settings on the KVM host:

Find the NUMA domain of the NIC using the `ls topo` command:

Make sure that memory for the VPX and the CPU is pinned to the same location.

In the following output, the 10G NIC “ens2” is tied to NUMA domain #1.

```
[root@localhost ~]# lstopo-no-graphics
Machine (128GB)
  NUMANode L#0 (P#0 64GB)
    Socket L#0 + L3 L#0 (20MB)
      L2 L#0 (256KB) + L1d L#0 (32KB) + L1l L#0 (32KB) + Core L#0 + PU L#0 (P#0)
      L2 L#1 (256KB) + L1d L#1 (32KB) + L1l L#1 (32KB) + Core L#1 + PU L#1 (P#1)
      L2 L#2 (256KB) + L1d L#2 (32KB) + L1l L#2 (32KB) + Core L#2 + PU L#2 (P#2)
      L2 L#3 (256KB) + L1d L#3 (32KB) + L1l L#3 (32KB) + Core L#3 + PU L#3 (P#3)
      L2 L#4 (256KB) + L1d L#4 (32KB) + L1l L#4 (32KB) + Core L#4 + PU L#4 (P#4)
      L2 L#5 (256KB) + L1d L#5 (32KB) + L1l L#5 (32KB) + Core L#5 + PU L#5 (P#5)
      L2 L#6 (256KB) + L1d L#6 (32KB) + L1l L#6 (32KB) + Core L#6 + PU L#6 (P#6)
      L2 L#7 (256KB) + L1d L#7 (32KB) + L1l L#7 (32KB) + Core L#7 + PU L#7 (P#7)
    HostBridge L#0
      PCI 8086:1521
        Net L#0 "eno1"
      PCI 8086:1521
        Net L#1 "eno2"
      PCI 8086:1584
        Net L#2 "ens3"
      PCI 8086:1584
        Net L#3 "ens4"
      PCI 8086:8d52
        Block L#4 "sda"
        Block L#5 "sdb"
      PCI 8086:8d52
        GPU L#6 "card0"
        GPU L#7 "control064"
      PCI 8086:8d82
      NUMANode L#1 (P#1 64GB)
        Socket L#1 + L3 L#1 (20MB)
          L2 L#8 (256KB) + L1d L#8 (32KB) + L1l L#8 (32KB) + Core L#8 + PU L#8 (P#8)
          L2 L#9 (256KB) + L1d L#9 (32KB) + L1l L#9 (32KB) + Core L#9 + PU L#9 (P#9)
          L2 L#10 (256KB) + L1d L#10 (32KB) + L1l L#10 (32KB) + Core L#10 + PU L#10 (P#10)
          L2 L#11 (256KB) + L1d L#11 (32KB) + L1l L#11 (32KB) + Core L#11 + PU L#11 (P#11)
          L2 L#12 (256KB) + L1d L#12 (32KB) + L1l L#12 (32KB) + Core L#12 + PU L#12 (P#12)
          L2 L#13 (256KB) + L1d L#13 (32KB) + L1l L#13 (32KB) + Core L#13 + PU L#13 (P#13)
          L2 L#14 (256KB) + L1d L#14 (32KB) + L1l L#14 (32KB) + Core L#14 + PU L#14 (P#14)
          L2 L#15 (256KB) + L1d L#15 (32KB) + L1l L#15 (32KB) + Core L#15 + PU L#15 (P#15)
        HostBridge L#6
          PCI 8086:1584
            Net L#8 "ens2"
          PCI 8086:10fb
            Net L#9 "ens1f0"
          PCI 8086:10fb
            Net L#10 "ens1f1"
          PCI ffff:ffff
            Net L#11 "enp131s16"
    [root@localhost ~]# modprobe kvm-intel acpienv=N
```

Allocate the VPX memory from the NUMA domain.

The `numactl` command indicates the NUMA domain from which the memory is allocated. In the following output, around 10 GB RAM is allocated from NUMA node #0.

```
[root@localhost ~]# numactl --hardware
available: 2 nodes (0-1)
node 0 cpus: 0 1 2 3 4 5 6 7
node 0 size: 65429 MB
node 0 free: 55854 MB
node 1 cpus: 8 9 10 11 12 13 14 15
node 1 size: 65536 MB
node 1 free: 52388 MB
node distances:
node  0  1
  0:  10 21
  1:  21 10
[root@localhost ~]#
```

To change the NUMA node mapping, follow these steps.

1. Edit the .xml of the VPX on the host.

```
1 /etc/libvirt/qemu/<VPX_name>.xml
```

2. Add the following tag:

```

1 <numatune>
2 <memory mode="strict" nodeset="1"/>   ☒ This is the NUMA domain
   name
3 </numatune>

```

3. Shut down the VPX.
4. Run the following command:

```
1 virsh define /etc/libvirt/qemu/<VPX_name>.xml
```

This command updates the configuration information for the VM with the NUMA node mappings.

5. Power on the VPX. Then check the `numactl --hardware` command output on the host to see the updated memory allocations for the VPX.

```

[root@localhost ~]# numactl --hardware
available: 2 nodes (0-1)
node 0 cpus: 0 1 2 3 4 5 6 7
node 0 size: 65429 MB
node 0 free: 65429 MB
node 1 cpus: 8 9 10 11 12 13 14 15
node 1 size: 65536 MB
node 1 free: 55854 MB
node distances:
node 0 1
  0: 10 21
  1: 21 10
[root@localhost ~]#

```

Pin vCPUs of VPX to physical cores.

- To view the vCPU to pCPU mappings of a VPX, type the following command

```
1 virsh vcpupin <VPX name>
```

```

root@localhost qemu]# virsh vcpupin NS-VPX-DVR
CPU: CPU Affinity
-----
0: 8
1: 9
2: 10
3: 11

```

The vCPUs 0–4 are mapped to physical cores 8–11.

- To view the current pCPU usage, type the following command:

```
1 mpstat -P ALL 5
```

```
[root@localhost qemu]# mpstat -P ALL 5
Linux 3.10.0-123.el7.x86_64 (localhost.localdomain) 05/17/2016 _x86_64_ (16 CPU)

02:26:20 PM CPU      %usr   %nice    %sys %iowait    %irq   %soft  %steal  %guest  %gnice   %idle
02:26:25 PM all      0.24    0.00    1.67    0.00    0.00    0.00    0.00    17.32    0.00    80.78
02:26:25 PM 0        0.20    0.00    1.00    0.00    0.00    0.00    0.00    0.00    0.00    98.80
02:26:25 PM 1        0.20    0.00    0.20    0.00    0.00    0.00    0.00    0.00    0.00    99.60
02:26:25 PM 2        0.20    0.00    0.40    0.00    0.00    0.00    0.00    0.00    0.00    99.40
02:26:25 PM 3        0.00    0.00    0.20    0.00    0.00    0.00    0.00    0.00    0.00    99.80
02:26:25 PM 4        0.20    0.00    0.20    0.00    0.00    0.00    0.00    0.00    0.00    99.60
02:26:25 PM 5        0.60    0.00    0.20    0.00    0.00    0.00    0.00    0.00    0.00    99.20
02:26:25 PM 6        0.40    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    99.60
02:26:25 PM 7        1.62    0.00    1.42    0.00    0.00    0.00    0.00    0.00    0.00    96.96
02:26:25 PM 8        0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    100.00
02:26:25 PM 9        0.00    0.00    7.60    0.00    0.00    0.00    0.00    92.40    0.00    0.00
02:26:25 PM 10       0.20    0.00    7.00    0.00    0.00    0.00    0.00    92.80    0.00    0.00
02:26:25 PM 11       0.00    0.00    8.60    0.00    0.00    0.00    0.00    91.40    0.00    0.00
02:26:25 PM 12       0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    100.00
02:26:25 PM 13       0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    100.00
02:26:25 PM 14       0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    100.00
02:26:25 PM 15       0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    0.00    100.00
```

In this output, 8 is management CPU, and 9–11 are packet engines.

- To change the vCPU to pCPU pinning, there are two options.
 - Change it at runtime after the VPX boots up using the following command:

```
1 virsh vcpupin <VPX name> <vCPU id> <pCPU number>
2 virsh vcpupin NetScaler-VPX-XML 0 8
3 virsh vcpupin NetScaler-VPX-XML 1 9
4 virsh vcpupin NetScaler-VPX-XML 2 10
5 virsh vcpupin NetScaler-VPX-XML 3 11
```

- To make static changes to the VPX, edit the `.xml` file as before with the following tags:

1. Edit the `.xml` file of the VPX on the host

```
1 /etc/libvirt/qemu/<VPX_name>.xml
```

2. Add the following tag:

```
1 <vcpu placement='static' cpuset='8-11'>4</vcpu>
2   <cputune>
3     <vcpupin vcpu='0' cpuset='8' />
4     <vcpupin vcpu='1' cpuset='9' />
5     <vcpupin vcpu='2' cpuset='10' />
6     <vcpupin vcpu='3' cpuset='11' />
7   </cputune>
```

3. Shut down the VPX.
4. Update the configuration information for the VM with the NUMA node mappings using the following command:

```
1 virsh define /etc/libvirt/qemu/ <VPX_name>.xml
```


5. Power on the VPX. Then check the `virsh vcpupin <VPX name>` command output on the host to see the updated CPU pinning.

Eliminate host interrupt overhead.

- Detect VM_EXITS using the `kvm_stat` command.

At the hypervisor level, host interrupts are mapped to the same pCPUs on which the vCPUs of the VPX are pinned. This might cause vCPUs on the VPX to get kicked out periodically.

To find the VM exits done by VMs running the host, use the `kvm_stat` command.

```
1 [root@localhost ~]# kvm_stat -1 | grep EXTERNAL
2   kvm_exit(EXTERNAL_INTERRUPT) 1728349 27738
3 [root@localhost ~]#
```

A higher value in the order of 1+M indicates an issue.

If a single VM is present, the expected value is 30–100 K. Anything more than that can indicate that there are one or more host interrupt vectors mapped to the same pCPU.

- Detect host interrupts and migrate host interrupts.

When you run the `concatenate` command for the “/proc/interrupts” file, it displays all the host interrupt mappings. If one or more active IRQs map to the same pCPU, its corresponding counter increments.

Move any interrupts that overlap with your Citrix ADC VPX’s pCPUs to unused pCPUs:

```
1 echo 0000000f > /proc/irq/55/smp_affinity
2 0000000f - - > it is a bitmap, LSBs indicates that IRQ 55 can
   only be scheduled on pCPUs 0 - 3
```

- Disable IRQ balance.

Disable IRQ balance daemon, so that no rescheduling happens on the fly.

```
1 service irqbalance stop
2 service irqbalance show - To check the status
3 service irqbalance start - Enable if needed
```

Make sure you run the `kvm_stat` command to ensure that there are not many counters.

Citrix ADC VPX with PV network interfaces

You can configure para-virtualization (PV), SR-IOV, and PCIe passthrough network interfaces as a **Two vNICs per pNIC** deployment. For more information, see [Two vNICs per pNIC deployment](#).


For optimal performance of PV (virtio) interfaces, follow these steps:

- Identify the NUMA domain to which the PCIe slot/NIC is tied to.

- The Memory and vCPU for the VPX must be pinned to the same NUMA domain.
- Vhost thread must be bound to the CPUs in the same NUMA domain.

Bind the virtual host threads to the corresponding CPUs:

1. Once the traffic is started, run the `top` command on the host.



2. Identify the virtual host process (named as `vhost-<pid-of-qemu>`) affinity.
3. Bind the vHost processes to the physical cores in the NUMA domain identified earlier using the following command:

```
1 taskset -pc <core-id> <process-id>
```

Example:

```
1 taskset -pc 12 29838
```

4. The processor cores corresponding to the NUMA domain can be identified with the following command:

```
1 [root@localhost ~]# virsh capabilities | grep cpu
2 <cpu>
3 </cpu>
4 <cpus num='8'>
5 <cpu id='0' socket_id='0' core_id='0' siblings='0' />
6 <cpu id='1' socket_id='0' core_id='1' siblings='1' />
7 <cpu id='2' socket_id='0' core_id='2' siblings='2' />
8 <cpu id='3' socket_id='0' core_id='3' siblings='3' />
9 <cpu id='4' socket_id='0' core_id='4' siblings='4' />
10 <cpu id='5' socket_id='0' core_id='5' siblings='5' />
11 <cpu id='6' socket_id='0' core_id='6' siblings='6' />
12 <cpu id='7' socket_id='0' core_id='7' siblings='7' />
```

```

13     </cpus>
14
15     <cpus num='8'>
16     <cpu id='8' socket_id='1' core_id='0' siblings='8'/>
17     <cpu id='9' socket_id='1' core_id='1' siblings='9'/>
18     <cpu id='10' socket_id='1' core_id='2' siblings='10'/>
19     <cpu id='11' socket_id='1' core_id='3' siblings='11'/>
20     <cpu id='12' socket_id='1' core_id='4' siblings='12'/>
21     <cpu id='13' socket_id='1' core_id='5' siblings='13'/>
22     <cpu id='14' socket_id='1' core_id='6' siblings='14'/>
23     <cpu id='15' socket_id='1' core_id='7' siblings='15'/>
24     </cpus>
25
26     <cpuselection/>
27     <cpuselection/>

```

Bind the QEMU process to the corresponding physical core:

1. Identify the physical cores on which the QEMU process is running. For more information, see the preceding output.
2. Bind the QEMU process to the same physical cores to which you bind the vCPUs, using the following command:

```
1 taskset -pc 8-11 29824
```

Citrix ADC VPX with SR-IOV and Fortville PCIe passthrough network interfaces

For optimal performance of the SR-IOV and Fortville PCIe passthrough network interfaces, follow these steps:

- Identify the NUMA domain to which the PCIe slot/NIC is tied to.
- The Memory and vCPU for the VPX must be pinned to the same NUMA domain.

Sample VPX XML file for vCPU and memory pinning for Linux KVM:

```

1     <domain type='kvm'>
2     <name>NetScaler-VPX</name>
3     <uuid>138f7782-1cd3-484b-8b6d-7604f35b14f4</uuid>
4     <memory unit='KiB'>8097152</memory>
5     <currentMemory unit='KiB'>8097152</currentMemory>
6     <vcpu placement='static'>4</vcpu>
7
8     <cputune>
9     <vcupin vcpu='0' cpuset='8'/>
10    <vcupin vcpu='1' cpuset='9'/>
11    <vcupin vcpu='2' cpuset='10'/>
12    <vcupin vcpu='3' cpuset='11'/>
13    </cputune>
14

```

```
15     <numatune>
16     <memory mode='strict' nodeset='1' />
17     </numatune>
18
19     </domain>
```

Citrix ADC VPX instance on Citrix Hypervisors

This section contains details of configurable options and settings, and other suggestions that help you achieve optimal performance of Citrix ADC VPX instance on Citrix Hypervisors.

- [Performance settings for Citrix Hypervisors](#)
- [Citrix ADC VPX with SR-IOV network interfaces](#)
- [Citrix ADC VPX with para-virtualized interfaces](#)

Performance settings for Citrix Hypervisors

Find the NUMA domain of the NIC using the “xl” command:

```
1 xl info -n
```

Pin vCPUs of VPX to physical cores.

```
1 xl vcpu-pin <Netsclaer VM Name> <vCPU id> <physical CPU id>
```

Check binding of vCPUs.

```
1 xl vcpu-list
```

Allocate more than 8 vCPUs to Citrix ADC VMs.

For configuring more than 8 vCPUs, run the following commands from the Citrix Hypervisor console:

```
1 xe vm-param-set uuid=your_vms_uuid VCPUs-max=16
2 xe vm-param-set uuid=your_vms_uuid VCPUs-at-startup=16
```

Citrix ADC VPX with SR-IOV network interfaces

For optimal performance of the SR-IOV network interfaces, follow these steps:

- Identify the NUMA domain to which the PCIe slot or NIC is tied to.
- Pin the Memory and vCPU for the VPX to the same NUMA domain.
- Bind the Domain-0 vCPU to the remaining CPU.

Citrix ADC VPX with para-virtualized interfaces

For optimal performance, two vNICs per pNIC and one vNIC per pNIC configurations are advised, as in other PV environments.

To achieve optimal performance of para-virtualized (netfront) interfaces, follow these steps:

- Identify the NUMA domain to which the PCIe slot or NIC is tied to.
- Pin the memory and vCPU for the VPX to the same NUMA domain.
- Bind the Domain-0 vCPU to the remaining CPU of the same NUMA domain.
- Pin host Rx/Tx threads of vNIC to Domain-0 vCPUs.

Pin host threads to Domain-0 vCPUs:

1. Find Xen-ID of the VPX by using the `xl list` command on the Citrix Hypervisor host shell.
2. Identify host threads by using the following command:

```
1 ps -ax | grep vif <Xen-ID>
```

In the following example, these values indicate:

- **vif5.0** - The threads for first interface allocated to VPX in XenCenter (management interface).
- **vif5.1** - The threads for second interface assigned to VPX and so on.

```
[root@xenserver-uuffyqlx ~]# xl list
Name                               ID   Mem VCPUs   State   Time(s)
Domain-0                            0   4092    8   r----- 633321.0
Sai_VPX                              5   8192    4   r----- 1529471.0
[root@xenserver-uuffyqlx ~]#
[root@xenserver-uuffyqlx ~]#
[root@xenserver-uuffyqlx ~]# ps -ax | grep "vif5"
Warning: bad syntax, perhaps a bogus '-'? See /usr/share/doc/procps-3.2.7/FAQ
20447 pts/6      S+    0:00 grep vif5
29187 ?          S     1:09 [vif5.0-guest-rx]
29188 ?          S     0:00 [vif5.0-dealloc]
29189 ?          S    201:33 [vif5.1-guest-rx]
29190 ?          S     80:51 [vif5.1-dealloc]
29191 ?          S     0:20 [vif5.2-guest-rx]
29192 ?          S     0:00 [vif5.2-dealloc]
[root@xenserver-uuffyqlx ~]#
```

3. Pin the threads to Domain-0 vCPUs using the following command:

```
1 taskset -pc <core-id> <process-id>
```

Example:

```
1 taskset -pc 1 29189
```

Apply Citrix ADC VPX configurations at the first boot of the Citrix ADC appliance in cloud

September 12, 2024

You can apply the Citrix ADC VPX configurations during the first boot of the Citrix ADC appliance in a cloud environment. This stage is addressed as the **preboot** stage in this document. Therefore in certain cases like ADC pooled licensing, a specific VPX instance is brought up in much lesser time. This feature is available in Microsoft Azure, Google Cloud platform, and AWS clouds.

What is user data

When you provision a VPX instance in a cloud environment, you have the option of passing user data to the instance. The user data allows you to perform common automated configuration tasks, customize the startup behaviors of instances, and run scripts after the instance starts. At the first boot, the Citrix ADC VPX instance performs the following tasks:

- Reads the user data.
- Interprets the configuration provided in user data.
- Applies the newly added configuration as it boots up.

How to provide preboot user data in cloud instance

You can provide preboot user data to the cloud instance in XML format. Different clouds have different interfaces for providing user data.

Provide preboot user data using the AWS console

When you provision a Citrix ADC VPX instance using the AWS console, navigate to **Configure Instance Details > Advanced Details**, and provide the preboot user data configuration in the **User data** field.

For detailed instructions on each of the steps, see [Deploy a Citrix ADC VPX instance on AWS by using the AWS web console](#).

For more information, see AWS documentation on [Launching an instance](#).

The screenshot shows the AWS console interface for configuring an EC2 instance. The page is titled "Step 3: Configure Instance Details" and includes a progress bar at the top with steps: 1. Choose AMI, 2. Choose Instance Type, 3. Configure Instance (active), 4. Add Storage, 5. Add Tags, 6. Configure Security Group, and 7. Review.

Key configuration options visible include:

- Domain join directory:** No directory (with "Create new directory" button)
- IAM role:** None (with "Create new IAM role" button)
- Shutdown behavior:** Stop
- Stop - Hibernate behavior:** Enable hibernation as an additional stop behavior
- Enable termination protection:** Protect against accidental termination
- Monitoring:** Enable CloudWatch detailed monitoring (with "Additional charges apply" link)
- Tenancy:** Shared - Run a shared hardware instance (with "Additional charges will apply for dedicated tenancy" link)
- Credit specification:** Unlimited (with "Additional charges may apply" link)
- File systems:** Add file system (with "Create new file system" button)

The "Advanced Details" section is expanded, showing:

- Metadata accessible:** Enabled
- Metadata version:** V1 and V2 (token optional)
- Metadata token response hop limit:** 1
- User data:** As text, As file, Input is already base64 encoded. Below this are radio buttons for "(Optional)" and a text input field.

The "User data" section, including the radio buttons and the text input field, is highlighted with a yellow rectangular box.

Provide preboot user data using AWS CLI

Type the following command in the AWS CLI:

```

1 aws ec2 run-instances \
2   --image-id ami-0abcdef1234567890 \
3   --instance-type t2.micro \
4   --count 1 \
5   --subnet-id subnet-08fc749671b2d077c \
6   --key-name MyKeyPair \
7   --security-group-ids sg-0b0384b66d7d692f9 \
8   --user-data file://my_script.txt

```

For more information, see AWS documentation on [Running instances](#).

For more information, see AWS documentation on [Using instance user data](#)

Provide preboot user data using the Azure console

When you provision a Citrix ADC VPX instance using Azure console, navigate to **Create a virtual machine > Advanced** tab. In the **Custom data** field, provide preboot user data configuration.

[Home](#) > [Virtual machines](#) >

Create a virtual machine

Basics Disks Networking Management **Advanced** Tags Review + create

Add additional configuration, agents, scripts or applications via virtual machine extensions or cloud-init.

Extensions

Extensions provide post-deployment configuration and automation.

Extensions ⓘ

[Select an extension to install](#)

Custom data

Pass a script, configuration file, or other data into the virtual machine while it is being provisioned. The data will be saved on the VM in a known location. [Learn more about custom data for VMs](#) ⓘ

Custom data

i Custom data on the selected image will be processed by cloud-init. [Learn more about custom data and cloud init](#) ⓘ

Host

Azure Dedicated Hosts allow you to provision and manage a physical server within our data centers that are dedicated to your Azure subscription. A dedicated host gives you assurance that only VMs from your subscription are on the host, flexibility to choose VMs from your subscription that will be provisioned on the host, and the control of platform maintenance at the level of the host. [Learn more](#)

Host group ⓘ

No host group found

Provide preboot user data using the Azure CLI

Type the following command in the Azure CLI:

```
1 az vm create \  
2   --resource-group myResourceGroup \  
3   --name MyVm \  
4   --image debian \  
5   --custom-data MyCloudInitScript.txt \
```

Example:

```
1 az vm create --resource-group MyResourceGroup -name MyVm --image debian \  
   --custom-data MyCloudInitScript.txt
```

You can pass your custom data or preboot configuration as a file to “--custom-data”parameter. In this

example, the file name is **MyCloudInitScript.txt**.

For more information, see [Azure CLI documentation](#).

Provide preboot user data using the GCP console

When you provision a Citrix ADC VPX instance using GCP console, fill in the properties of instance. Expand **Management, security, disks, networking, sole tenancy**. Navigate to the **Management** tab. In the **Automation** section, provide preboot user data configuration in the **Startup Script** field.

For detailed information on creating the VPX instance using GCP, see [Deploy a Citrix ADC VPX instance on Google Cloud Platform](#).

The screenshot shows the GCP console configuration page for a VM instance. The 'Management' tab is selected. The 'Automation' section is highlighted with a yellow border. It contains the following fields:

- Description (Optional)**: A text input field.
- Deletion protection**: A checkbox labeled 'Enable deletion protection' with a note: 'When deletion protection is enabled, instance cannot be deleted. [Learn more](#)'
- Reservations**: A dropdown menu labeled 'Use an existing reservation when creating this VM instance' with the selected option 'Automatically use created reservation'.
- Automation**: A section containing:
 - Startup script (Optional)**: A text input field with a yellow border. The text below it reads: 'You can choose to specify a startup script that will run when your instance boots up or restarts. Startup scripts can be used to install software and updates, and to ensure that services are running within the virtual machine. [Learn more](#)'
 - Metadata (Optional)**: A section with a text input field and a note: 'You can set custom metadata for an instance or project outside of the server-defined metadata. This is useful for passing in arbitrary values to your project or instance that can be queried by your code on the instance. [Learn more](#)'
 - A table with columns 'Key' and 'Value' and an '+ Add item' button.

Provide preboot user data using the gcloud CLI

Type the following command in the GCP CLI:

```
1 gcloud compute instances create INSTANCE_NAMES --metadata-from-file=
  startup-script=LOCAL_FILE_PATH
```

metadata-from-file - Reads the value or user data from a file stored at the .

For more information, see [gcloud CLI documentation](#)

Preboot user data format

The preboot user data must be provided to the cloud instance in XML format. The Citrix ADC preboot user data that you provide through the cloud infrastructure during boot can comprise the following four sections:

- Citrix ADC configuration represented with the `<NS-CONFIG>` tag.
- Custom bootstrapping the Citrix ADC represented with the `<NS-BOOTSTRAP>` tag.
- Storing user-scripts in Citrix ADC represented with the `<NS-SCRIPTS>` tag.
- Pooled licensing configuration represented with the `<NS-LICENSE-CONFIG>` tag.

You can provide the preceding four sections in any order within the ADC preboot configuration. Ensure to strictly follow the formatting shown in the following sections while providing the preboot user data.

Note:

The entire preboot user data configuration must be enclosed in the `<NS-PRE-BOOT-CONFIG>` tag as shown in the following examples.

Example 1:

```
1 <NS-PRE-BOOT-CONFIG>
2     <NS-CONFIG>           </NS-CONFIG>
3     <NS-BOOTSTRAP>       </NS-BOOTSTRAP>
4     <NS-SCRIPTS>         </NS-SCRIPTS>
5     <NS-LICENSE-CONFIG>  </NS-LICENSE-CONFIG>
6 </NS-PRE-BOOT-CONFIG>
```

Example 2:

```
1 <NS-PRE-BOOT-CONFIG>
2     <NS-LICENSE-CONFIG> </NS-LICENSE-CONFIG>
3     <NS-SCRIPTS>       </NS-SCRIPTS>
4     <NS-BOOTSTRAP>     </NS-BOOTSTRAP>
5     <NS-CONFIG>        </NS-CONFIG>
6 </NS-PRE-BOOT-CONFIG>
```

Use the `<NS-CONFIG>` tag to provide the specific Citrix ADC VPX configurations that needs to be applied to the VPX instance at the preboot stage.

Note:

The `<NS-CONFIG>` section must have valid ADC CLI commands. The CLIs are not verified for the syntactic errors or format.

Citrix ADC configurations

Use the `<NS-CONFIG>` tag to provide the specific Citrix ADC VPX configurations that needs to be applied to the VPX instance at the preboot stage.

Note:

The `<NS-CONFIG>` section must have valid ADC CLI commands. The CLIs are not verified for the syntactic errors or format.

Example:

In the following example, the `<NS-CONFIG>` section has the details of the configurations. A VLAN of ID '5' is configured and bound to the SNIP (5.0.0.1). A load balancing virtual server (4.0.0.101) is also configured.

```

<NS-CONFIG>
  add vlan 5
  add ns ip 5.0.0.1 255.255.255.0

  bind vlan 5 -IPAddress 5.0.0.1 255.255.255.0
  enable ns feature WL SP LB RESPONDER
  add server 5.0.0.201 5.0.0.201
  add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -maxClient 0 -maxReq 0 -cip
D SABLED -usip
NO -useproxyport YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -TCPB NO -CMP NO
  add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -persistenceType NONE -cltTimeout 180
</NS-CONFIG>
</NS-CONFIG>

```

You can copy the configuration shown in the preceding screenshot from here:

```

1 <NS-CONFIG>
2   <NS-CONFIG>
3     add vlan 5
4     add ns ip 5.0.0.1 255.255.255.0
5     bind vlan 5 -IPAddress 5.0.0.1 255.255.255.0

```

```

6      enable ns feature WL SP LB RESPONDER
7      add server 5.0.0.201 5.0.0.201
8      add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -
          maxClient 0 -maxReq 0 -cip DISABLED -usip
9      NO -useproxyport YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -
          TCPB NO -CMP NO
10     add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -
          persistenceType NONE -cltTimeout 180
11     </NS-CONFIG>
12 </NS-PRE-BOOT-CONFIG>

```

The Citrix ADC VPX instance comes up with the configuration applied in the <NS-CONFIG> section as shown in the following illustrations.

```

> sh ns ip
-----
1) 10.160.0.72      0      NetScaler IP      Active  Enabled  Enabled  NA      Enabled
2) 5.0.0.1         0      SNIP               Active  Enabled  Enabled  NA      Enabled
3) 4.0.0.101       0      VIP                Active  Enabled  Enabled  Enabled Enabled
Done
> sh vlan
1)  VLAN ID: 1
   Link-local IPv6 addr: fe80::4001:aff:fea0:48/64
   Interfaces : 1/1 1/2 LO/1
2)  VLAN ID: 5      VLAN Alias Name:
   IPs :
      5.0.0.1      Mask: 255.255.255.0
3)  VLAN ID: 10     VLAN Alias Name:
   Interfaces : 0/1
   IPs :
      10.160.0.72  Mask: 255.255.240.0
Done

```

```

> sh server
1)  Name:      5.0.0.201      State:ENABLED
    IPAddress: 5.0.0.201
2)  Name:      169.254.169.254  State:ENABLED
    IPAddress: 169.254.169.254
Done
> stat service

Service(s) Summary
      IP  port      Type      State      Req/s
preb...s_201      5.0.0.201      80      HTTP      DOWN      0/s
gcpl...vice0 169.254.169.254      53      DNS      UP      0/s
Done
> sh service preboot_s5_201
preboot_s5_201 (5.0.0.201:80) - HTTP
State: DOWN
Last state change was at Tue Dec 29 07:18:28 2020
Time since last state change: 0 days, 00:05:02.820
Server Name: 5.0.0.201
Server ID : None      Monitor Threshold : 0
Max Conn: 0      Max Req: 0      Max Bandwidth: 0 kbits
Use Source IP: NO
Client Keepalive (CKA): NO
Monitoring Owner: 0
Access Down Service: NO
TCP Buffering (TCPB): NO
HTTP Compression (CMP): NO
Idle timeout: Client: 180 sec      Server: 360 sec
Client IP: DISABLED
Cacheable: NO
SC: OFF
SP: OFF
Down state flush: ENABLED
Monitor Connection Close : NONE
Appflow logging: ENABLED
Process Local: DISABLED

```

User scripts

Use the `<NS-SCRIPTS>` tag to provide any script that must be stored and ran in Citrix ADC VPX instance.

You can include many scripts within the `<NS-SCRIPTS>` tag. Each script must be included within the `<SCRIPT>` tag.

Each `<SCRIPT>` section corresponds to one script and contains all the details of the script using the following sub tags.

- **<SCRIPT-NAME>**: Indicates the name of the script file that must be stored.
- **<SCRIPT-CONTENT>**: Indicates the content of the file that must be stored.
- **<SCRIPT-TARGET-LOCATION>**: Indicates the designated target location where this file must be stored. If the target location is not provided, by default, the file, or script is saved in the “/nsconfig” directory.
- **<SCRIPT-NS-BOOTUP>**: Specify the commands that you use to run the script.

- If you use the `<SCRIPT-NS-BOOTUP>` section, the commands provided in the section are stored in “/nsconfig/nsafter.sh”, and the commands are run after the packet engine boots up as part of “nsafter.sh” execution.
- If you do not use the `<SCRIPT-NS-BOOTUP>` section, the script file is stored in the target location that you specify.

Example 1:

In this example, the `<NS-SCRIPTS>` tag contains details of only one script: script-1.sh. The “script-1.sh” script is saved at the “/var” directory. The script is populated with the specified contents, and is run with the “sh /var/script-1.sh” command after packet engine boots up.

```

<NS-PRE-BOOT-CONFIG>
  <NS-SCRIPTS>
    <SCRIPT>
      <SCRIPT-CONTENT>
        #Shell script
        echo "Running script 1" > /var/script-1.output
        date >> /var/script-1.output
      </SCRIPT-CONTENT>
      <SCRIPT-NAME> script-1.sh </SCRIPT-NAME>
      <SCRIPT-TARGET-LOCATION> /var/ </SCRIPT-TARGET-LOCATION>
      <SCRIPT-NS-BOOTUP>sh /var/script-1.sh</SCRIPT-NS-BOOTUP>
    </SCRIPT>
  </NS-SCRIPTS>
</NS-PRE-BOOT-CONFIG>

```

You can copy the configuration shown in the preceding screenshot from here:

```

1 <NS-PRE-BOOT-CONFIG>
2   <NS-SCRIPTS>
3     <SCRIPT>
4       <SCRIPT-CONTENT>
5         #Shell script
6         echo "Running script 1" > /var/script-1.output
7         date >> /var/script-1.output
8       </SCRIPT-CONTENT>
9
10      <SCRIPT-NAME> script-1.sh </SCRIPT-NAME>
11      <SCRIPT-TARGET-LOCATION> /var/ </SCRIPT-TARGET-LOCATION>
12      <SCRIPT-NS-BOOTUP>sh /var/script-1.sh</SCRIPT-NS-BOOTUP>
13    </SCRIPT>
14  </NS-SCRIPTS>
15 </NS-PRE-BOOT-CONFIG>

```

In the following snapshot, you can verify that “script-1.sh” script is saved in the “/var/” directory. The “Script-1.sh” script is run, and the output file is created appropriately.

```

root@ns#
root@ns# ls /var/
.monit.id          core              gui               nsinstall        pubkey
.monit.state      crash            install          nslog            python
.snap             cron             krb              nsproflog       run
AAA               db               learnt_data      nssynclog       safenet
app_catalog       dev             log              nstemplates     script-1.output
cloudhadaemon     download        mastools        nstmp           script-1.sh
cloudhadaemon.tgz empty           netscaler       nstrace         tmp
clusterd          file-2.txt      ns_gui          opt             vpn
configdb          gcfl           ns_sys_backup  osr_compliance  vpns
root@ns#
root@ns# cat /var/script-1.sh
#Shell script
echo "Running script 1" > /var/script-1.output
date >> /var/script-1.output
root@ns#
root@ns# cat /var/script-1.output
Running script 1
Wed Jan  6 05:25:33 UTC 2021
root@ns#
root@ns#

```

Example 2:

In the following example, the <NS-SCRIPTS> tag contains details of two scripts.

- The first script is saved as “script-1.sh” at the “/var” directory. The script is populated with the specified contents, and is run with command “sh /var/script-1.sh” after packet engine boots up.
- The second script is saved as “file-2.txt” at the “/var” directory. This file is populated with the specified contents. But it is not run because the bootup execution command <SCRIPT-NS-BOOTUP> is not provided.

```

<NS-PRE-BOOT-CONFIG>
  <NS-SCRIPTS>
    <SCRIPT>
      <SCRIPT-CONTENT>
      #Shell script
      echo "Running script 1" > /var/script-1.output
      date >> /var/script-1.output
      </SCRIPT-CONTENT>
      <SCRIPT-NAME> script-1.sh </SCRIPT-NAME>
      <SCRIPT-TARGET-LOCATION> /var/ </SCRIPT-TARGET-LOCATION>
      <SCRIPT-NS-BOOTUP>sh /var/script-1.sh</SCRIPT-NS-BOOTUP>
    </SCRIPT>
    <SCRIPT>
      <SCRIPT-CONTENT>
      This script has no execution point. It will just be saved at the target location. NS Consumer module should consume this script/file.
      </SCRIPT-CONTENT>
      <SCRIPT-NAME>file-2.txt</SCRIPT-NAME>
      <SCRIPT-TARGET-LOCATION>/var/</SCRIPT-TARGET-LOCATION>
    </SCRIPT>
  </NS-SCRIPTS>
</NS-PRE-BOOT-CONFIG>

```

You can copy the configuration shown in the preceding screenshot from here:

```

1 <NS-PRE-BOOT-CONFIG>
2   <NS-SCRIPTS>
3     <SCRIPT>
4       <SCRIPT-CONTENT>
5         #Shell script
6         echo "Running script 1" > /var/script-1.output
7         date >> /var/script-1.output
8       </SCRIPT-CONTENT>
9
10      <SCRIPT-NAME> script-1.sh </SCRIPT-NAME>
11      <SCRIPT-TARGET-LOCATION> /var/ </SCRIPT-TARGET-LOCATION>
12      <SCRIPT-NS-BOOTUP>sh /var/script-1.sh</SCRIPT-NS-BOOTUP>
13    </SCRIPT>
14
15    <SCRIPT>
16      <SCRIPT-CONTENT>
17        This script has no execution point.
18        It will just be saved at the target location
19        NS Consumer module should consume this script/file
20      </SCRIPT-CONTENT>
21      <SCRIPT-NAME>file-2.txt</SCRIPT-NAME>
22      <SCRIPT-TARGET-LOCATION>/var/</SCRIPT-TARGET-LOCATION>
23    </SCRIPT>
24  </NS-SCRIPTS>
25 </NS-PRE-BOOT-CONFIG>

```

In the following snapshot, you can verify that script-1.sh and file-2.txt are created in the “/var/” directory. The Script-1.sh is run, and the output file is created appropriately.

```

root@ns# ls /var/
.monit.id          core              gui               nsinstall        pubkey
.monit.state      crash            install          nslog            python
.snap            cron             krb              nsproflog        run
AAA              db              learnt_data      nssynclog        safenet
app_catalog      dev             log              nstemplates     script-1.output
cloudhadaemon    download        mastools         nstmp            script-1.sh
cloudhadaemon.tgz empty           netscaler       nstrace          tmp
clusterd         file-2.txt      ns_gui          opt              vpn
configdb         gcfl           ns_sys_backup   osr_compliance  vpns
root@ns#
root@ns# cat /var/script-1.sh
#Shell script
echo "Running script 1" > /var/script-1.output
date >> /var/script-1.output
root@ns#
root@ns# cat /var/script-1.output
Running script 1
Wed Jan  6 05:08:56 UTC 2021
root@ns#
root@ns# cat /var/file-2.txt
This script has no execution point.
It will just be saved at the target location
NS Consumer module should consume this script/file
root@ns#
root@ns#

```


Licensing

Use the `<NS-LICENSE-CONFIG>` tag to apply Citrix ADC pooled licensing while booting up the VPX instance. Use the `<LICENSE-COMMANDS>` tag within `<NS-LICENSE-CONFIG>` section to provide the pooled license commands. These commands must be syntactically valid.

You can specify the pooled licensing details such as, license type, capacity, and license server in the `<LICENSE-COMMANDS>` section using the standard pooled licensing commands. For more information, see [Configure Citrix ADC pooled capacity licensing](#).

After applying the `<NS-LICENSE-CONFIG>`, the VPX comes up with the requested edition upon boot, and VPX tries to check out the configured licenses from the license server.

- If the license checkout is successful, the configured bandwidth is applied to VPX.
- If the license checkout fails, the license is not retrieved from license server within 10–12 minutes approximately. As a result, the system reboots and enters an unlicensed state.

Example:

In the following example, after applying the `<NS-LICENSE-CONFIG>`, the VPX comes up with the Premium edition upon boot, and VPX tries to check out the configured licenses from the license server (10.102.38.214).

```
<NS-PRE-BOOT-CONFIG>
<NS-LICENSE-CONFIG>
  <LICENSE-COMMANDS>

  add ns licenseserver 10.102.38.214 -port 2800
  set ns capacity -unit gbps -bandwidth 3 edition platinum
  </LICENSE-COMMANDS>
</NS-LICENSE-CONFIG>
</NS-PRE-BOOT-CONFIG>
```

You can copy the configuration shown in the preceding screenshot from here:

```
1 <NS-PRE-BOOT-CONFIG>
2   <NS-LICENSE-CONFIG>
3     <LICENSE-COMMANDS>
4       add ns licenseserver 10.102.38.214 -port 2800
5       set ns capacity -unit gbps -bandwidth 3 edition platinum
6     </LICENSE-COMMANDS>
7   </NS-LICENSE-CONFIG>
8 </NS-PRE-BOOT-CONFIG>
```

As shown in the following illustration, you can run the “show license server” command, and verify that the license server (10.102.38.214) is added to the VPX.

```

Done
> sh licenseServer
      License Server: 10.102.38.214      Port: 2800      Status:
Done
>
>
>

```

Bootstrapping

Use the `<NS-BOOTSTRAP>` tag to provide the custom bootstrapping information. You can use the `<SKIP-DEFAULT-BOOTSTRAP>` and `<NEW-BOOTSTRAP-SEQUENCE>` tags within the `<NS-BOOTSTRAP>` section. This section informs Citrix ADC appliance whether to avoid the default bootstrap or not. If the default bootstrapping is avoided, this section provides you an option to provide a new bootstrapping sequence.

Default bootstrap configuration

The default bootstrap configuration in Citrix ADC appliance follows these interface assignments:

- **Eth0** - Management interface with a certain NSIP address.
- **Eth1** - Client-facing interface with a certain VIP address.
- **Eth2** - Server-facing interface with a certain SNIP address.

Customize bootstrap configuration

You can skip the default bootstrap sequence and provide a new bootstrap sequence for the Citrix ADC VPX instance. Use the `<NS-BOOTSTRAP>` tag to provide the custom bootstrapping information. For example, you can change the default bootstrapping, where the Management interface (NSIP), Client-facing interface (VIP), and server-facing interface (SNIP) are always provided in certain order.

The following table indicates the bootstrapping behavior with the different values that are allowed for `<SKIP-DEFAULT-BOOTSTRAP>` and `<NEW-BOOTSTRAP-SEQUENCE>` tags.

<code>SKIP-DEFAULT-BOOTSTRAP</code>	<code>NEW-BOOTSTRAP-SEQUENCE</code>	Bootstrap behavior
YES	YES	The default bootstrapping behavior is skipped, and a new custom bootstrap sequence provided in the <code><NS-BOOTSTRAP></code> section is run.

SKIP-DEFAULT- BOOTSTRAP	NEW-BOOTSTRAP- SEQUENCE	Bootstrap behavior
YES	NO	The default bootstrapping behavior is skipped. The bootstrap commands provided in the <NS-CONFIG> section is run.

You can customize the bootstrap configuration by the following three methods:

- Provide only the interface details
- Provide the interface details along with IP addresses and subnet mask
- Provide bootstrap related commands in the <NS-CONFIG> section

Method 1: Custom bootstrap by specifying only the interface details

You specify the management, client-facing and server-facing interfaces but not their IP addresses and subnet masks. The IP addresses and subnet masks are populated by querying the cloud infrastructure.

Custom bootstrap example for AWS

You provide the custom bootstrap sequence as shown in the following example. For more information, see [How to provide preboot user data in cloud instance](#). Eth1 interface is assigned as the management interface (NSIP), Eth0 interface as the client interface (VIP), and Eth2 interface as the server interface (SNIP). The <NS-BOOTSTRAP> section contains only the interface details and not the details of IP addresses and subnet masks.

```

<NS-PRE-BOOT-CONFIG>
  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>

    <MGMT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth1</INTERFACE-NUM>
    </MGMT-INTERFACE-CONFIG>

    <CLIENT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth0</INTERFACE-NUM>
    </CLIENT-INTERFACE-CONFIG>

    <SERVER-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth2</INTERFACE-NUM>
    </SERVER-INTERFACE-CONFIG>

  </NS-BOOT STRAP>
</NS-PRE-BOOT-CONFIG>
  
```

After the VM instance is created, in the AWS portal, you can verify the network interface properties as follows:

1. Navigate to the **AWS Portal > EC2 instances**, and select the instance that you have created by providing the custom bootstrap information.
2. In the **Description** tab, you can verify the properties of each network interface as shown in the following illustrations.

Network Interface eth1	
Interface ID	eni-021961099be6815eb
VPC ID	vpc-6b258c02
Attachment Owner	566658252593
Attachment Status	attached
Attachment Time	Fri Jan 01 11:11:23 GMT+530 2021
Delete on Terminate	false
Private IP Address	172.31.52.88
Private DNS Name	ip-172-31-52-88.ap-south-1.compute.internal

```
Network Interface eth0
Interface ID   eni-039e5f3329cd879e9
VPC ID        vpc-6b258c02
Attachment Owner 566658252593
Attachment Status attached
Attachment Time  Fri Jan 01 10:58:28 GMT+530 2021
Delete on Terminate true
Private IP Address 172.31.5.155
Private DNS Name  ip-172-31-5-155.ap-south-1.compute.internal
```

```
Network Interface eth2
Interface ID   eni-09e55a6cfb791e68d
VPC ID        vpc-6b258c02
Attachment Owner 566658252593
Attachment Status attached
Attachment Time  Fri Jan 01 11:11:33 GMT+530 2021
Delete on Terminate false
Private IP Address 172.31.76.177
Private DNS Name  ip-172-31-76-177.ap-south-1.compute.internal
```

You can run the `show ns ip` command in **ADC CLI**, and verify the network interfaces applied to the ADC VPX instance during the first boot of the ADC appliance.

```

> sh ns ip
  Ippaddress      Traffic Domain  Type           Mode   Arp   Icmp   Vserver  State
  -----
1)  172.31.52.88    0               NetScaler IP   Active Enabled Enabled NA       Enabled
2)  172.31.76.177  0               SNIP          Active Enabled Enabled NA       Enabled
3)  172.31.5.155   0               VIP           Active Enabled Enabled Enabled Enabled
Done
> sh vlan
1)  VLAN ID: 1
    Link-local IPv6 addr: fe80::839:e2ff:feaf:4a9e/64
    Interfaces : 1/1 1/3 LO/1
2)  VLAN ID: 10    VLAN Alias Name:
    Interfaces : 1/2
    IPs :
        172.31.52.88      Mask: 255.255.240.0
Done
> sh route
  Network      Netmask      Gateway/OwnedIP  VLAN  State  Traffic Domain  Type
  -----
1)  0.0.0.0      0.0.0.0      172.31.48.1      0      UP     0               STATIC
2)  127.0.0.0    255.0.0.0    127.0.0.1        0      UP     0               PERMANENT
3)  172.31.0.0    255.255.240.0  172.31.5.155     0      UP     0               DIRECT
4)  172.31.48.0   255.255.240.0  172.31.52.88     0      UP     0               DIRECT
5)  172.31.64.0   255.255.240.0  172.31.76.177    0      UP     0               DIRECT
6)  172.31.0.2    255.255.255.255  172.31.48.1      0      UP     0               STATIC
Done

```

Custom bootstrap example for Azure

You provide the custom bootstrap sequence as shown in the following example. For more information, see [How to provide preboot user data in cloud instance](#). Eth2 interface is assigned as the management interface (NSIP), Eth1 interface as the client interface (VIP), and Eth0 interface as the server interface (SNIP). The <NS-BOOTSTRAP> section contains only the interface details and not the details of IP addresses and subnet masks.

```


<NS-PRE-BOOT-CONFIG>
  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>

    <MGMT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth2</INTERFACE-NUM>
    </MGMT-INTERFACE-CONFIG>

    <CLIENT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth1</INTERFACE-NUM>
    </CLIENT-INTERFACE-CONFIG>

    <SERVER-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth0</INTERFACE-NUM>
    </SERVER-INTERFACE-CONFIG>
  </NS-BOOTSTRAP>
</NS-PRE-BOOT-CONFIG>
  
```

You can see that the Citrix ADC VPX instance is created with three network interfaces. Navigate to the **Azure portal > VM instance > Networking**, and verify the networking properties of the three NICs as shown in the following illustrations.



You can run the “show nsip” command in the ADC CLI, and verify that the new bootstrap sequence

specified in the <NS-BOOTSTRAP> section is applied. You can run the “show route” command to verify the subnet mask.

```

> sh ns ip
      Ippaddress      Traffic Domain  Type
-----
1)    172.27.2.53      0              NetScaler IP
2)    172.27.0.53      0              SNIP
3)    172.27.1.53      0              VIP
      Mode           Arp           Icmp          Vserver      State
-----
1)    Active         Enabled       Enabled       NA           Enabled
2)    Active         Enabled       Enabled       NA           Enabled
3)    Active         Enabled       Enabled       Enabled      Enabled
Done
> sh vlan
1)    VLAN ID: 1
      Link-local IPv6 addr: fe80::20d:3aff:fec9:c26c/64
      Interfaces : 0/1 1/1 LO/1
2)    VLAN ID: 10     VLAN Alias Name:
      Interfaces : 1/2
      IPs :
          172.27.2.53      Mask: 255.255.255.0
Done
> sh route
      Network          Netmask          Gateway/OwnedIP  VLAN  State  Traffic Domain  Type
-----
1)    0.0.0.0           0.0.0.0          172.27.2.1      0     UP     0              STATIC
2)    127.0.0.0         255.0.0.0        127.0.0.1      0     UP     0              PERMANENT
3)    172.27.0.0        255.255.255.0    172.27.0.53    0     UP     0              DIRECT
4)    172.27.1.0        255.255.255.0    172.27.1.53    0     UP     0              DIRECT
5)    172.27.2.0        255.255.255.0    172.27.2.53    0     UP     0              DIRECT
6)    169.254.0.0       255.255.0.0      172.27.0.1     0     UP     0              STATIC
7)    168.63.129.16     255.255.255.255  172.27.0.1     0     UP     0              STATIC
8)    169.254.169.254   255.255.255.255  172.27.0.1     0     UP     0              STATIC
Done
>

```

Custom bootstrap examples for GCP

You provide the custom bootstrap sequence as shown in the following example. For more information, see [How to provide preboot user data in cloud instance](#). Eth1 interface is assigned as the management interface (NSIP), Eth0 interface as the client interface (VIP), and Eth2 interface as the server interface (SNIP). The <NS-BOOTSTRAP> section contains only the interface details and not the details of IP addresses and subnet masks.


```

<NS-PRE-BOOT-CONFIG>
  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>

    <MGMT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth1</INTERFACE-NUM>
    </MGMT-INTERFACE-CONFIG>

    <CLIENT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth0</INTERFACE-NUM>
    </CLIENT-INTERFACE-CONFIG>

    <SERVER-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth2</INTERFACE-NUM>
    </SERVER-INTERFACE-CONFIG>

  </NS-BOOT STRAP>
</NS-PRE-BOOT-CONFIG>
  
```

After the VM instance is created in the GCP portal, you can verify the network interface properties as follows:

1. Select the instance that you have created by providing the custom bootstrap information.
2. Navigate to the Network interface properties and verify the NIC details as follows:

Network interfaces									
Name	Network	Subnetwork	Primary internal IP	Alias IP ranges	External IP	Network Tier	IP forwarding	Network details	
nic0	default	default	10.160.0.71	–	35.244.56.180 (ephemeral)	Premium	Off	View details	
nic1	vsk-vm-network-1	asia-south1-subnet-1	10.128.0.40	–	35.244.40.113 (ephemeral)	Premium		View details	
nic2	vsk-vm-network-2	asia-south1-subnet-5	10.128.4.27	–	34.93.241.147 (ephemeral)	Premium		View details	
Public DNS PTR Record									
None									

You can run the `show nsip` command in **ADC CLI**, and verify the network interfaces applied to the ADC VPX instance during the first boot of the ADC appliance.

```

> sh ns ip
      Ippaddress      Traffic Domain  Type
-----
1)    10.128.4.27      0               NetScaler IP
2)    10.160.0.71      0               SNIP
3)    10.128.0.40      0               VIP
Done
> sh vlan
1)    VLAN ID: 1
      Link-local IPv6 addr: fe80::4001:aff:fea0:47/64
      Interfaces : 0/1 1/1 LO/1
2)    VLAN ID: 10     VLAN Alias Name:
      Interfaces : 1/2
      IPs :
          10.128.4.27      Mask: 255.255.255.0
Done
> sh route
      Network      Netmask      Gateway/OwnedIP  VLAN  State  Traffic Domain  Type
-----
1)    0.0.0.0       0.0.0.0      10.128.4.1       0      UP     0               STATIC
2)    127.0.0.0     255.0.0.0    127.0.0.1        0      UP     0               PERMANENT
3)    10.128.0.0     255.255.255.0 10.128.0.40      0      UP     0               DIRECT
4)    10.128.4.0     255.255.255.0 10.128.4.27      0      UP     0               DIRECT
5)    10.160.0.0     255.255.240.0 10.160.0.71      0      UP     0               DIRECT
Done
> █

```

Method 2: Custom bootstrap by specifying the interfaces, IP addresses, and subnet masks

You specify the management, client-facing and server-facing interfaces along with their IP addresses and subnet mask.

Custom bootstrap examples for AWS

In the following example, you skip the default bootstrap and run a new bootstrap sequence for the Citrix ADC appliance. For the new bootstrap sequence, you specify the following details:

- **Management interface:** Interface - Eth1, NSIP - 172.31.52.88, and subnet mask - 255.255.240.0
- **Client facing interface:** Interface - Eth0, VIP - 172.31.5.155, and subnet mask - 255.255.240.0.
- **Server facing interface:** Interface - Eth2, SNIP - 172.31.76.177, and subnet mask - 255.255.240.0.

```
<NS-PRE-BOOT-CONFIG>
  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>
    <MGMT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth1 </INTERFACE-NUM>
      <IP>172.31.52.88 </IP>
      <SUBNET-MASK>255.255.240.0 </SUBNET-MASK>
    </MGMT-INTERFACE-CONFIG>
    <CLIENT-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth0 </INTERFACE-NUM>
      <IP>172.31.5.155 </IP>
      <SUBNET-MASK>255.255.240.0 </SUBNET-MASK>
    </CLIENT-INTERFACE-CONFIG>
    <SERVER-INTERFACE-CONFIG>
      <INTERFACE-NUM>eth2 </INTERFACE-NUM>
      <IP>172.31.76.177 </IP>
      <SUBNET-MASK>255.255.240.0 </SUBNET-MASK>
    </SERVER-INTERFACE-CONFIG>
  </NS-BOOTSTRAP>
</NS-PRE-BOOT-CONFIG>
```

You can run the `show ns ip` command in the ADC CLI, and verify that the new bootstrap sequence specified in the `<NS-BOOTSTRAP>` section is applied. You can run the “show route” command to verify the subnet mask.

```

> sh ns ip
-----
Ipaddress      Traffic Domain  Type           Mode   Arp   Icmp   Vserver  State
-----
1) 172.31.52.88   0               NetScaler IP   Active Enabled Enabled NA      Enabled
2) 172.31.76.177 0               SNIP          Passive Enabled Enabled NA      Enabled
3) 172.31.5.155  0               VIP           Passive Enabled Enabled Enabled Enabled
Done
> sh vlan
1)  VLAN ID: 1
   Link-local IPv6 addr: fe80::839:e2ff:feaf:4a9e/64
   Interfaces : 1/1 1/3 LO/1
2)  VLAN ID: 10   VLAN Alias Name:
   Interfaces : 1/2
   IPs :
       172.31.52.88      Mask: 255.255.240.0
Done
> sh route
-----
Network        Netmask        Gateway/OwnedIP  VLAN   State  Traffic Domain  Type
-----
1) 0.0.0.0      0.0.0.0        172.31.48.1     0      UP     0               STATIC
2) 127.0.0.0   255.0.0.0      127.0.0.1       0      UP     0               PERMANENT
3) 172.31.0.0   255.255.240.0  172.31.5.155    0      UP     0               DIRECT
4) 172.31.48.0  255.255.240.0  172.31.52.88    0      UP     0               DIRECT
5) 172.31.64.0  255.255.240.0  172.31.76.177   0      UP     0               DIRECT
6) 172.31.0.2   255.255.255.255 172.31.48.1     0      UP     0               STATIC
Done

```

Custom bootstrap example for Azure

In the following example, a new bootstrap sequence for ADC is mentioned and default bootstrap is skipped. You provide the interface details along with the IP addresses and subnet masks as follows:

- Management interface (eth2), NSIP (172.27.2.53), and subnet mask (255.255.255.0)
- Client facing interface (eth1), VIP (172.27.1.53), and subnet mask (255.255.255.0)
- Server facing interface (eth0), SNIP (172.27.0.53), and subnet mask (255.255.255.0)

```

<NS-PRE-BOOT-CONFIG>
  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>

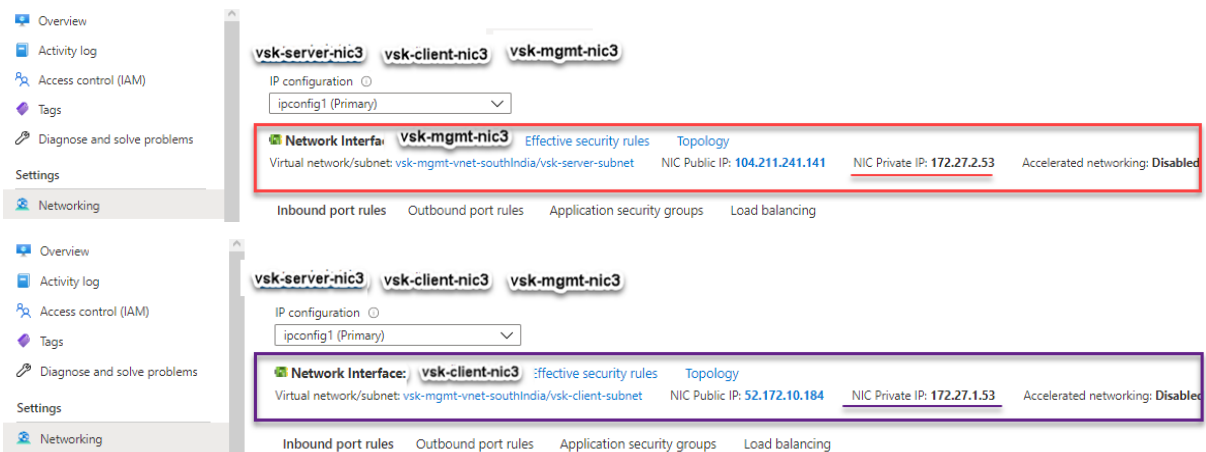
    <MGMT-INTERFACE-CONFIG>
      <INTERFACE-NUM> eth2 </INTERFACE-NUM>
      <IP> 172.27.2.53 </IP>
      <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
    </MGMT-INTERFACE-CONFIG>


    <CLIENT-INTERFACE-CONFIG>
      <INTERFACE-NUM> eth1 </INTERFACE-NUM>
      <IP> 172.27.1.53 </IP>
      <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
    </CLIENT-INTERFACE-CONFIG>

    <SERVER-INTERFACE-CONFIG>
      <INTERFACE-NUM> eth0 </INTERFACE-NUM>
      <IP> 172.27.0.53 </IP>
      <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
    </SERVER-INTERFACE-CONFIG>

  </NS-BOOTSTRAP>
</NS-PRE-BOOT-CONFIG>
  
```

You can see that the Citrix ADC VPX instance is created with three network interfaces. Navigate to the **Azure portal > VM instance > Networking**, and verify the networking properties of the three NICs as shown in the following illustrations.





You can run the `show ns ip` command in the ADC CLI, and verify that the new bootstrap sequence specified in the `<NS-BOOTSTRAP>` section is applied. You can run the “show route” command to verify the subnet mask.

```
> sh ns ip
-----
Ipaddress      Traffic Domain  Type           Mode   Arp   Icmp   Vserver  State
-----
1) 172.27.2.53  0              NetScaler IP  Active Enabled Enabled NA      Enabled
2) 172.27.0.53  0              SNIP          Active Enabled Enabled NA      Enabled
3) 172.27.1.53  0              VIP           Active Enabled Enabled Enabled Enabled
Done
> sh vlan
1) VLAN ID: 1
   Link-local IPv6 addr: fe80::20d:3aff:fec9:c26c/64
   Interfaces : 0/1 1/1 LO/1
2) VLAN ID: 10   VLAN Alias Name:
   Interfaces : 1/2
   IPs :
      172.27.2.53      Mask: 255.255.255.0
Done
> sh route
-----
Network      Netmask      Gateway/OwnedIP  VLAN  State  Traffic Domain  Type
-----
1) 0.0.0.0    0.0.0.0      172.27.2.1      0     UP     0               STATIC
2) 127.0.0.0  255.0.0.0    127.0.0.1      0     UP     0               PERMANENT
3) 172.27.0.0 255.255.255.0 172.27.0.53    0     UP     0               DIRECT
4) 172.27.1.0 255.255.255.0 172.27.1.53    0     UP     0               DIRECT
5) 172.27.2.0 255.255.255.0 172.27.2.53    0     UP     0               DIRECT
6) 169.254.0.0 255.255.0.0  172.27.0.1     0     UP     0               STATIC
7) 168.63.129.16 255.255.255.255 172.27.0.1 0     UP     0               STATIC
8) 169.254.169.254 255.255.255.255 172.27.0.1 0     UP     0               STATIC
Done
```

Custom bootstrap example for GCP

In the following example, a new bootstrap sequence for ADC is mentioned and default bootstrap is skipped. You provide the interface details along with the IP addresses and subnet masks as follows:

- Management interface (eth2), NSIP (10.128.4.31), and subnet mask (255.255.255.0)
- Client facing interface (eth1), VIP (10.128.0.43), and subnet mask (255.255.255.0)
- Server facing interface (eth0), SNIP (10.160.0.75), and subnet mask (255.255.255.0)

```

<NS-PRE-BOOT-CONFIG>
  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>

    <MGMT-INTERFACE-CONFIG>
      <INTERFACE-NUM> eth2 </INTERFACE-NUM>
      <IP> 10.128.4.31 </IP>
      <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
    </MGMT-INTERFACE-CONFIG>

    <CLIENT-INTERFACE-CONFIG>
      <INTERFACE-NUM> eth1 </INTERFACE-NUM>
      <IP> 10.128.0.43 </IP>
      <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
    </CLIENT-INTERFACE-CONFIG>

    <SERVER-INTERFACE-CONFIG>
      <INTERFACE-NUM> eth0 </INTERFACE-NUM>
      <IP> 10.160.0.75 </IP>
      <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
    </SERVER-INTERFACE-CONFIG>

  </NS-BOOTSTRAP>
</NS-PRE-BOOT-CONFIG>

```

After the VM instance is created in the GCP portal with the custom bootstrap, you can verify the network interface properties as follows:

1. Select the instance that you have created by providing the custom bootstrap information.
2. Navigate to the Network interface properties and verify the NIC details as follows.

Network interfaces								
Name	Network	Subnetwork	Primary internal IP	Alias IP ranges	External IP	Network Tier	IP forwarding	Network details
nic0	default	default	vsk-defnw-st-ip1 (10.160.0.75)	—	34.93.216.90 (ephemeral)	Premium	Off	View details
nic1	vsk-vpc-network-1	asia-south1-subnet-1	vsk-vpc-nw1-st-ip1 (10.128.0.43)	—	35.244.40.113 (ephemeral)	Premium		View details
nic2	vsk-vpc-network-2	asia-south1-subnet-5	vsk-nw2-st-ip-1 (10.128.4.31)	—	34.93.202.214 (ephemeral)	Premium		View details

You can run the `show nsip` command in the ADC CLI, and verify that the new bootstrap sequence specified in the `<NS-BOOTSTRAP>` section is applied. You can run the “show route” command to verify the subnet mask.

```

> sh ns ip
-----
Ipaddress      Traffic Domain  Type           Mode   Arp   Icmp   Vserver  State
-----
1) 10.128.4.31   0              NetScaler IP   Active Enabled Enabled NA      Enabled
2) 10.160.0.75   0              SNIP          Passive Enabled Enabled NA      Enabled
3) 10.128.0.43   0              VIP           Passive Enabled Enabled Enabled Enabled
Done
> sh vlan
1)  VLAN ID: 1
   Link-local IPv6 addr: fe80::4001:aff:fea0:4b/64
   Interfaces : 0/1 1/1 LO/1
2)  VLAN ID: 10   VLAN Alias Name:
   Interfaces : 1/2
   IPs :
      10.128.4.31      Mask: 255.255.255.0
Done
> sh route
-----
Network      Netmask      Gateway/OwnedIP  VLAN  State  Traffic Domain  Type
-----
1) 0.0.0.0    0.0.0.0      10.128.4.1      0      UP     0               STATIC
2) 127.0.0.0  255.0.0.0    127.0.0.1      0      UP     0               PERMANENT
3) 10.128.0.0  255.255.255.0  10.128.0.43    0      UP     0               DIRECT
4) 10.128.4.0  255.255.255.0  10.128.4.31    0      UP     0               DIRECT
5) 10.160.0.0  255.255.255.0  10.160.0.75    0      UP     0               DIRECT
Done
>

```

Method 3: Custom bootstrap by providing bootstrap related commands in the <NS-CONFIG> section

You can provide the bootstrap related commands in the <NS-CONFIG> section. In the <NS-BOOTSTRAP> section, you must specify the <NEW-BOOTSTRAP-SEQUENCE> as “No” to run the bootstrapping commands in the <NS-CONFIG> section. You must also provide the commands to assign NSIP, default route, and NSVLAN. In addition, provide the commands relevant for the cloud that you use.

Before providing a custom bootstrap, ensure that your cloud infrastructure supports a particular interface configuration.

Custom bootstrap example for AWS

In this example, bootstrap related commands are provided in the <NS-CONFIG> section. The <NS-BOOTSTRAP> section indicates that the default bootstrapping is skipped, and the custom bootstrap information provided in the <NS-CONFIG> section is run. You must also provide the commands to create NSIP, add default route, and add NSVLAN.


```

<NS-PRE-BOOT-CONFIG>
  <NS-CONFIG>
    set ns config -IPAddress 172.31.52.88 -netmask 255.255.240.0
    add route 0.0.0.0 0.0.0.0 172.31.48.1
    set ns config -nsvlan 10 -ifnum 1/2 -tagged NO
    add route 172.31.0.2 255.255.255.255 172.31.48.1

    enable ns feature WL SP LB RESPONDER
    add server 5.0.0.201 5.0.0.201
    add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -maxClient 0 -maxReq 0 -cip DISABLED -usip NO -
useproxyport YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -TCPB NO -CMP NO
    add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -persistenceType NONE -cltTimeout 180

  </NS-CONFIG>

  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE> NO </NEW-BOOTSTRAP-SEQUENCE>
  </NS-BOOTSTRAP>

</NS-PRE-BOOT-CONFIG>

```

Bootstrap related commands

route to DNS server is added through default gateway

You can copy the configuration shown in the preceding screenshot from here:

```

1 <NS-PRE-BOOT-CONFIG>
2   <NS-CONFIG>
3
4     set ns config -IPAddress 172.31.52.88 -netmask 255.255.240.0
5     add route 0.0.0.0 0.0.0.0 172.31.48.1
6     set ns config -nsvlan 10 -ifnum 1/2 -tagged NO
7     add route 172.31.0.2 255.255.255.255 172.31.48.1
8
9     enable ns feature WL SP LB RESPONDER
10    add server 5.0.0.201 5.0.0.201
11    add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -
maxClient 0 -maxReq 0 -cip DISABLED -usip NO - useproxyport
YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -TCPB NO
-CMP NO
12    add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -
persistenceType NONE -cltTimeout 180
13
14  </NS-CONFIG>
15
16  <NS-BOOTSTRAP>
17    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
18    <NEW-BOOTSTRAP-SEQUENCE> NO </NEW-BOOTSTRAP-SEQUENCE>
19  </NS-BOOTSTRAP>
20
21
22 </NS-PRE-BOOT-CONFIG>

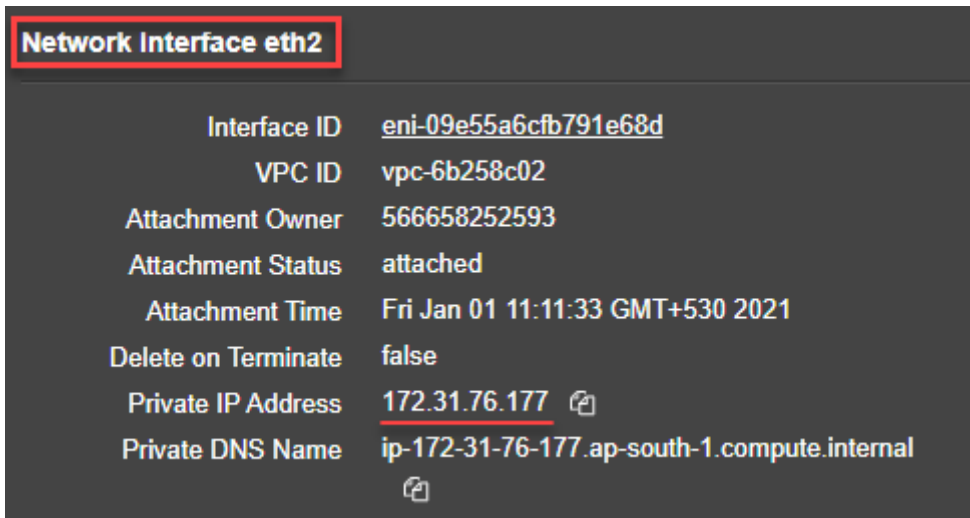
```

After the VM instance is created, in the AWS portal, you can verify the network interface properties as follows:

1. Navigate to the **AWS Portal > EC2 instances**, and select the instance that you have created by providing the custom bootstrap information.
2. In the **Description** tab, you can verify the properties of each network interface as shown in the following illustrations.

Network Interface eth1	
Interface ID	eni-021961099be6815eb
VPC ID	vpc-6b258c02
Attachment Owner	566658252593
Attachment Status	attached
Attachment Time	Fri Jan 01 11:11:23 GMT+530 2021
Delete on Terminate	false
Private IP Address	172.31.52.88
Private DNS Name	ip-172-31-52-88.ap-south-1.compute.internal

Network Interface eth0	
Interface ID	eni-039e5f3329cd879e9
VPC ID	vpc-6b258c02
Attachment Owner	566658252593
Attachment Status	attached
Attachment Time	Fri Jan 01 10:58:28 GMT+530 2021
Delete on Terminate	true
Private IP Address	172.31.5.155
Private DNS Name	ip-172-31-5-155.ap-south-1.compute.internal



You can run the `show ns ip` command in **ADC CLI**, and verify the network interfaces applied to the ADC VPX instance during the first boot of the ADC appliance.

```
> sh ns ip
-----
1) 172.31.52.88 0 NetScaler IP Active Enabled Enabled NA Enabled
2) 4.0.0.101 0 VIP Active Enabled Enabled Enabled Enabled
Done
> sh vlan
1) VLAN ID: 1
   Link-local IPv6 addr: fe80::839:e2ff:feaf:4a9e/64
   Interfaces : 1/1 1/3 LO/1
2) VLAN ID: 10 VLAN Alias Name:
   Interfaces : 1/2
   IPs :
       172.31.52.88 Mask: 255.255.240.0
Done
> sh route
-----
1) Network Netmask Gateway/OwnedIP VLAN State Traffic Domain Type
2) 0.0.0.0 0.0.0.0 172.31.48.1 0 UP 0 STATIC
3) 127.0.0.0 255.0.0.0 127.0.0.1 0 UP 0 PERMANENT
4) 172.31.48.0 255.255.240.0 172.31.52.88 0 UP 0 DIRECT
5) 172.31.0.2 255.255.255.255 172.31.48.1 0 UP 0 STATIC
Done
>
```

Custom bootstrap example for Azure

In this example, bootstrap related commands are provided in the `<NS-CONFIG>` section. The `<NS-BOOTSTRAP>` section indicates that the default bootstrapping is skipped, and the custom bootstrap information provided in the `<NS-CONFIG>` section is run.

Note:

For Azure cloud, Instance Metadata Server (IMDS) and DNS servers are accessible only through primary interface (Eth0). Therefore, if Eth0 interface is not used as management interface (NSIP),

Eth0 interface must at least be configured as SNIP for IMDS or DNS access to work. The route to IMDS endpoint (169.254.169.254) and DNS endpoint (168.63.129.16) through Eth0's gateway must also be added.

```

<NS-PRE-BOOT-CONFIG>
  <NS-CONFIG>
    set ns config -IPAddress 172.27.2.61 -netmask 255.255.255.0
    add route 0.0.0.0 0.0.0.0 172.27.2.1
    set ns config -nsvlan 10 -ifnum 1/2 -tagged NO
    add ns ip 172.27.0.61 255.255.255.0 -type SNIP
    add route 169.254.169.254 255.255.255.255 172.27.0.1
    add route 168.63.129.16 255.255.255.255 172.27.0.1

    add vlan 5
    bind vlan 5 -IPAddress 5.0.0.1 255.255.255.0
    enable ns feature WL SP LB RESPONDER
    add server 5.0.0.201 5.0.0.201
    add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -maxClient 0 -maxReq 0 -cip DISABLED -usip
    NO -useproxyport YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -TCPB NO -CMP NO
    add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -persistenceType NONE -cltTimeout 180

  </NS-CONFIG>

  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE> NO </NEW-BOOTSTRAP-SEQUENCE>

  </NS-BOOTSTRAP>
  
```

```

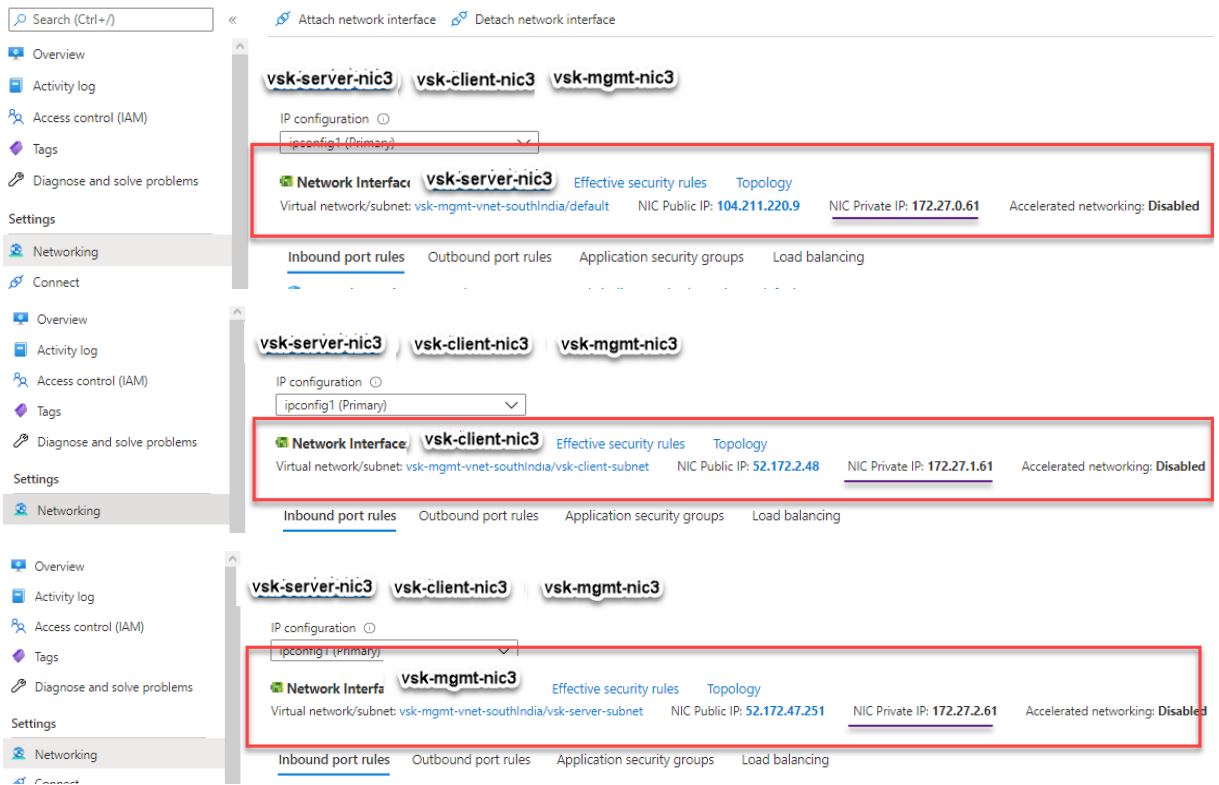
1 <NS-PRE-BOOT-CONFIG>
2
3 <NS-CONFIG>
4
5   set ns config -IPAddress 172.27.2.61 -netmask 255.255.255.0
6   add route 0.0.0.0 0.0.0.0 172.27.2.1
7   set ns config -nsvlan 10 -ifnum 1/2 -tagged NO
8   add ns ip 172.27.0.61 255.255.255.0 -type SNIP
9   add route 169.254.169.254 255.255.255.255 172.27.0.1
10  add route 168.63.129.16 255.255.255.255 172.27.0.1
11
12  add vlan 5
13  bind vlan 5 -IPAddress 5.0.0.1 255.255.255.0
14  enable ns feature WL SP LB RESPONDER
15  add server 5.0.0.201 5.0.0.201
16  add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -
      maxClient 0 -maxReq 0 -cip DISABLED -usip NO -useproxyport
      YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -TCPB NO
  
```

```

-CMP NO
17     add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -
      persistenceType NONE -cltTimeout 180
18
19     </NS-CONFIG>
20
21     <NS-BOOTSTRAP>
22
23     <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
24     <NEW-BOOTSTRAP-SEQUENCE> NO </NEW-BOOTSTRAP-SEQUENCE>
25
26     </NS-BOOTSTRAP>
27
28 </NS-PRE-BOOT-CONFIG>

```

You can see that the Citrix ADC VPX instance is created with three network interfaces. Navigate to the **Azure portal > VM instance > Networking**, and verify the networking properties of the three NICs as shown in the following illustrations.



You can run the `show ns ip` command in the ADC CLI, and verify that the new bootstrap sequence specified in the `<NS-BOOTSTRAP>` section is applied. You can run the “show route” command to verify the subnet mask.

```

> sh ns ip
-----
Ipaddress      Traffic Domain  Type           Mode   Arp   Icmp   Vserver  State
-----
1) 172.27.2.61   0              NetScaler IP   Active Enabled Enabled NA      Enabled
2) 172.27.0.61   0              SNIP           Active Enabled Enabled NA      Enabled
3) 4.0.0.101     0              VIP            Active Enabled Enabled Enabled Enabled
Done
> sh vlan
1) VLAN ID: 1
   Link-local IPv6 addr: fe80::20d:3aff:fec9:9076/64
   Interfaces : 0/1 1/1 LO/1
2) VLAN ID: 5    VLAN Alias Name:
3) VLAN ID: 10  VLAN Alias Name:
   Interfaces : 1/2
   IPs :
     172.27.2.61      Mask: 255.255.255.0
Done
> sh route
-----
Network      Netmask      Gateway/OwnedIP  VLAN  State  Traffic Domain  Type
-----
1) 0.0.0.0    0.0.0.0      172.27.2.1      0     UP     0              STATIC
2) 127.0.0.0  255.0.0.0    127.0.0.1      0     UP     0              PERMANENT
3) 172.27.0.0 255.255.255.0 172.27.0.61    0     UP     0              DIRECT
4) 172.27.2.0 255.255.255.0 172.27.2.61    0     UP     0              DIRECT
5) 169.254.0.0 255.255.0.0  172.27.0.1     0     UP     0              STATIC
6) 168.63.129.16 255.255.255.255 172.27.0.1  0     UP     0              STATIC
7) 169.254.169.254 255.255.255.255 172.27.0.1  0     UP     0              STATIC
Done

```

Custom bootstrap example for GCP

In this example, bootstrap related commands are provided in the <NS-CONFIG> section. The <NS-BOOTSTRAP> section indicates that the default bootstrapping is skipped, and the custom bootstrap information provided in the <NS-CONFIG> section is applied.

```

<NS-PRE-BOOT-CONFIG>

  <NS-CONFIG>
    set ns config -IPAddress 10.128.0.2 -netmask 255.255.255.0
    add route 0.0.0.0 0.0.0.0 10.128.0.1
    set ns config -nsvlan 10 -ifnum 1/1 -tagged NO

    enable ns feature WL SP LB RESPONDER
    add server 5.0.0.201 5.0.0.201
    add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -maxClient 0 -maxReq 0 -cip
    DISABLED -usip NO -useproxyport YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -TCPB NO -CMP NO
    add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -persistenceType NONE -cltTimeout 180

  </NS-CONFIG>

  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE> NO </NEW-BOOTSTRAP-SEQUENCE>
  </NS-BOOTSTRAP>

</NS-PRE-BOOT-CONFIG>

```

You can copy the configuration shown in the preceding screenshot from here:

```

1 <NS-PRE-BOOT-CONFIG>
2
3   <NS-CONFIG>
4
5     set ns config -IPAddress 10.128.0.2 -netmask 255.255.255.0
6     add route 0.0.0.0 0.0.0.0 10.128.0.1
7     set ns config -nsvlan 10 -ifnum 1/1 -tagged NO
8
9     enable ns feature WL SP LB RESPONDER
10    add server 5.0.0.201 5.0.0.201
11    add service preboot_s5_201 5.0.0.201 HTTP 80 -gslb NONE -
12      maxClient 0 -maxReq 0 -cip DISABLED -usip NO -useproxyport
13      YES -sp OFF -cltTimeout 180 -svrTimeout 360 -CKA NO -TCPB NO
14      -CMP NO
15    add lb vserver preboot_v4_101 HTTP 4.0.0.101 80 -
16      persistenceType NONE -cltTimeout 180
17
18   </NS-CONFIG>
19
20   <NS-BOOTSTRAP>
21     <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
22     <NEW-BOOTSTRAP-SEQUENCE> NO </NEW-BOOTSTRAP-SEQUENCE>
23   </NS-BOOTSTRAP>
24
25 </NS-PRE-BOOT-CONFIG>

```

After the VM instance is created in the GCP portal with the custom bootstrap, you can verify the network interface properties as follows:

1. Select the instance that you have created by providing the custom bootstrap information.
2. Navigate to the Network interface properties and verify the NIC details as shown in the illustration.

Network interfaces					
Name	Network	Subnetwork	Primary internal IP	Alias IP ranges	External IP
nic0	default	default	10.160.0.74	–	34.93.9.79 (ephemeral)
nic1	vsk-vpc-network-1	asia-south1-subnet-1	asia-south1-subnet1-10-128-0-2 (10.128.0.2)	–	34.93.245.110 (ephemeral)
nic2	vsk-vpc-network-2	asia-south1-subnet-5	10.128.4.30	–	34.93.146.248 (ephemeral)

You can run the `show nsip` command in **ADC CLI**, and verify that the configurations provided in the preceding `<NS-CONFIG>` section are applied at the first boot of the ADC appliance.

```
1 ! [Show NSIP output] (/en-us/vpx/media/gcp-show-nsip-method3.png)
```

Impact of attaching and detaching NICs in AWS and Azure

AWS and Azure provide the option to attach a network interface to an instance, and detach a network interface from an instance. Attaching or detaching interfaces might alter interface positions. Hence, Citrix recommends you to refrain from detaching interfaces from the ADC VPX instance. If you detach or attach an interface when custom bootstrapping is configured, Citrix ADC VPX instance reassigns the primary IP of the newly available interface in the management interface’s position as NSIP. If no further interfaces are available after the one you detached, then the first interface is made the management interface for the ADC VPX instance.

For example, a Citrix ADC VPX instance is brought up with 3 interfaces: Eth0 (SNIP), Eth1 (NSIP), and Eth2 (VIP). If you detach Eth1 interface from the instance, which is a management interface, ADC configures the next available interface (Eth2) as the management interface. Thereby, the ADC VPX instance is still accessed through the primary IP of Eth2 interface. If Eth2 is also not available, then the remaining interface (Eth0) is made the management interface. Therefore, the access to ADC VPX instance continues to exist.

Let’s consider a different assignment of interfaces as follows: Eth0 (SNIP), Eth1 (VIP), and Eth2 (NSIP). If you detach Eth2 (NSIP), because no new interface is available after Eth2, the first interface (Eth0) is made the management interface.

Install a Citrix ADC VPX instance on a bare metal server

September 12, 2024

A bare metal is a fully dedicated physical server that delivers physical isolation, fully integrated into the cloud environment. It is also known as a single-tenant server. Single tenancy allows you to avoid the noisy neighbor effect. With bare metal, you do not witness the noisy neighbor effect because you are the sole user.

A bare metal server installed with a hypervisor provides you a management suite to create virtual machines on the server. The hypervisor does not run applications natively. Its purpose is to virtualize your workloads into separate virtual machines to gain the flexibility and reliability of virtualization.

Prerequisites for installing Citrix ADC VPX instance on bare metal servers

A bare metal server must be obtained from a cloud vendor that meets all the system requirements for the respective hypervisor.

Install the Citrix ADC VPX instance on bare metal servers

To install Citrix ADC VPX instances on a bare metal server, you must first obtain a bare metal server with adequate system resources from a cloud vendor. On that bare metal server, any of the supported hypervisors such as Linux KVM, VMware ESX, Citrix Hypervisor, or Microsoft Hyper-V must be installed and configured before deploying the ADC VPX instance.

For more information on the list of different hypervisors and features supported on a Citrix ADC VPX instance, see [Support matrix and usage guidelines](#).

For more information on installing Citrix ADC VPX instances on different hypervisors, see the respective documentation.

- **Citrix Hypervisor:** See [Install a Citrix ADC VPX instance on Citrix Hypervisor](#).
- **VMware ESX:** See [Install a Citrix ADC VPX instance on VMware ESX](#).
- **Microsoft Hyper-V:** See [Install a Citrix ADC VPX instance on Microsoft Hyper-V server](#).
- **Linux KVM platform:** See [Install a Citrix ADC VPX instance on Linux-KVM platform](#).

Install a Citrix ADC VPX instance on Citrix Hypervisor/XenServer


December 17, 2024

To install VPX instances on the Citrix Hypervisor/XenServer, you must first install the Hypervisor on a machine with adequate system resources. To perform the Citrix ADC VPX instance installation, you use Citrix XenCenter, which must be installed on a remote machine that can connect to the Hypervisor host through the network.

For more information about Hypervisor, see [Citrix Hypervisor documentation](#).

The following figure shows the bare-metal solution architecture of Citrix ADC VPX instance on Hypervisor.

Figure. A Citrix ADC VPX instance on Citrix Hypervisor/XenServer



Prerequisites for installing a Citrix ADC VPX instance on Hypervisor

Before you begin installing a virtual appliance, do the following:

- Install Hypervisor version 6.0 or later on hardware that meets the minimum requirements.
- Install XenCenter on a management workstation that meets the minimum system requirements.
- Obtain virtual appliance license files. For more information about virtual appliance licenses, see the [Citrix ADC Licensing Guide](#).

Hypervisor hardware requirements

The following table describes the minimum hardware requirements for a Hypervisor platform running a Citrix ADC VPX instance.

Table 1. Minimum system requirements for Hypervisor running a nCore VPX instance

Component	Requirement
CPU	2 or more 64-bit x86 CPUs with virtualization assist (Intel-VT) enabled. AMD processor is not supported. To run the Citrix ADC VPX instance, hardware support for virtualization must be enabled on the Hypervisor host. Make sure that the BIOS option for virtualization support is not disabled. For more details, see BIOS documentation.
RAM	3 GB
Disk space	Locally attached storage (PATA, SATA, SCSI) with 40 GB of disk space.
NIC	<p>Note: Hypervisor installation creates a 4 GB partition for the Hypervisor host control domain. One 1-Gbps NIC is recommended; two 1-Gbps NICs are recommended for high availability.</p> <p>The remaining space is available for NetScaler VPX instance and other virtual machines.</p>

For more information, see the [XenServer documentation](#).

The following table lists the virtual computing resources that Hypervisor must provide for each nCore VPX virtual appliance.

Table 2. Minimum virtual computing resources required for running a nCore VPX instance

Component	Requirement
Memory	2 GB
Virtual CPU (vCPU)	2
Virtual network interfaces	2

Note:

For production use of Citrix ADC VPX instance, Citrix recommends that CPU priority (in virtual machine properties) be set to the highest level, to improve scheduling behavior and network

latency.

XenCenter system requirements

XenCenter is a Windows client application. It cannot run on the same machine as the Hypervisor host. For more information about minimum system requirements and installing XenCenter, see the following Hypervisor documents:

- [System requirements](#)
- [Install](#)

Install Citrix ADC VPX instances on Hypervisor by using XenCenter

After you have installed and configured Hypervisor and XenCenter, you can use XenCenter to install virtual appliances on Hypervisor. The number of virtual appliances that you can install depends on the amount of memory available on the hardware that is running Hypervisor.

To install Citrix ADC VPX instances on Hypervisor by using XenCenter, follow these steps:

1. Start XenCenter on your workstation.
2. On the Server menu, click **Add**.
3. In the Add New Server dialog box, in the host name text box, type the IP address or DNS name of the Hypervisor that you want to connect to.
4. In the User Name and Password text boxes, type the administrator credentials, and then click Connect. The Hypervisor name appears in the navigation pane with a green circle, which indicates that the Hypervisor is connected.
5. In the navigation pane, click the name of the Hypervisor on which you want to install the Citrix ADC VPX instance.
6. On the VM menu, click **Import**.
7. In the Import dialog box, in the Import file name, browse to the location at which you saved the Citrix ADC VPX instance .xva image file. Make sure that the Exported VM option is selected, and then click **Next**.
8. Select the Hypervisor on which you want to install the virtual appliance, and then click **Next**.
9. Select the local storage repository in which to store the virtual appliance, and then click **Import** to begin the import process.
10. You can add, modify, or delete the virtual network interfaces as required. When finished, click **Next**.

11. Click **Finish** to complete the import process.

Note:

To view the status of the import process, click the **Log** tab.

12. If you want to install another virtual appliance, repeat steps 5 through 11.

Note:

After the initial configuration of the VPX instance, if you want to upgrade the appliance to the latest software release, see [Upgrading or Downgrading the System Software](#).

Configure VPX instances to use single root I/O virtualization (SR-IOV) network interfaces

January 27, 2025

After you have installed and configured a Citrix ADC VPX instance on Citrix Hypervisor, you can configure the virtual appliance to use SR-IOV network interfaces.

The following NICs are supported:

- Intel 82599 10G
- Intel X710 10G
- Intel XL710 40G

Limitations

Citrix Hypervisor does not support some features on SR-IOV interfaces. The limitations with Intel 82599, Intel X710, and Intel XL710 NICs are listed in the following sections.

Limitations for Intel 82599 NIC

Intel 82599 NIC does not support the following features:

- L2 mode switching
- Clustering
- Admin partitioning [Shared VLAN mode]
- High Availability [Active - Active mode]
- Jumbo frames
- IPv6 protocol in Cluster environment

Limitations for Intel X710 10G and Intel XL710 40G NICs

Intel X710 10G and Intel XL710 40G NICs have the following limitations:

- L2 mode switching is not supported.
- Admin partitioning (shared VLAN mode) is not supported.
- In a cluster, Jumbo frames are not supported when the XL710 NIC is used as a data interface.
- Interface list reorders when interfaces are disconnected and reconnected.
- Interface parameter configurations such as speed, duplex, and auto negotiations are not supported.
- For both Intel X710 10G and Intel XL710 40G NICs, the interface comes up as 40/x interface.
- Up to only 16 Intel X710/XL710 SR-IOV interfaces can be supported on a VPX instance.

Note:

For Intel X710 10G and Intel XL710 40G NICs to support IPv6, enable trust mode on the virtual functions (VFs) by typing the following command on the Citrix Hypervisor host:

```
# ip link set <PNIC> <VF> trust on
```

Example:

```
# ip link set ens785f1 vf 0 trust on
```

Prerequisites for Intel 82599 NIC

On the Citrix Hypervisor host, ensure that you:

- Add the Intel 82599 NIC (NIC) to the host.
- Block list the `ixgbevf` driver by adding the following entry to the `/etc/modprobe.d/blacklist.conf` file:

```
blacklist ixgbevf
```

- Enable SR-IOV Virtual Functions (VFs) by adding the following entry to the `/etc/modprobe.d/ixgbe` file:

```
options ixgbe max_vfs=<number_of_VFs>
```

where `<number_VFs>` is the number of SR-IOV VFs that you want to create.

- Verify that SR-IOV is enabled in BIOS.

Note:

IXGBE driver version 3.22.3 is recommended.

Assign Intel 82599 SR-IOV VFs to the Citrix ADC VPX instance by using the Citrix Hypervisor host

To assign an Intel 82599 SR-IOV VFs to Citrix ADC VPX instance, follow these steps:

1. On the Citrix Hypervisor host, use the following command to assign the SR-IOV VFs to the Citrix ADC VPX instance:

```
xe host-call-plugin plugin=iovirt host-uuid=<Xen host UUID> fn=assign_free_vf args:uuid=<NetScaler VM UUID> args:ethdev=<interface name> args:mac=<Mac addr>
```

Where:

- <Xen host UUID> is the UUID of the Citrix Hypervisor host.
- <NetScaler VM UUID> is the UUID of the Citrix ADC VPX instance.
- <interface name> is the interface for the SR-IOV VFs.
- <MAC address > is the MAC address of the SR-IOV VF.

Note:

Specify the MAC address that you want use in the args:Mac= parameter, if not specified, the `iovirt` script randomly generates and assigns a MAC address. Also, if you want to use the SR-IOV VFs in Link Aggregation mode, make sure that you specify the MAC address as 00:00:00:00:00:00.

2. Boot the Citrix ADC VPX instance.

Unassign Intel 82599 SR-IOV VFs to the ADC VPX instance by using the Citrix Hypervisor host

If you have assigned an incorrect SR-IOV VFs or if you want to modify an assigned SR-IOV VFs, you need to unassign and reassign the SR-IOV VFs to the Citrix ADC VPX instance.

To unassign SR-IOV network interface assigned to a Citrix ADC VPX instance, follow these steps:

1. On the Citrix Hypervisor host, use the following command to assign the SR-IOV VFs to the Citrix ADC VPX instance and reboot the Citrix ADC VPX instance:

```
xe host-call-plugin plugin=iovirt host-uuid=<Xen_host_UUID> fn=unassign_all args:uuid=<Netscaler_VM_UUID>
```

Where:

- <Xen_host_UUID> - The UUID of the Citrix Hypervisor host.
- <Netscaler_VM_UUID> - The UUID of the Citrix ADC VPX instance

2. Boot the Citrix ADC VPX instance.

Assign Intel X710/XL710 SR-IOV VFs to the Citrix ADC VPX instance by using the Citrix Hypervisor host

To assign an Intel X710/XL710 SR-IOV VF to the Citrix ADC VPX instance, follow these steps:

1. Run the following command on the Citrix Hypervisor host to create a network.

```
1 xe network-create name=label=<network-name>
```

Example:

```
1 xe network-create name=label=SR-IOV-NIC-18 8ee59b73-7319-6998-cd69
  -b9fa3e8d7503
```

2. Determine the PIF Universal Unique Identifier (UUID) of the NIC on which the SR-IOV network is to be configured.

```
1 xe pif-list
2
3         uuid ( R0) : e2874343-f1de-1fa7-8fef-98547c348783
4         device ( R0): eth18
5 currently-attached ( R0): true
6         VLAN ( R0): -1
7         network-uuid ( R0): f865bd85-44dd-b865-ab65-dcd6ae28c16e
```

3. Configure the network as an SR-IOV network. The following command also returns the UUID of the newly created SR-IOV network:

```
1 xe network-sriov-create network-uuid=<network-uuid> pif-uuid=<
  physical-pif-uuid>
```

Example:

```
1 xe network-sriov-create network-uuid=8ee59b73-7319-6998-cd69-
  b9fa3e8d7503 pif-uuid=e2874343-f1de-1fa7-8fef-98547
  c3487831629b44f-832a-084e-d67d-5d6d314d5e0f
```

To get more information on the SR-IOV network parameters, run the following command:

```
1 [root@citrix-XS82-TOP0 ~]# xe network-sriov-param-list uuid=1629
  b44f-832a-084e-d67d-5d6d314d5e0f
2
3         uuid ( R0): 1629b44f-832a-084e-d67d-5d6d314d5e0f
4         physical-PIF ( R0): e2874343-f1de-1fa7-8fef-98547c348783
5         logical-PIF ( R0): 85d52771-5814-c62d-45fa-f37b536144ff
6         requires-reboot ( R0): false
7         remaining-capacity ( R0): 32
```

4. Create a virtual interface (VIF) and attach it to the target VM.


```
1 xe vif-create device=0 mac=b2:61:fc:ae:00:1d network-uuid=8ee59b73
   -7319-6998-cd69-b9fa3e8d7503 vm-uuid=b507e8a6-f5ca-18eb-561d
   -308218a9dd68
2 3e1e2e58-b2ad-6dc0-61d4-1d149c9c6466
```

Note:

The NIC index number of the VM must start with 0.

Use the following command to find the VM UUID:

```
1 [root@citrix-XS82-TOP0 ~]# xe vm-list
2 uuid ( R0): b507e8a6-f5ca-18eb-561d-308218a9dd68
3 name-label ( RW): sai-vmx-1
4 power-state ( R0): halted
```

Remove Intel X710/XL710 SR-IOV VFs from the Citrix ADC instance by using the Citrix Hypervisor host

To remove an Intel X710/XL710 SR-IOV VF from a Citrix ADC VPX instance, follow these steps:

1. Copy the UUID for the VIF that you want to destroy.
2. Run the following command on the Citrix Hypervisor host to destroy the VIF.

```
1 xe vif-destroy uuid=<vif-uuid>
```

Example:

```
1 [root@citrix-XS82-TOP0 ~]# xe vif-destroy uuid=3e1e2e58-b2ad-6dc0
   -61d4-1d149c9c6466
```

Configure link aggregation on the SR-IOV interface

To use the SR-IOV virtual functions in link aggregation mode, you need to disable spoof checking for virtual functions that you have created. On the Citrix Hypervisor host, use the following command to disable spoof checking:

```
ip link set <interface_name> vf <VF_id> spoofchk off
```

Where:

- <interface_name> is the interface name.
- <VF_id> is the virtual function ID.

After disabling spoof checking for all the virtual functions that you have created, restart the Citrix ADC VPX instance and configure link aggregation. For instructions, see [Configure link aggregation](#).

Important:

While you are assigning the SR-IOV virtual functions (VFs) to the Citrix ADC VPX instance, make sure that you specify the MAC address 00:00:00:00:00:00 for the VFs.

Configure VLAN on the SR-IOV interface

You can configure VLAN on the SR-IOV virtual functions. For instructions, see [Configuring a VLAN](#).

Important:

Make sure that the Citrix Hypervisor host does not contain VLAN settings for the VF interface.

Other references

[SR-IOV enabled NICs](#)

[Add an SR-IOV Network](#)

Install a Citrix ADC VPX instance on VMware ESX

September 18, 2024

Before installing Citrix ADC VPX instances on VMware ESX, make sure that VMware ESX Server is installed on a machine with adequate system resources. To install a Citrix ADC VPX instance on VMware ESXi, you use the VMware vSphere client. The client or tool must be installed on a remote machine that can connect to VMware ESX through the network.

This section includes the following topics:

- Prerequisites
- Installing a Citrix ADC VPX instance on VMware ESX

Important:

You cannot install standard VMware Tools or upgrade the VMware Tools version available on a Citrix ADC VPX instance. VMware Tools for a Citrix ADC VPX instance are delivered as part of the Citrix ADC software release.

Prerequisites

Before you begin installing a virtual appliance, do the following:

- Install VMware ESX on hardware that meets the minimum requirements.
- Install VMware Client on a management workstation that meets the minimum system requirements.
- Download the Citrix ADC VPX appliance setup files.
- Create a virtual switch and attach the physical NIC to the virtual switch.
- Add port group and attach to the virtual switch.
- Attach the port group to the VM.
- Obtain the VPX license files. For more information about Citrix ADC VPX instance licenses, see [Licensing overview](#).

VMware ESX hardware requirements

The following table describes the minimum system requirements for VMware ESX servers running Citrix ADC VPX nCore virtual appliance.

Table 1. Minimum system requirements for a VMware ESX server running a Citrix ADC VPX instance

Component	Requirement
-	2 or more 64-bit x86 CPUs with virtualization assist (Intel-VT) enabled. To run Citrix ADC VPX instance, hardware support for virtualization must be enabled on the VMware ESX host. Make sure that the BIOS option for virtualization support is not disabled. For more information, see your BIOS documentation.
RAM	2 GB VPX. For critical deployments, we do not recommend 2 GB RAM for VPX because the system operates in a memory-constrained environment. This might lead to scale, performance, or stability related issues. Recommended is 4 GB RAM or 8 GB RAM.
Disk space	20 GB more than the minimum server requirements from VMware for setting up ESXi. See VMware documentation for minimum server requirements.
Network	One 1-Gbps NIC (NIC); Two 1-Gbps NICs recommended

For information about installing VMware ESX, see <http://www.vmware.com/>.

To enable SR-IOV or PCI passthrough support, ensure that the following processors are supported:

- Intel processors support Intel-VT.
- I/O Memory Management Unit (IOMMU) or SR-IOV is enabled in BIOS.

The following table lists the virtual computing resources that the VMware ESX server must provide for each VPX nCore virtual appliance.

Table 2. Minimum virtual computing resources required for running a Citrix ADC VPX instance

Component	Requirement
Memory	4 GB
Virtual CPU (vCPU)	2
Virtual network interfaces	1. In ESX, you can install a maximum of 10 virtual network interfaces if the VPX hardware is upgraded to version 7 or higher.
Disk space	20 GB

Note:

This is in addition to any disk requirements for the hypervisor.

For production use of VPX virtual appliance, the full memory allocation must be reserved. CPU cycles (in MHz) equal to at least the speed of one CPU core of the ESX must be reserved.

VMware vSphere client system requirements

VMware vSphere is a client application that can run on Windows and Linux operating systems. It cannot run on the same machine as the VMware ESX server. The following table describes the minimum system requirements.

Table 3. Minimum system requirements for VMware vSphere client installation

Component	Requirement
Operating system	For detailed requirements from VMware, search for the “vSphere Compatibility Matrixes” PDF file at http://kb.vmware.com/ .

Component	Requirement
CPU	750 MHz; 1 gigahertz (GHz) or faster recommended
RAM	1 GB; 2 GB recommended
NIC (NIC)	100 Mbps or faster NIC

OVF Tool 1.0 system requirements

OVF Tool is a client application that can run on Windows and Linux systems. It cannot run on the same machine as the VMware ESX server. The following table describes the minimum system requirements.

Table 4. Minimum system requirements for OVF tool installation

Component	Requirement
Operating system	For detailed requirements from VMware, search for the “OVF Tool User Guide” PDF file at http://kb.vmware.com/ .
CPU	750 MHz minimum, 1 GHz or faster recommended
RAM	1 GB Minimum, 2 GB recommended
NIC (NIC)	100 Mbps or faster NIC

For information about installing OVF, search for the “OVF Tool User Guide” PDF file at <http://kb.vmware.com/>.

Downloading the Citrix ADC VPX setup files

The Citrix ADC VPX instance setup package for VMware ESX follows the Open Virtual Machine (OVF) format standard. You can download the files from the Citrix website. You need a Citrix account to log on. If you do not have a Citrix account, access the home page at <http://www.citrix.com>, click the **New Users link**, and follow the instructions to create a new Citrix account.

Once logged on, navigate the following path from the Citrix home page:

Citrix.com > **Downloads > Citrix ADC > Virtual Appliances.**

Copy the following files to a workstation on the same network as the ESX server. Copy all three files into the same folder.

- NSVPX-ESX-<release number>-<build number>-disk1.vmdk (for example, NSVPX-ESX-13.0-71.44_nc_64-disk1.vmdk)
- NSVPX-ESX-<release number>-<build number>.ovf (for example, NSVPX-ESX-13.0-71.44_nc_64.ovf)
- NSVPX-ESX-<release number>-<build number>.mf (for example, NSVPX-ESX-13.0-71.44_nc_64.mf)

Install a Citrix ADC VPX instance on VMware ESX

After you have installed and configured VMware ESX, you can use the VMware vSphere client to install virtual appliances on the VMware ESX server. The number of virtual appliances that you can install depends on the amount of memory available on the hardware that is running VMware ESX.

To install Citrix ADC VPX instances on VMware ESX by using VMware vSphere Client, follow these steps:

1. Start the VMware vSphere client on your workstation.
2. In the **IP address / Name** text box, type the IP address of the VMware ESX server that you want to connect to.
3. In the **User Name** and **Password** text boxes, type the administrator credentials, and then click **Login**.
4. On the **File** menu, click **Deploy OVF Template**.
5. In the **Deploy OVF Template** dialog box, in **Deploy from file**, browse to the location at which you saved the Citrix ADC VPX instance setup files, select the .ovf file, and click **Next**.
6. Map the networks shown in the virtual appliance OVF template to the networks that you configured on the ESX host. Click **Next** to start installing a virtual appliance on VMware ESX. When installation is complete, a pop-up window informs you of the successful installation.
7. You are now ready to start the Citrix ADC VPX instance. In the navigation pane, select the Citrix ADC VPX instance that you have installed, and from the right-click menu, select **Power On**.
8. After the VM is booted, from the console, configure the Citrix ADC IP, Netmask and Gateway addresses. When you complete the configuration, select the **Save and Quit** option in the console.
9. To install another virtual appliance, repeat from Step 6 through step 8.

Note:

By default, the Citrix ADC VPX instance uses E1000 network interfaces.

After the installation, you can use vSphere client or vSphere Web Client to manage virtual appliances on VMware ESX.

To enable VLAN tagging on VMware ESX, configure the port group's VLAN ID to All (4095) on the vSwitch. For detailed instructions on setting a VLAN ID on the vSwitch, refer to the VMware documentation.

Migrate a Citrix ADC VPX instance by using VMware vMotion

You can migrate a Citrix ADC VPX instance by using VMware vSphere vMotion.

Follow these usage guidelines:

- VMware does not support the vMotion feature on virtual machines configured with PCI Passthrough and SR-IOV interfaces.
- Supported interfaces are E1000 and VMXNET3. To use vMotion on your VPX instance, ensure that the instance is configured with a supported interface.
- For more information about how to migrate an instance by using VMware vMotion, see the VMware documentation.

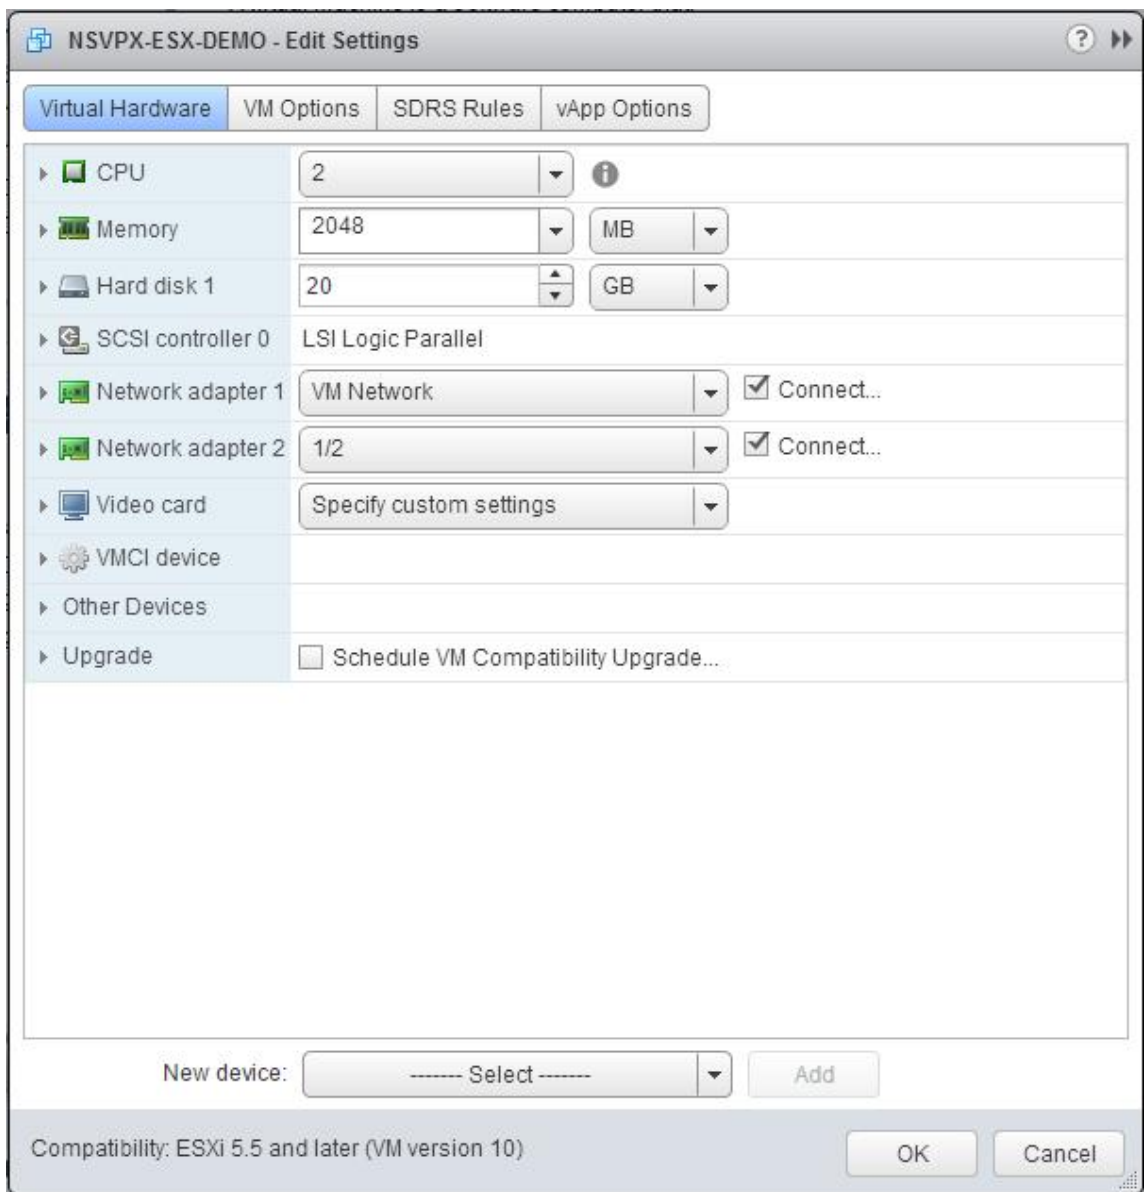
Configure a Citrix ADC VPX instance to use VMXNET3 network interface

September 12, 2024


After you have installed and configured the Citrix ADC VPX instance on the VMware ESX, you can use the VMware vSphere web client to configure the virtual appliance to use VMXNET3 network interfaces.

To configure Citrix ADC VPX instances to use VMXNET3 network interfaces by using the VMware vSphere Web Client:

1. In the vSphere Web Client, select Hosts and Clusters.
2. Upgrade the Compatibility setting of the Citrix ADC VPX instance to ESX, as follows:
 - a. Power off the Citrix ADC VPX instance.
 - b. Right-click the Citrix ADC VPX instance and select Compatibility > Upgrade VM Compatibility.
 - c. In the Configure VM Compatibility dialog box, select ESXi 5.5 and later from the Compatible with drop-down list and click OK.
3. Right-click on the Citrix ADC VPX instance and click Edit Settings.



4. In the <virtual_appliance> - Edit Settings dialog box, click the CPU section.



5. In the CPU section, update the following:

- Number of CPUs
- Number of Sockets
- Reservations
- Limit
- Shares

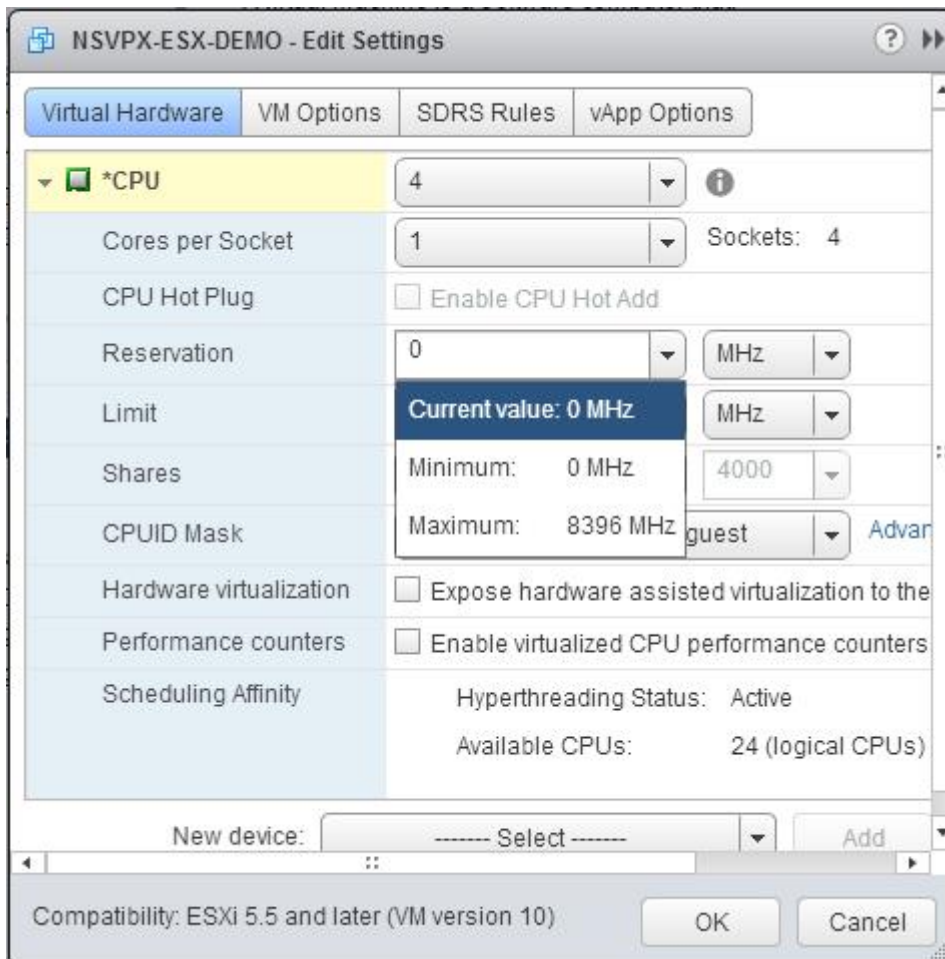
Set the values as follows:

- a. In the CPU drop-down list, select the number of CPUs to assign to the virtual appliance.
- b. In the Cores per Socket drop-down list, select the number of sockets.
- c. (Optional) In the CPU Hot Plug field, select or unselect the Enable CPU Hot Add check box.


Note:

Citrix recommends accepting the default (disabled).


d. In the Reservation drop-down list, select the number that is shown as the maximum value.



e. In the Limit drop-down list, select the number that is shown as the maximum value.



f. In the Shares drop-down lists, select Custom and the number that is shown as the maximum value.



6. In the Memory section, update the following:

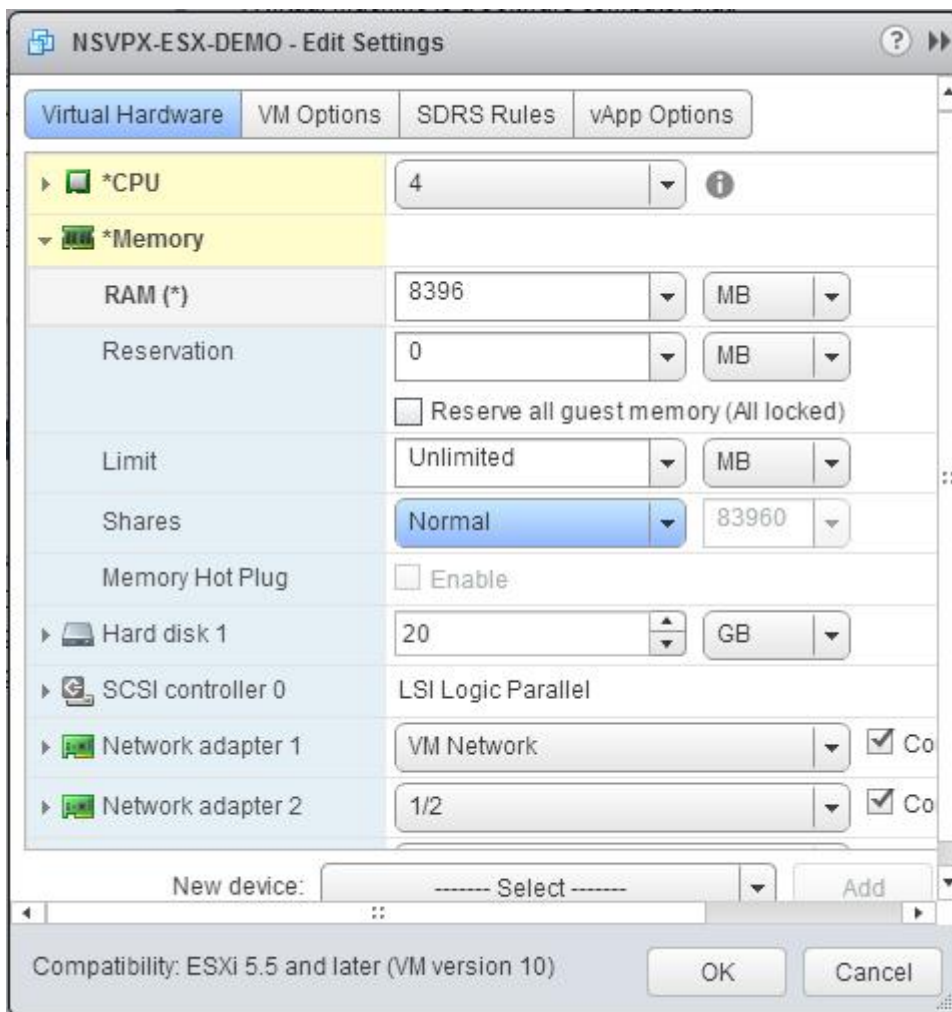
- Size of RAM
- Reservations
- Limit
- Shares

Set the values as follows:

a. In the RAM drop-down list, select the size of the RAM. It must be the number of vCPUs x 2 GB. For example, if the number of vCPUs is 4, the RAM must be 4 x 2 GB = 8 GB.

Note:


For an Advanced or Premium edition of the Citrix ADC VPX appliance, make sure that you allocate 4 GB of RAM to each vCPU. For example, if the number of vCPU is 4 then RAM = 4 x 4 GB = 16 GB.




b. In the Reservation drop-down list, enter the value for the memory reservation, and select the Reserve all guest memory (All locked) check box. The memory reservation must be the number of vCPUs x 2 GB. For example, if the number of vCPUs is 4, the memory reservation must be 4 x 2 GB = 8 GB.

Note:

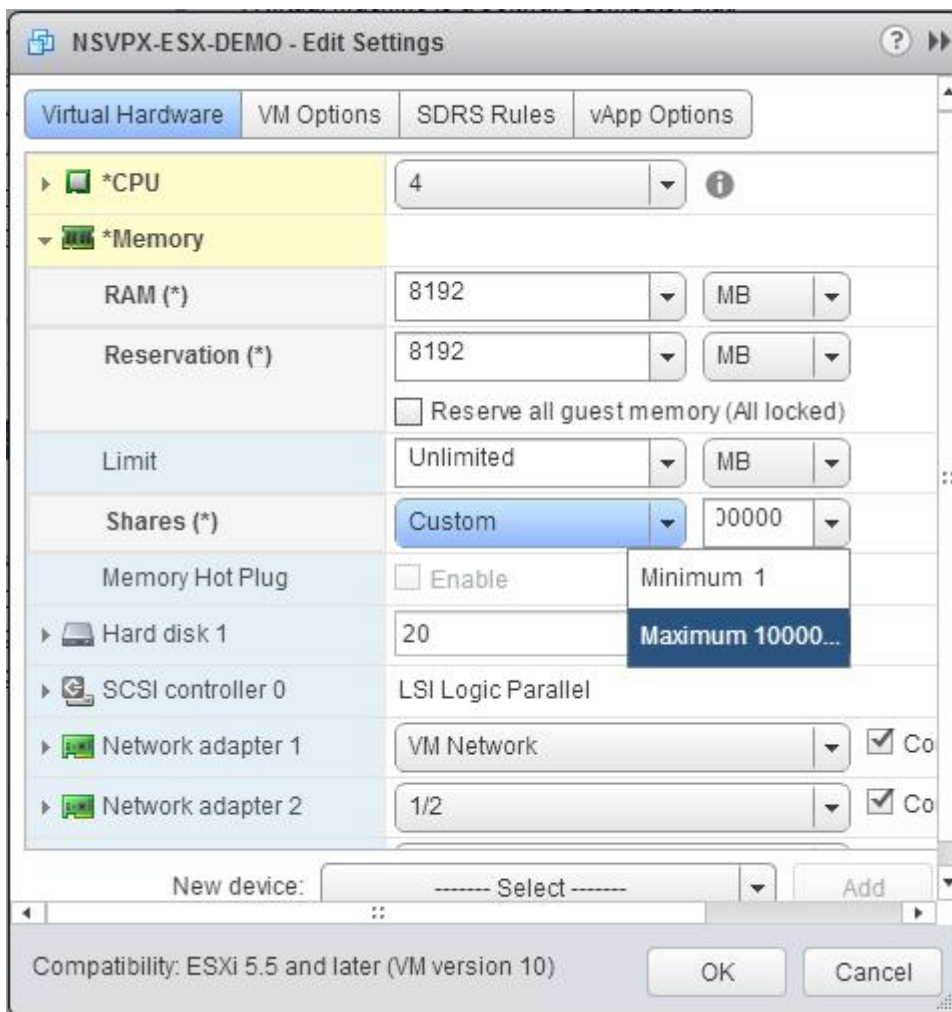
For an Advanced or Premium edition of the Citrix ADC VPX appliance, make sure that you allocate 4 GB of RAM to each vCPU. For example, if the number of vCPU is 4 then RAM = 4 x 4 GB = 16 GB.




c. In the Limit drop-down list, select the number that is shown as the maximum value.



d. In the Shares drop-down lists, select Custom and the number that is shown as the maximum value.



7. Add a VMXNET3 network interface. From the New device drop-down list, select Network and click **Add**.




8. In the New Network section, from the drop-down list, select the network interface, and do the following:

a. In the Adapter Type drop-down list, select VMXNET3.

Important:

The default E1000 network interface and VMXNET3 cannot coexist, make sure that you remove the E1000 network interface and use VMXNET3 (0/1) as the management interface.



9. Click **OK**.
10. Power on the Citrix ADC VPX instance.
11. Once the Citrix ADC VPX instance powers on, you can use the following command to verify the configuration:

```
show interface summary
```

The output must show all the interfaces that you configured:

```

1 > show interface summary
2 -----
3           Interface  MTU      MAC                               Suffix
4 -----
5 1      0/1          1500    00:0c:29:89:1d:0e               NetScaler Vir...rface,
6      VMXNET3
7 2      1/1          9000    00:0c:29:89:1d:18               NetScaler Vir...rface,
8      VMXNET3

```

7	3	1/2 VMXNET3	9000	00:0c:29:89:1d:22	NetScaler Vir...rface,
8	4	LO/1 interface	9000	00:0c:29:89:1d:0e	Netscaler Loopback

Note:

After you add a VMXNET3 interface and restart the Citrix ADC VPX appliance, the VMware ESX hypervisor might change the order in which the NIC is presented to the VPX appliance. So, network adapter 1 might not always remain 0/1, resulting in loss of management connectivity to the VPX appliance. To avoid this issue, change the virtual network of the network adapter accordingly.

This is a VMware ESX hypervisor limitation.

Configure a Citrix ADC VPX instance to use SR-IOV network interface

September 9, 2024

After you have installed and configured the Citrix ADC VPX instance on VMware ESX, you can use the VMware vSphere web client to configure the virtual appliance to use single root I/O virtualization (SR-IOV) network interfaces.

Limitations

A Citrix ADC VPX configured with SR-IOV network interface has the following limitations:

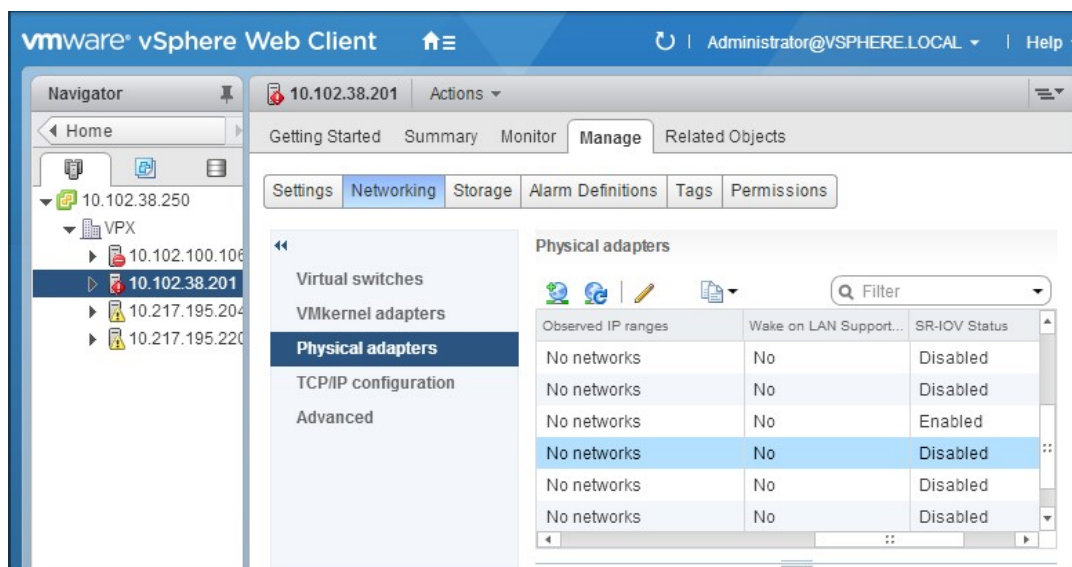
- The following features are not supported on SR-IOV interfaces using the Intel 82599 10G NIC on ESX VPX:
 - L2 mode switching
 - Static Link Aggregation and LACP
 - Clustering
 - Admin partitioning [Shared VLAN mode]
 - High Availability [Active - Active mode]
 - Jumbo frames
 - IPv6
- The following features are not supported on the SR-IOV interface with an Intel 82599 10G NIC on KVM VPX:
 - Static Link Aggregation and LACP
 - L2 mode switching

- Clustering
- Admin partitioning [Shared VLAN mode]
- High Availability [Active –Active mode]
- Jumbo frames
- IPv6
- VLAN configuration on Hypervisor for SR-IOV VF interface through `ip link` command is not supported

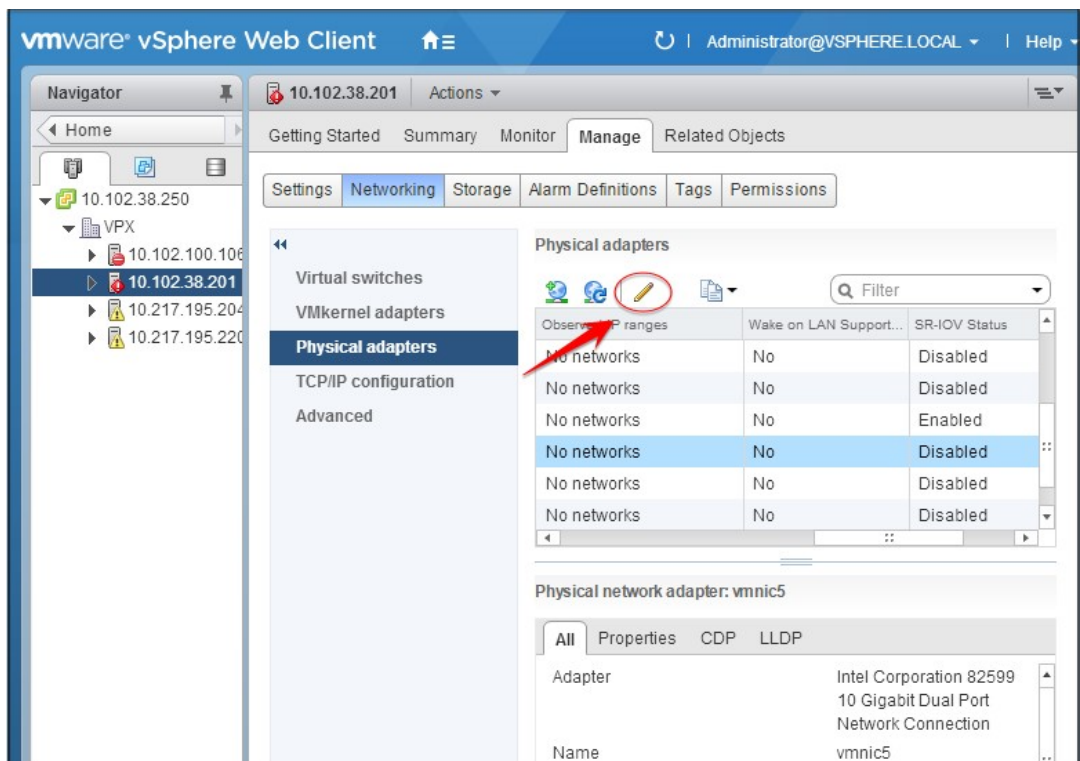
Prerequisite

Make sure that you:

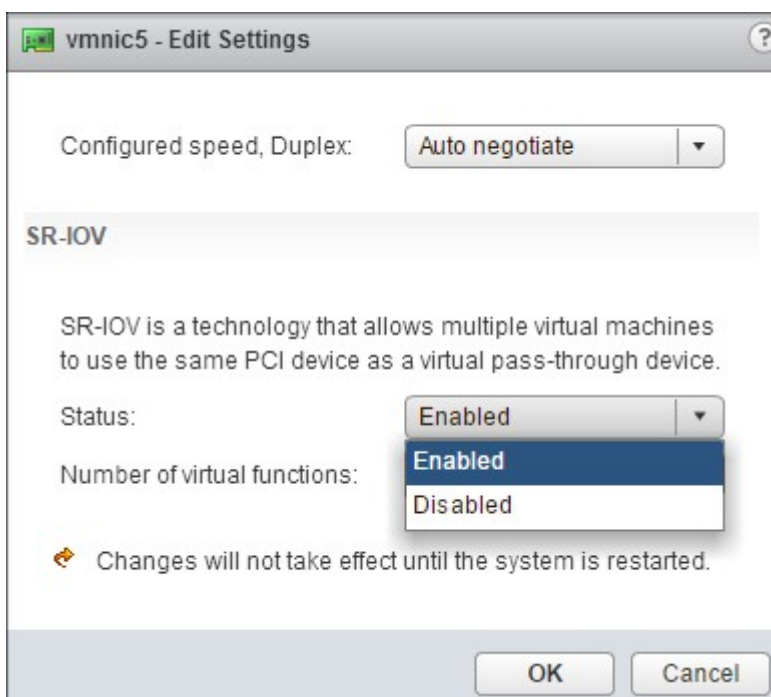
- Add the Intel 82599 NIC (NIC) to the ESX Host. IXGBE driver version 3.7.13.7.14iov is recommended.
- Enable SR-IOV on the host physical adapter, as follows:
 1. In the vSphere Web Client, navigate to the Host.
 2. On the **Manage > Networking** tab, select **Physical adapters**. The SR-IOV Status field shows whether a physical adapter supports SR-IOV.




3. Select the physical adapter, and then click the pencil icon to open the **Edit Settings** dialog box.



4. Under SR-IOV, select **Enabled** from the **Status** drop-down list.



5. In the **Number of virtual functions** field, enter the number of virtual functions that you want to configure for the adapter.




6. Click **OK**.
 7. Restart the host.
- Create a Distributed Virtual Switch (DVS) and [Portgroups](#). For instructions, see the VMware Documentation.

Note:

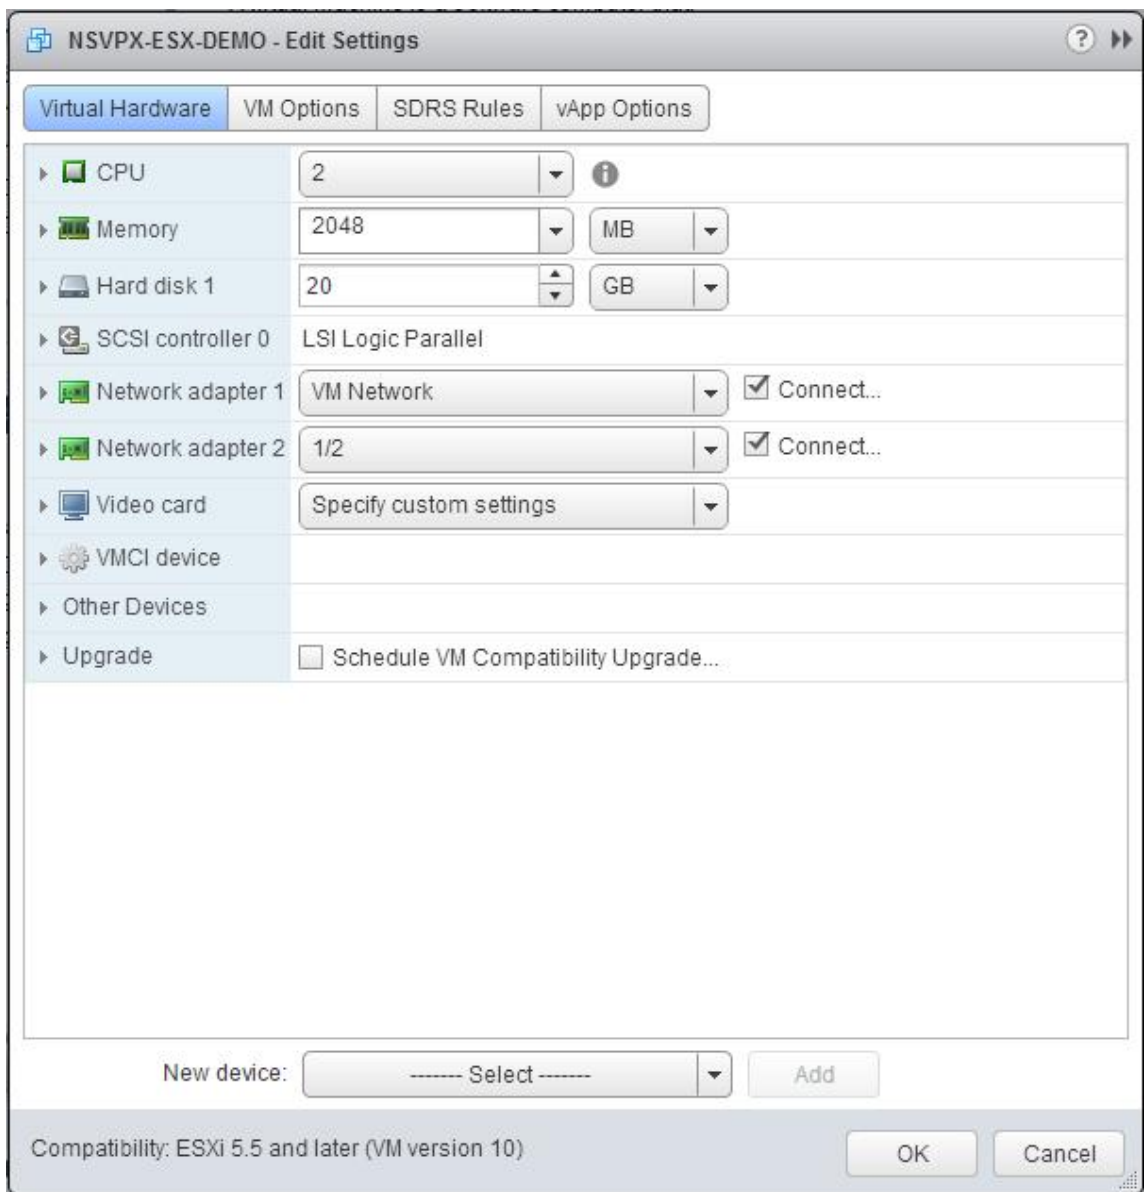
Citrix has qualified the SR-IOV configuration on DVS and [Portgroups](#) only.

To configure Citrix ADC VPX instances to use SR-IOV network interface by using VMware vSphere Web Client:


1. In the vSphere Web Client, select **Hosts and Clusters**.
2. Upgrade the Compatibility setting of the Citrix ADC VPX instance to ESX 5.5 or later, as follows:
 - a. Power off the Citrix ADC VPX instance.
 - b. Right-click the Citrix ADC VPX instance and select **Compatibility > Upgrade VM Compatibility**.
 - c. In the **Configure VM Compatibility** dialog box, select **ESXi 5.5 and later** from the **Compatible with** drop-down list and click **OK**.



3. Right-click on the Citrix ADC VPX instance and click **Edit Settings**.



4. In the **<virtual_appliance> - Edit Settings** dialog box, click the **CPU** section.



5. In the **CPU** section, update the following settings:

- Number of CPUs
- Number of Sockets
- Reservations
- Limit
- Shares

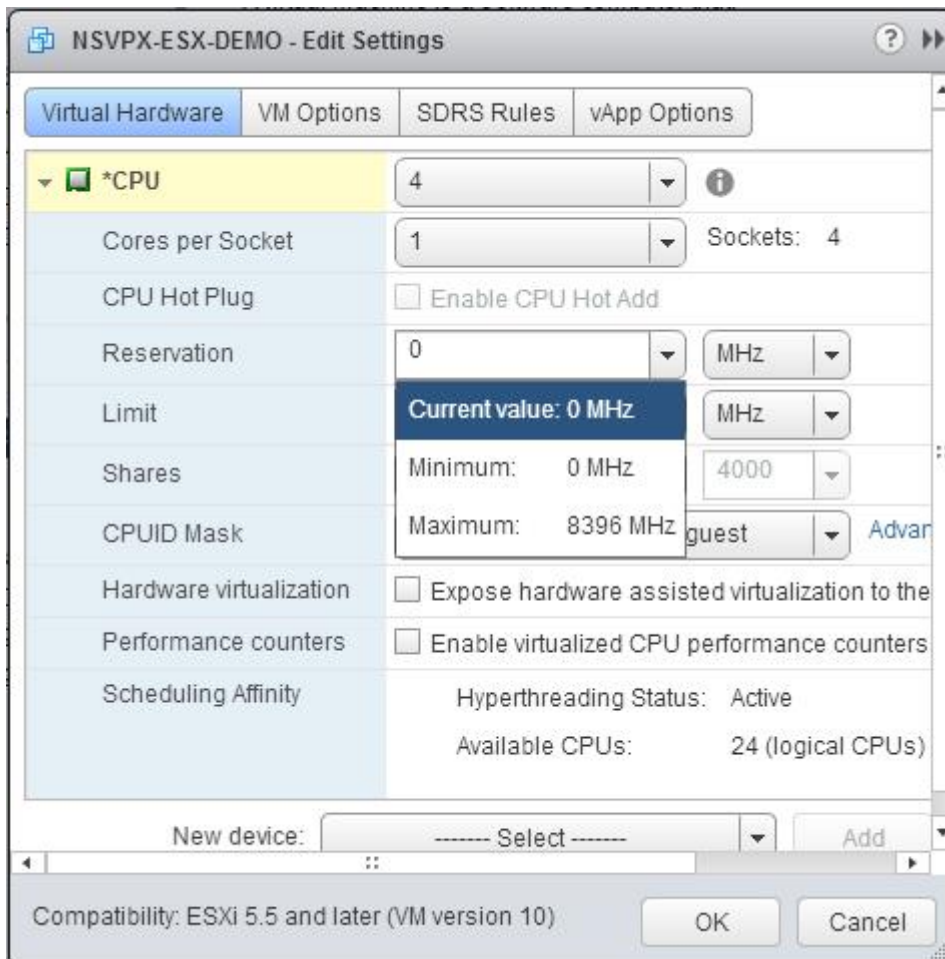
Set the values as follows:

- In the **CPU** drop-down list, select the number of CPUs to assign to the virtual appliance.
- In the **Cores per Socket** drop-down list, select the number of sockets.
- (Optional) In the **CPU Hot Plug** field, select or clear the **Enable CPU Hot Add** check box.


Note:

Citrix recommends accepting the default (disabled).

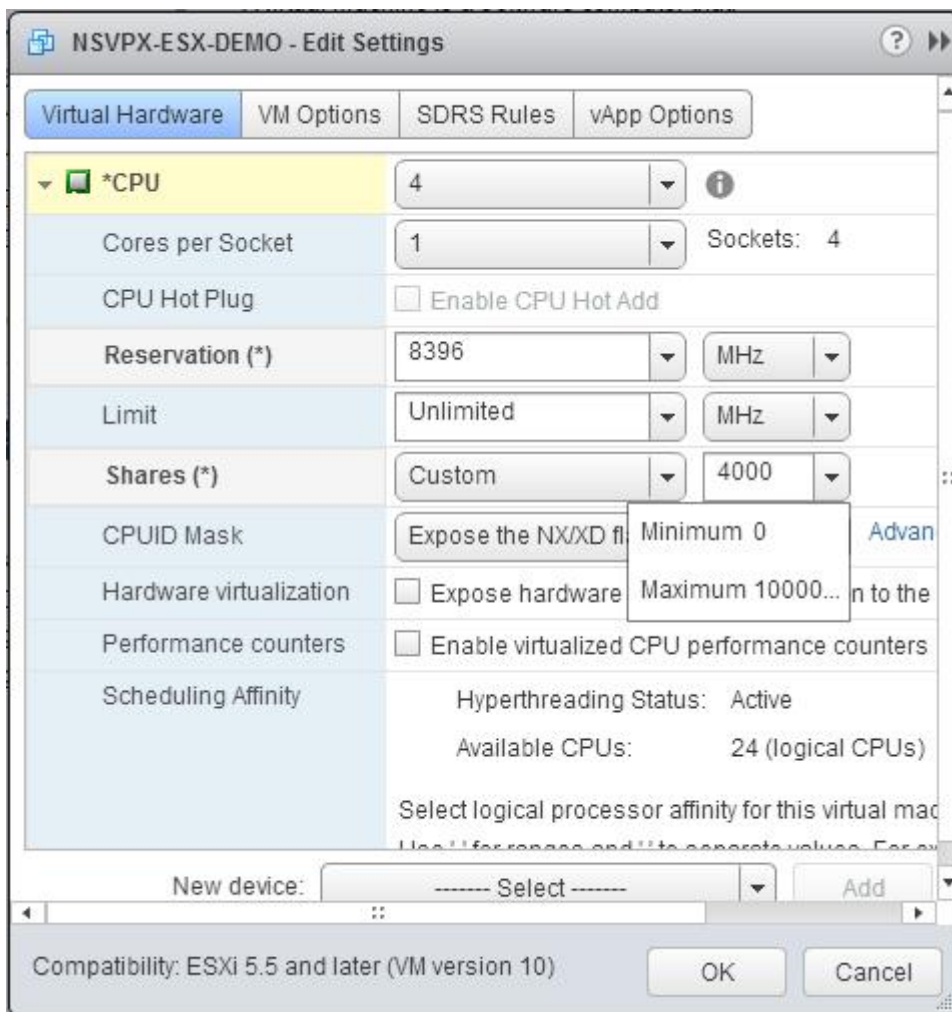
d. In the **Reservation** drop-down list, select the number that is shown as the maximum value.



e. In the **Limit** drop-down list, select the number that is shown as the maximum value.



f. In the **Shares** drop-down lists, select **Custom** and the number that is shown as the maximum value.




6. In the **Memory** section, update the following settings:

- Size of RAM
- Reservations
- Limit
- Shares

Set the values as follows:

a. In the **RAM** drop-down list, select the size of the RAM. It must be the number of vCPUs x 2 GB. For example, if the number of vCPU is 4 then RAM = 4 x 2 GB = 8 GB.


Note: For Advanced or Premium edition of the Citrix ADC VPX appliance, make sure that you allocate 4 GB of RAM to each vCPU. For example, if the number of vCPU is 4 then RAM = 4 x 4 GB = 16 GB.




b. In the **Reservation** drop-down list, enter the value for the memory reservation, and select the **Reserve all guest memory (All locked)** check box. The memory reservation must be number of vCPUs x 2 GB. For example, if the number of vCPUs is 4, the memory reservation must be 4 x 2 GB = 8 GB.

Note:


For Advanced or Premium edition of the Citrix ADC VPX appliance, make sure that you allocate 4 GB of RAM to each vCPU. For example, if the number of vCPU is 4 then RAM = 4 x 4 GB = 16 GB.




c. In the **Limit** drop-down list, select the number that is shown as the maximum value.




d. In the **Shares** drop-down lists, select **Custom**, and select the number that is shown as the maximum value.




7. Add an SR-IOV network interface. From the **New device** drop-down list, select **Network** and click **Add**.




8. In the **New Network** section. From the drop-down list, select the **Portgroup** that you created, and do the following:
 - a. In the **Adapter Type** drop-down list, select **SR-IOV passthrough**.



- b. In the **Physical function** drop-down list, select the physical adapter mapped with the Portgroup.



- c. In the **Guest OS MTU Change** drop-down list, select **Disallow**.
9. In the **<virtual_appliance> - Edit Settings** dialog box, click the **VM Options** tab.
10. On the **VM Options** tab, select the **Advanced** section. From the **Latency Sensitivity** drop-down list, select **High**.



11. Click **OK**.
12. Power on the Citrix ADC VPX instance.
13. Once the Citrix ADC VPX instance powers on, you can use the following command to verify the configuration:

```
show interface summary
```

The output must show all the interfaces that you configured:

```
1 > show interface summary
2 -----
3      Interface  MTU      MAC      Suffix
```

```

4 -----
5 1      0/1      1500      00:0c:29:1b:81:0b      NetScaler Virtual
      Interface
6 2      10/1     1500      00:50:56:9f:0c:6f      Intel 82599 10G VF
      Interface
7 3      10/2     1500      00:50:56:9f:5c:1e      Intel 82599 10G VF
      Interface
8 4      10/3     1500      00:50:56:9f:02:1b      Intel 82599 10G VF
      Interface
9 5      10/4     1500      00:50:56:9f:5a:1d      Intel 82599 10G VF
      Interface
10 6     10/5     1500      00:50:56:9f:4e:0b      Intel 82599 10G VF
      Interface
11 7     L0/1     1500      00:0c:29:1b:81:0b      Netscaler Loopback
      interface
12 Done
13 > show inter 10/1
14 1)      Interface 10/1 (Intel 82599 10G VF Interface) #1
15      flags=0xe460 <ENABLED, UP, UP, HAMON, 802.1q>
16      MTU=1500, native vlan=55, MAC=00:50:56:9f:0c:6f, uptime 0
      h21m53s
17      Actual: media FIBER, speed 10000, duplex FULL, fctl NONE,
      throughput 10000
18      LLDP Mode: NONE,                LR Priority: 1024
19
20      RX: Pkts(838020742) Bytes(860888485431) Errs(0) Drops(2527)
      Stalls(0)
21      TX: Pkts(838149954) Bytes(860895860507) Errs(0) Drops(0) Stalls
      (0)
22      NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
23      Bandwidth thresholds are not set.
24 Done

```

Migrating the Citrix ADC VPX from E1000 to SR-IOV or VMXNET3 Network Interfaces

September 6, 2024

May 24, 2018

You can configure your existing Citrix ADC VPX instances that use E1000 network interfaces to use SR-IOV or VMXNET3 network interfaces.

To configure an existing Citrix ADC VPX instance to use SR-IOV network interfaces, see [Configure a Citrix ADC VPX instance to use SR-IOV network interface](#).

To configure an existing Citrix ADC VPX instance to use VMXNET3 network interfaces, see [Configure a](#)

[Citrix ADC VPX instance to use VMXNET3 network interface.](#)

Configure a Citrix ADC VPX instance to use PCI passthrough network interface

September 12, 2024

Overview

After you have installed and configured a Citrix ADC VPX instance on VMware ESX Server, you can use the vSphere Web Client to configure the virtual appliance to use PCI passthrough network interfaces.

The PCI passthrough feature allows a guest virtual machine to directly access physical PCI and PCIe devices connected to a host.

Prerequisites

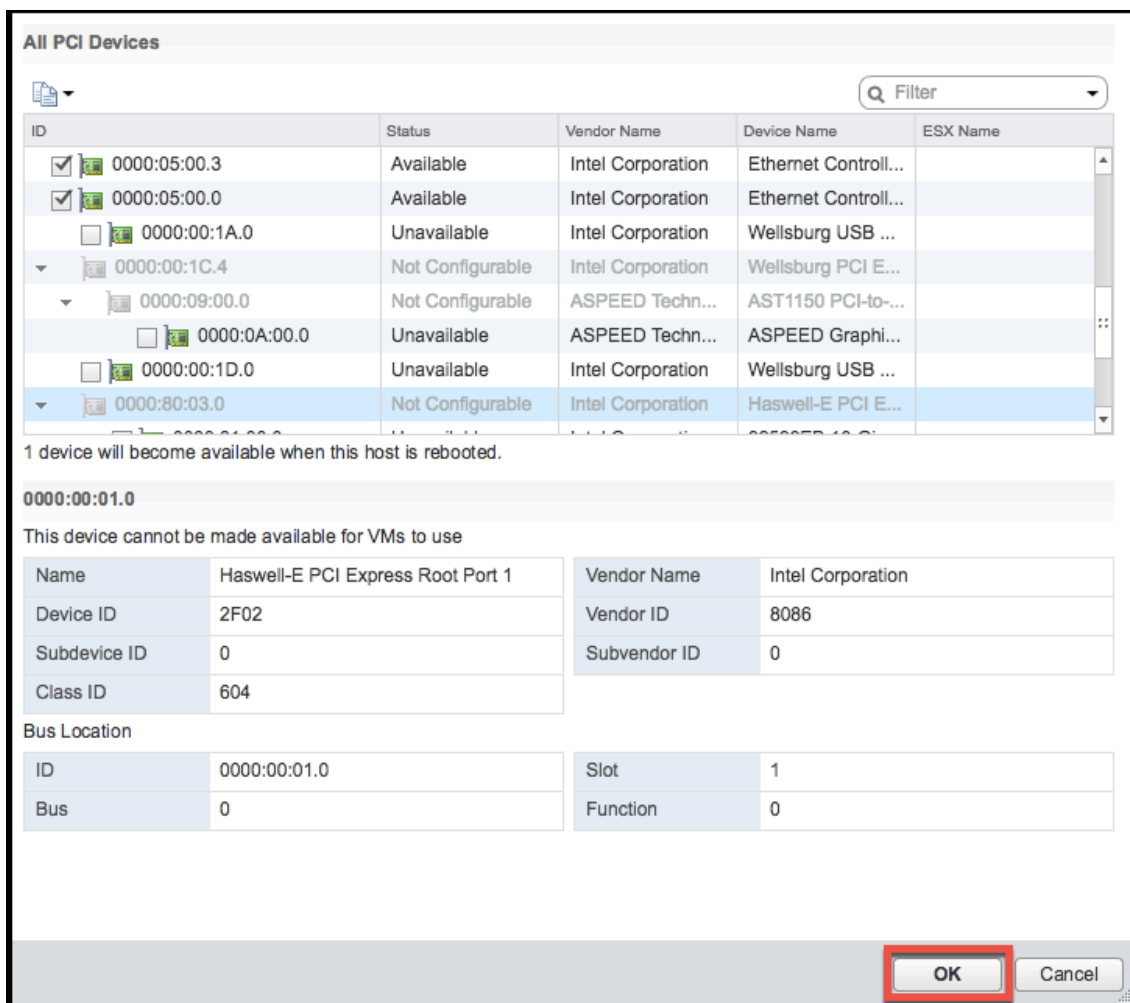
- The firmware version of the Intel XL710 NIC on the host is 5.04.
- A PCI passthrough device connected to and configured on the host
- Supported NICs:
 - Intel X710 10G NIC
 - Intel XL710 Dual Port 40G NIC
 - Intel XL710 Single Port 40G NIC

Configure passthrough devices on a host

Before configuring a passthrough PCI device on a virtual machine, you must configure it on the host machine. Follow these steps to configure passthrough devices on a host.

1. Select the host from the Navigator panel of the vSphere Web Client.
2. Click **Manage > Settings > PCI Devices**. All available passthrough devices are displayed.
3. Right-click the device that you want to configure and click **Edit**.
4. The **Edit PCI Device Availability** window appears.

5. Select the devices to be used for passthrough and click **OK**.




6. Restart the host machine.

Configure passthrough devices on a Citrix ADC VPX instance

Follow these steps to configure a passthrough PCI device on a Citrix ADC VPX instance.

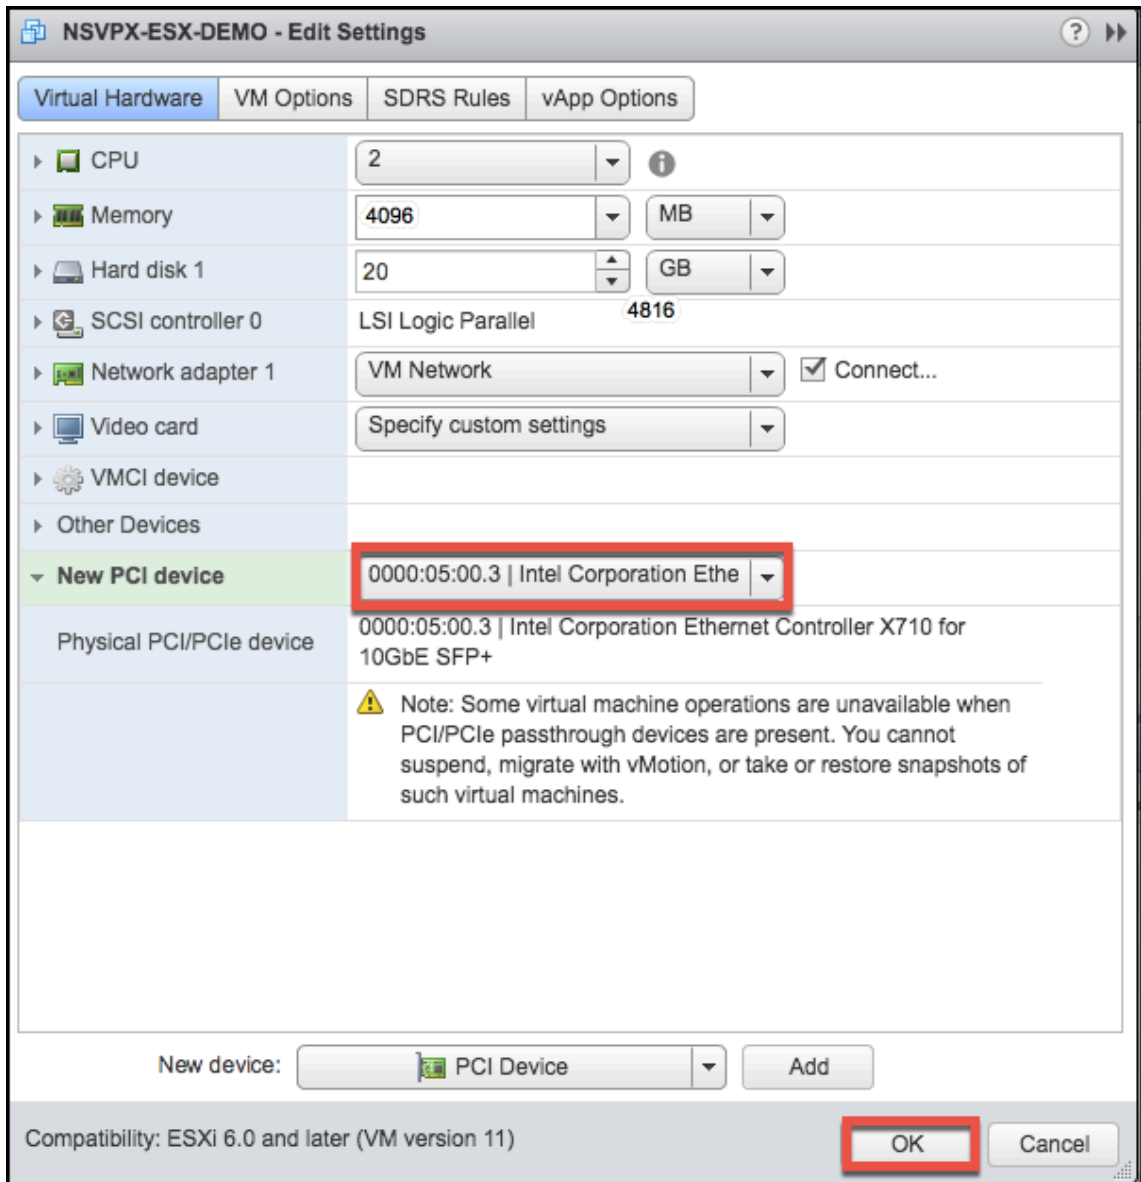
1. Power off the virtual machine.
2. Right-click the virtual machine and select **Edit Settings**.
3. On the **Virtual Hardware** tab, select **PCI Device** from the **New Device** drop-down menu, and click **Add**.



4. Expand **New PCI device** and select the passthrough device to connect to the virtual machine from the drop-down list and click **OK**.

Note:

VMXNET3 network interface and PCI Passthrough Network Interface cannot coexist.



5. Power on the guest virtual machine.

You have completed the steps to configuring Citrix ADC VPX to use PCI passthrough network interfaces.

Apply Citrix ADC VPX configurations at the first boot of the Citrix ADC appliance on VMware ESX hypervisor

September 6, 2024

You can apply the Citrix ADC VPX configurations during the first boot of the Citrix ADC appliance on

the VMware ESX hypervisor. Therefore in certain cases, a specific setup or VPX instance is brought up in much lesser time.

For more information on Preboot user data and its format, see [Apply Citrix ADC VPX configurations at the first boot of the Citrix ADC appliance in cloud](#).

Note:

To bootstrap using preboot user data in ESX, default gateway config must be passed in `<NS-CONFIG>` section. For more information on the content of the `<NS-CONFIG>` tag, see [Sample-`<NS-CONFIG>`-section](#).

Sample `<NS-CONFIG>` section:

```
1 <NS-PRE-BOOT-CONFIG>
2
3   <NS-CONFIG>
4     add route 0.0.0.0 0.0.0.0 10.102.38.1
5   </NS-CONFIG>
6
7   <NS-BOOTSTRAP>
8     <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
9     <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>
10
11   <MGMT-INTERFACE-CONFIG>
12     <INTERFACE-NUM> eth0 </INTERFACE-NUM>
13     <IP> 10.102.38.216 </IP>
14     <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
15   </MGMT-INTERFACE-CONFIG>
16 </NS-BOOTSTRAP>
17
18 </NS-PRE-BOOT-CONFIG>
```

How to provide preboot user data on ESX hypervisor

You can provide preboot user data on ESX hypervisor in the following two ways:

- Using CD/DVD ISO
- Using OVF Property

Provide user data using CD/DVD ISO

You can use VMware vSphere client to inject user data into the VM as an ISO image using the CD/DVD drive.

Follow these steps to provide user data using CD/DVD ISO:

1. Create a file with file name `userdata` that contains the preboot user data content. For more information on the content of the `<NS-CONFIG>` tag, see Sample `<NS-CONFIG>` section.

Note:

File name must be strictly used as `userdata`.

2. Store the `userdata` file in a folder, and build an ISO image using the folder.


You can build an ISO image with `userdata` file by the following two methods:

- Using any image processing tool such as PowerISO.
- Using `mkisofs` command in Linux.


The following sample configuration shows how to generate an ISO image using the `mkisofs` command in Linux.

```
1 root@ubuntu:~/sai/14jul2021# ls -l total 4
2 drwxr-xr-x 2 root root 4096 Jul 14 12:32 esx_preboot_userdata
3 root@ubuntu:~/sai/14jul2021#
4 root@ubuntu:~/sai/14jul2021# ls -l esx_preboot_userdata/total 4
5 -rw-r--r-- 1 root root 3016 Jul 14 12:32 userdata
6 root@ubuntu:~/sai/14jul2021# mkisofs -o esx_preboot_userdata.iso
  ./esx_preboot_userdata
7 I: -input-charset not specified, using utf-8 (detected in locale
  settings)
8 Total translation table size: 0
9 Total rockridge attributes bytes: 0
10 Total directory bytes: 112
11 Path table size(bytes): 10
12 Max brk space used 0
13 176 extents written (0 MB)
14 root@ubuntu:~/sai/14jul2021# ls -lh
15 total 356K
16 drwxr-xr-x 2 root root 4.0K Jul 14 12:32 esx_preboot_userdata
17 -rw-r--r-- 1 root root 352K Jul 14 12:34 esx_preboot_userdata.iso
18
19 root@ubuntu:~/sai# ls preboot_userdata_155_193 userdata
20 root@ubuntu:~/sai# mkisofs -o preboot_userdata_155_193.iso ./
  preboot_userdata_155_193
21 I: -input-charset not specified, using utf-8 (detected in locale
  settings)
22 Total translation table size: 0
23 Total rockridge attributes bytes: 0
24 Total directory bytes: 112
25 Path table size(bytes): 10
26 Max brk space used 0
27 176 extents written (0 MB)
```


3. Provision the Citrix ADC VPX instance using standard deployment process to create the VM. But do not power on the VM automatically.



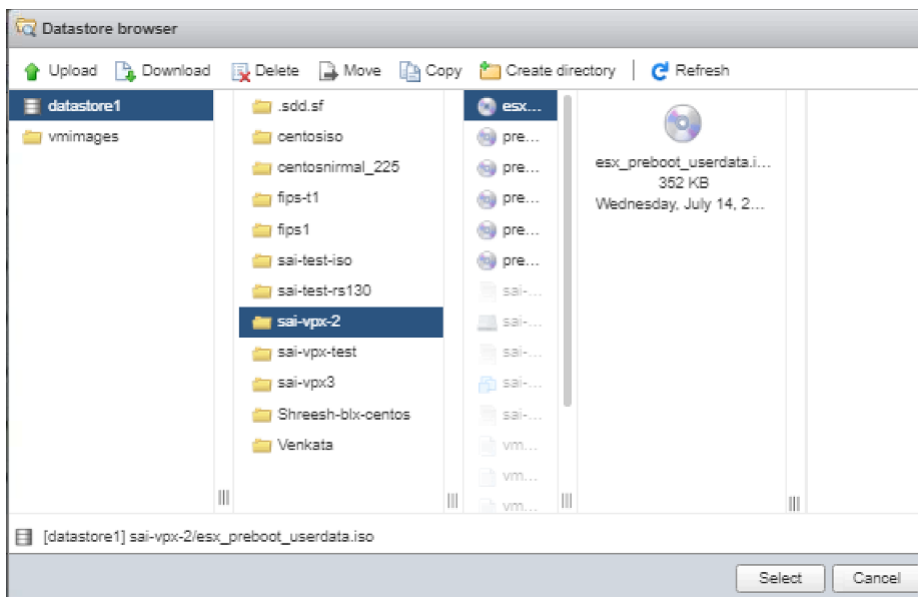
4. After the VM is successfully created, attach the ISO file as CD/DVD drive to the VM.



5. Navigate to **New CD/DVD Drive** and choose **Datastore ISO file** from the drop-down menu.



6. Select a Datastore in the vSphere Client.



7. Power on the VM.

Provide user data using OVF property

Follow these steps to provide user data using OVF property.

1. Create a file with user data content.

```

root@ubuntu:~/sai/14jul2021# cat esx_userdata.xml
<NS-PRE-BOOT-CONFIG>
  <NS-CONFIG>
    add route 0.0.0.0 0.0.0.0 10.102.38.1
  </NS-CONFIG>

  <NS-BOOTSTRAP>
    <SKIP-DEFAULT-BOOTSTRAP>YES</SKIP-DEFAULT-BOOTSTRAP>
    <NEW-BOOTSTRAP-SEQUENCE>YES</NEW-BOOTSTRAP-SEQUENCE>

    <MGMT-INTERFACE-CONFIG>
      <INTERFACE-NUM> eth0 </INTERFACE-NUM>
      <IP> 10.102.38.219 </IP>
      <SUBNET-MASK> 255.255.255.0 </SUBNET-MASK>
    </MGMT-INTERFACE-CONFIG>
  </NS-BOOTSTRAP>
</NS-PRE-BOOT-CONFIG>

```

2. Encode the user data content with Base64 encoding. You can perform the Base64 encoding using the following two methods:

- In Linux, use the following command:

```
1 base64 <userdata-filename> > <output-file>
```

Example:

```
1 base64 esx_userdata.xml > esx_userdata_b64
```

```
root@ubuntu:~/sai/14jul2021# base64 esx_userdata.xml > esx_userdata_b64
root@ubuntu:~/sai/14jul2021#
root@ubuntu:~/sai/14jul2021# cat esx_userdata_b64
PE5TLVBSRS1CT09ULUNPTkZJRz4KICAgIDx0Uy1DT05GSUc+CglhZGQgcm91dGUgMC4wLjAuMCAw
LjAuMCAwIDEwLjEwMi4zOC4xCiAgICA8L05TLUNPTkZJRz4KICAgICA8TlMtQk9PFVNUUkFQgog
ICAgICA8U0tJUC1ERUZBVUxULUJPT1RTVFJBUD5ZRVM8L1NLSVAtREVGVQVVMVC1CT09U
U1RSQVA+CiAgICA8ICA8ICA8IDxORVetQk9PFVNUUkFQLVNFUUVVFTkNFP1lFUzwwTkVXLUJPT1RT
VFJBUC1TRVFRU5DRT4KICAgICA8ICA8PE1HTVQtSU5URVJGQUNFLUNPTkZJRz4KICAgICA8ICA8
ICA8ICA8IDxJTlRFUkZBQ0UtTlVNPiBlbGgwIDwvSU5URVJGQUNFLU5VTT4KICAgICA8ICA8ICA8
ICA8IDxJUD4gICA8MTAUMTAyLjM4LjIxOSA8L01QPgogICAgICA8ICA8ICA8ICA8PFNVQk5FVC1N
QVNLPlAyNTUuMjU1LjI1NS4wIDwvU1VCTkVULU1BU0s+CiAgICA8ICA8PC9NR01ULU1OVEVSRkFD
RS1DT05GSUc+CiAgICA8L05TLUJPT1RTVFJBUD4KPC90Uy1QUkUtQk9PVC1DT05GSUc+Cg==
```

- Use online tools to encode user data content, for example, Base64 Encode and Decode.

3. Include a **Product** section in the OVF template of a Citrix ADC VPX instance on ESX hypervisor.

Sample Product section:

```
1 <ProductSection>
2
3 <Info>Information about the installed software</Info>
4 <Product>NSVPX-VSK Template</Product>
5 <Vendor>Citrix</Vendor>
6 <VendorUrl>www.citrix.com</VendorUrl>
7 <Category> Preboot Userdata </Category>
8
9 <Property ovf:key="guestinfo.userdata" ovf:type="string" ovf:
   userConfigurable="true" ovf:value="">
10
11 <Label>Userdata</Label>
12 <Description> Userdata for ESX VPX </Description>
13 </Property>
14
15 </ProductSection>
```

4. Provide the base64 encoded user data as the `ovf:value` for `guestinfo.userdata` property in the Product section.

```
1 <ProductSection>
2
3 <Info>Information about the installed software</Info>
4 <Product>NSVPX-VSK Template</Product>
5 <Vendor>Citrix</Vendor>
6 <VendorUrl>www.citrix.com</VendorUrl>
7 <Category> Preboot Userdata </Category>
8 <Property ovf:key="guestinfo.userdata" ovf:type="string" ovf:
   userConfigurable="true"
9   ovf:value="PE5TLVBSRS1CT09ULUNPTkZJRz4KICAgIDx0Uy1DT05GSUc+
   CglhZGQgcm91dGUgMC4wLjAuMCAw
10   LjAuMCAwIDEwLjEwMi4zOC4xCiAgICA8L05TLUNPTkZJRz4KICAgICA8TlMtQk9PFVNUUkFQgog
   ICAgICA8U0tJUC1ERUZBVUxULUJPT1RTVFJBUD5ZRVM8L1NLSVAtREVGVQVVMVC1CT09U
   U1RSQVA+CiAgICA8ICA8ICA8IDxORVetQk9PFVNUUkFQLVNFUUVVFTkNFP1lFUzwwTkVXLUJPT1RT
   VFJBUC1TRVFRU5DRT4KICAgICA8ICA8PE1HTVQtSU5URVJGQUNFLUNPTkZJRz4KICAgICA8ICA8ICA8
   ICA8IDxJUD4gICA8MTAUMTAyLjM4LjIxOSA8L01QPgogICAgICA8ICA8ICA8ICA8PFNVQk5FVC1N
   QVNLPlAyNTUuMjU1LjI1NS4wIDwvU1VCTkVULU1BU0s+CiAgICA8ICA8PC9NR01ULU1OVEVSRkFD
   RS1DT05GSUc+CiAgICA8L05TLUJPT1RTVFJBUD4KPC90Uy1QUkUtQk9PVC1DT05GSUc+Cg==
```

```

11     ICaGICaGICaGICa8U0tJUC1ERUZBVUxULUJPT1RTVFJBUD5ZRVm8L1NLSVAtREVGQVVMVC1C
12     U1RSQVA+
        CiAgICaGICaGICaGIDxORVctQk9PVFNuUkFQLVNFUVVFTkNFPlFUzWvTkVXLUJPT1RT
13     VFJBUC1TRVFVRU5DRT4KCiAgICaGICaGPE1HTVQtSU5URVJGQUNFLUNPTkZJRz4KICaGICaG
14     ICaGICaGIDxJTLRFUkZBQ0UtTlVNPiBlDGgwIDWvSU5URVJGQUNFLU5VTT4KICaGICaGICaG
15     ICaGIDxJUD4gICaGMTAuMTAyLjM4LjIxOSA8L0lQPgogICaGICaGICaGICaGICaGPFNVQk5F
16     QVNLPiAyNTUuMjU1LjI1NS4wIDWvU1VCTkVULU1BU0s+
        CiAgICaGICaGPC9NR01ULU1OVEVSRkFD
17     RS1DT05GSUc+
        CiAgICA8L05TLUJPT1RTVFJBUD4KPC90Uy1QUkUtQk9PVC1DT05GSUc+Cg
        ==">
18
19     <Label>Userdata</Label>
20     <Description> Userdata for ESX VPX </Description>
21 </Property>
22
23 </ProductSection>

```

5. Use the modified OVF template with Product section for the VM deployment.

```

Please change the default NSROOT password.
Enter new password:
Please re-enter your password:
Done
> sh ns ver
NetScaler NS13.0: Build 83.9005.nc, Date: Jul 13 2021, 02:56:05 (64-bit)
Done
> sh ns ip
state
-----
1) 10.102.38.219 0 NetScaler IP Active Enabled Enabled NA E
nabled
Done
> sh route
Network Netmask Gateway/OwnedIP VLAN State Traffic Domain Type
-----
1) 0.0.0.0 0.0.0.0 10.102.38.1 0 UP 0 STATI
C
2) 127.0.0.0 255.0.0.0 127.0.0.1 0 UP 0 PERMA
NENT
3) 10.102.38.0 255.255.255.0 10.102.38.219 0 UP 0 DIREC
T
Done

```

Install a Citrix ADC VPX instance on VMware cloud on AWS

September 12, 2024

The VMware Cloud (VMC) on AWS enables you to create cloud software-defined data centers (SDDC) on AWS with the desired number of ESX hosts. The VMC on AWS supports Citrix ADC VPX deployments.

VMC provides a user interface same as on-prem vCenter. It functions identical to the ESX-based Citrix ADC VPX deployments.

Prerequisites

Before you begin installing a virtual appliance, do the following:

- One VMware SDDC must be present with at least one host.
- Download the Citrix ADC VPX appliance setup files.
- Create appropriate network segments on VMware SDDC to which the virtual machines connect.
- Obtain VPX license files. For more information about Citrix ADC VPX instance licenses, see the *Citrix ADC VPX Licensing Guide* at </en-us/licensing/licensing-guide-for-netscaler>.

VMware cloud hardware requirements

The following table lists the virtual computing resources that the VMware SDDC must provide for each VPX nCore virtual appliance.

Table 1. Minimum virtual computing resources required for running a Citrix ADC VPX instance

Component	Requirement
Memory	2 GB
Virtual CPU (vCPU)	2
Virtual network interfaces	In VMware SDDC, you can install a maximum of 10 virtual network interfaces if the VPX hardware is upgraded to version 7 or higher.
Disk space	20 GB

Note:

This is in addition to any disk requirements for the hypervisor.

For production use of the VPX virtual appliance, the full memory allocation must be reserved.

OVF Tool 1.0 system requirements

OVF Tool is a client application that can run on Windows and Linux systems. The following table describes the minimum system requirements.

Table 2. Minimum system requirements for OVF tool installation

Component	Requirement
Operating system	For detailed requirements from VMware, search for the “OVF Tool User Guide” PDF file at http://kb.vmware.com/ .
CPU	750 MHz minimum, 1 GHz or faster recommended
RAM	1 GB Minimum, 2 GB recommended
NIC	100 Mbps or faster NIC

For information about installing OVF, search for the “OVF Tool User Guide” PDF file at <http://kb.vmware.com/>.

Downloading the Citrix ADC VPX setup files

The Citrix ADC VPX instance setup package for VMware ESX follows the Open Virtual Machine (OVF) format standard. You can download the files from the Citrix website. You need a Citrix account to log on. If you do not have a Citrix account, access the home page at <http://www.citrix.com>. Click the **New Users link**, and follow the instructions to create a new Citrix account.

Once logged on, navigate the following path from the Citrix home page:

Citrix.com > **Downloads > Citrix ADC > Virtual Appliances.**

Copy the following files to a workstation on the same network as the ESX server. Copy all three files into the same folder.

- NSVPX-ESX-<release number>-<build number>-disk1.vmdk (for example, NSVPX-ESX-13.0-79.64-disk1.vmdk)
- NSVPX-ESX-<release number>-<build number>.ovf (for example, NSVPX-ESX-13.0-79.64.ovf)
- NSVPX-ESX-<release number>-<build number>.mf (for example, NSVPX-ESX-13.0-79.64.mf)

Install a Citrix ADC VPX instance on VMware cloud

After you have installed and configured VMware SDDC, you can use the SDDC to install virtual appliances on the VMware cloud. The number of virtual appliances that you can install depends on the amount of memory available on the SDDC.

To install Citrix ADC VPX instances on VMware cloud, follow these steps:

1. Open VMware SDDC on your workstation.
2. In the **User Name** and **Password** text boxes, type the administrator credentials, and then click **Login**.
3. On the **File** menu, click **Deploy OVF Template**.
4. In the **Deploy OVF Template** dialog box, in **Deploy from file**, browse to the location at which you saved the Citrix ADC VPX instance setup files, select the .ovf file, and click **Next**.

Note:

By default, the Citrix ADC VPX instance uses E1000 network interfaces. To deploy ADC with the VMXNET3 interface, modify the OVF to use VMXNET3 interface instead of E1000.

5. Map the networks shown in the virtual appliance OVF template to the networks that you configured on the VMware SDDC. Click **Next** to start installing a virtual appliance on VMware SDDC.
6. You are now ready to start the Citrix ADC VPX instance. In the navigation pane, select the Citrix ADC VPX instance that you have installed and, from the right-click menu, select **Power On**. Click the **Console** tab to emulate a console port.
7. If you want to install another virtual appliance, repeat from Step 6.
8. Specify the management IP address from the same segment that is selected to be the management network. The same subnet is used for the Gateway.
9. The VMware SDDC requires that NAT and firewall rules are created explicitly for all private IP addresses belonging to network segments.

Install a Citrix ADC VPX instance on Microsoft Hyper-V server

September 12, 2024

To install Citrix ADC VPX instances on Microsoft Windows Server, you must first install Windows Server, with the Hyper-V role enabled, on a machine with adequate system resources. While installing the Hyper-V role, be sure to specify the NICs on the server that Hyper-V uses to create the virtual networks. You can reserve some NICs for the host. Use Hyper-V Manager to perform the Citrix ADC VPX instance installation.

Citrix ADC VPX instance for Hyper-V is delivered in virtual hard disk (VHD) format. It includes the default configuration for elements such as CPU, network interfaces, and hard-disk size and format. After you install Citrix ADC VPX instance, you can configure the network adapters on virtual appliance, add virtual NICs, and then assign the Citrix ADC IP address, subnet mask, and gateway, and complete the basic configuration of the virtual appliance.

After the initial configuration of the VPX instance, if you want to upgrade the appliance to the latest software release, see [Upgrade a Citrix ADC VPX standalone appliance](#)

Note:

Intermediate System-to-Intermediate System (ISIS) protocol is not supported on the Citrix ADC VPX virtual appliance hosted on the HyperV-2012 platform.

Prerequisites for installing Citrix ADC VPX instance on Microsoft servers

Before you begin installing a virtual appliance, do the following:

- Enable the Hyper-V role on Windows Servers. For more information, see [http://technet.microsoft.com/en-us/library/ee344837\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/ee344837(WS.10).aspx).
- Download the virtual appliance setup files.
- Obtain Citrix ADC VPX instance license files. For more information about Citrix ADC VPX instance licenses, see the *Citrix ADC VPX Licensing Guide* at https://support.citrix.com/s/article/CTX255959-how-to-allocate-and-install-citrix-netscaler-vpx-licenses?language=en_US.

Microsoft server hardware requirements

The following table describes the minimum system requirements for Microsoft servers.

Table 1. Minimum system requirements for Microsoft servers

Component	Requirement
CPU	1.4 GHz 64-bit processor
RAM	8 GB
Disk Space	32 GB or greater

The following table lists the virtual computing resources for each Citrix ADC VPX instance.

Table 2. Minimum virtual computing resources required for running a Citrix ADC VPX instance

Component	Requirement
RAM	4 GB
Virtual CPU	2

Component	Requirement
Disk Space	20 GB
Virtual Network Interfaces	1

Download the Citrix ADC VPX setup files

The Citrix ADC VPX instance for Hyper-V is delivered in virtual hard disk (VHD) format. You can download the files from the Citrix website. You need a Citrix account to log in. If you do not have a Citrix account, access the home page at <http://www.citrix.com>, click **Sign In > My account > Create Citrix Account**, and follow the instructions to create a Citrix account.

To download the Citrix ADC VPX instance setup files, follow these steps:

1. In a web browser, go to <http://www.citrix.com/>.
2. Sign in with your user name and password.
3. Click **Downloads**.
4. In **Select a Product** drop-down menu, select **Citrix ADC (NetScaler ADC)**.
5. Under **Citrix ADC Release X.X > Virtual Appliances**, click **Citrix ADC VPX Release X.X**
6. Download the compressed file to your server.

Install the Citrix ADC VPX instance on Microsoft servers

After you have enabled the Hyper-V role on Microsoft Server and extracted the virtual appliance files, you can use Hyper-V Manager to install Citrix ADC VPX instance. After you import the virtual machine, you need to configure the virtual NICs by associating them to the virtual networks created by Hyper-V.

You can configure a maximum of eight virtual NICs. Even if the physical NIC is DOWN, the virtual appliance assumes that the virtual NIC is UP, because it can still communicate with the other virtual appliances on the same host (server).

Note:

You cannot change any settings while the virtual appliance is running. Shut down the virtual appliance and then make changes.

To install Citrix ADC VPX instance on Microsoft Server by using Hyper-V Manager:

1. To start Hyper-V Manager, click **Start**, point to **Administrative Tools**, and then click **Hyper-V Manager**.
2. In the navigation pane, under **Hyper-V Manager**, select the server on which you want to install Citrix ADC VPX instance.
3. On the **Action** menu, click **Import Virtual Machine**.
4. In the **Import Virtual Machine** dialog box, in **Location**, specify the path of the folder that contains the Citrix ADC VPX instance software files, and then select **Copy the virtual machine (create a new unique ID)**. This folder is the parent folder that contains the Snapshots, Virtual Hard Disks, and Virtual Machines folders.

Note:

If you received a compressed file, make sure that you extract the files into a folder before you specify the path to the folder.

1. Click **Import**.
2. Verify that the virtual appliance that you imported is listed under **Virtual Machines**.
3. To install another virtual appliance, repeat steps **2** through **6**.

Important

Make sure that you extract the files to a different folder in step **4**.

Auto-provision a Citrix ADC VPX instance on Hyper-V

Auto-provisioning of Citrix ADC VPX instance is optional. If auto-provisioning is not done, the virtual appliance provides an option to configure the IP address and so on.

To auto-provision Citrix ADC VPX instance on Hyper-V, follow these steps.

1. Create an ISO9660 compliant ISO image using the xml file as depicted in the example. Make sure that the name of the xml file is **userdata**.

You can create an ISO file from XML file using:

- Any image processing tool such as PowerISO.
- `mkisofs` command in Linux.

```
1 <?xml version="1.0" encoding="UTF-8" standalone="no"?>
2
3 <Environment xmlns:oe="http://schemas.dmtf.org/ovf/environment/1"
4
5 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
6
7 oe:id=""
```

```
8
9  xmlns="http://schemas.dmtf.org/ovf/environment/1"
10
11 <PlatformSection>
12
13 <Kind>HYPER-V</Kind>
14
15 <Version>2013.1</Version>
16
17 <Vendor>CITRIX</Vendor>
18
19 <Locale>en</Locale>
20
21 </PlatformSection>
22
23 <PropertySection>
24
25 <Property oe:key="com.citrix.netscaler.ovf.version" oe:value="1.0"
26   />
27 <Property oe:key="com.citrix.netscaler.platform" oe:value="NS1000V
28   "/>
29 <Property oe:key="com.citrix.netscaler.orch\_env" oe:value="cisco-
30   orch-env"/>
31 <Property oe:key="com.citrix.netscaler.mgmt.ip" oe:value="
32   10.102.100.122"/>
33 <Property oe:key="com.citrix.netscaler.mgmt.netmask" oe:value="
34   255.255.255.128"/>
35 <Property oe:key="com.citrix.netscaler.mgmt.gateway" oe:value="
36   10.102.100.67"/></PropertySection>
37 </Environment>
```

2. Copy the ISO image to hyper-v server.
3. Select the virtual appliance that you imported, and then on the **Action** menu, select **Settings**. You can also select the virtual appliance and then right click and select **Settings**. The **Settings** window for the selected virtual appliance is displayed.
4. In the **Settings** window, under the hardware section, click **IDE Controller**.
5. In the right window pane, select **DVD Drive** and click **Add**. The DVD Drive is added under the **IDE Controller** section in the left window pane.
6. Select the **DVD Drive** added in step 5. In the right window pane, select the **Image file radio** button and click **Browse** and select the ISO image that you copied on Hyper-V server, in step 2.
7. Click **Apply**.

Note:

The virtual appliance instance comes up in the default IP address, when:

- The DVD drive is attached and the ISO file is not provided.
- The ISO file does not include the user data file.
- The user data file name or format is not correct.

To configure virtual NICs on the Citrix ADC VPX instance, follow these steps:

1. Select the virtual appliance that you imported, and then on the **Action** menu, select **Settings**.
2. In the **Settings for <virtual appliance name>** dialog box, click **Add Hardware** in the left pane.
3. In the right pane, from the list of devices, select **Network Adapter**.
4. Click **Add**.
5. Verify that **Network Adapter (not connected)** appears in the left pane.
6. Select the network adapter in the left pane.
7. In the right pane, from the **Network** menu, select the virtual network to connect the adapter to.
8. To select the virtual network for other network adapters that you want to use, repeat steps **6** and **7**.
9. Click **Apply**, and then click **OK**.

To configure the Citrix ADC VPX instance:

1. Right-click the virtual appliance that you previously installed, and then select **Start**.
2. Access the console by double-clicking the virtual appliance.
3. Type the Citrix ADC IP address, subnet mask, and gateway for your virtual appliance.

You have completed the basic configuration of your virtual appliance. Type the IP address in a Web browser to access the virtual appliance.

Note:

You can also use virtual machine (VM) template to provision Citrix ADC VPX instance using SCVMM.

If you use Microsoft Hyper-V NIC teaming solution with NetScaler VPX instances, see article [CTX224494](#) for more information.

Install a Citrix ADC VPX instance on Linux-KVM platform

June 20, 2024

To set up a Citrix ADC VPX for the Linux-KVM platform, you can use the graphical Virtual Machine Manager (Virtual Manager) application. If you prefer the Linux-KVM command line, you can use the `virsh` program.

The host Linux operating system must be installed on suitable hardware by using virtualization tools such as KVM Module and QEMU. The number of virtual machines (VMs) that can be deployed on the hypervisor depends on the application requirement and the chosen hardware.

After you provision a Citrix ADC VPX instance, you can add more interfaces.

Limitations and usage guidelines

General recommendations

To avoid unpredictable behavior, apply the following recommendations:

- Do not change the MTU of the VNet interface associated with the VPX VM. Shut down the VPX VM before modifying any configuration parameters, such as Interface modes or CPU.
- Do not force a shutdown of the VPX VM. That is, do not use the **Force off** command.
- Any configurations done on the host Linux might or might not be persistent, depending on your Linux distribution settings. You can choose to make these configurations persistent to ensure consistent behavior across reboots of host Linux operating system.
- The Citrix ADC package has to be unique for each of the Citrix ADC VPX instance provisioned.

Limitations

- Live migration of a VPX instance that runs on KVM is not supported.

Prerequisites for installing a Citrix ADC VPX instance on Linux-KVM platform

September 12, 2024

Check the minimum system requirements for a Linux-KVM server running on a Citrix ADC VPX instance.

CPU requirement:

- 64-bit x86 processors with the hardware virtualization feature included in Intel VT-X processors.

To test whether your CPU supports the Linux host, enter the following command at the host Linux shell prompt:

```
1 *.egrep '^flags.*(vmx|svm)' /proc/cpuinfo*
```


If the **BIOS** settings for the preceding extension are disabled, you must enable them in the BIOS.

- Provide at least 2 CPU cores to Host Linux.
- There is no specific recommendation for processor speed, but higher the speed, the better the performance of the VM application.

Memory (RAM) requirement:

Minimum 4 GB for the host Linux kernel. Add more memory as required by the VMs.

Hard disk requirement:

Calculate the space for Host Linux kernel and VM requirements. A single Citrix ADC VPX VM requires 20 GB of disk space.

Software requirements

The Host kernel used must be a 64-bit Linux kernel, release 2.6.20 or later, with all virtualization tools. Citrix recommends newer kernels, such as 3.6.11-4 and later.

Many Linux distributions such as Red Hat, CentOS, and Fedora, have tested kernel versions and associated virtualization tools.

Guest VM hardware requirements

Citrix ADC VPX supports IDE and virtIO hard disk type. The Hard Disk Type has been configured in the XML file, which is a part of the Citrix ADC package.

Networking requirements

Citrix ADC VPX supports virtIO para-virtualized, SR-IOV, and PCI Passthrough network interfaces.

For more information about the supported network interfaces, see:

- [Provision the Citrix ADC VPX instance by using the Virtual Machine Manager](#)
- [Configure a Citrix ADC VPX instance to use SR-IOV network interfaces](#)
- [Configure a Citrix ADC VPX instance to use PCI passthrough network interfaces](#)

Source Interface and Modes

The source device type can be either Bridge or MacVTap. In MacVTap, four modes are possible - VEPA, Bridge, Private, and Pass-through. Check the types of interfaces that you can use and the supported traffic types, as per the following:

Bridge:

- Linux Bridge.
- `Ebtables` and `iptables` settings on host Linux might filter the traffic on the bridge if you do not choose the correct setting or disable `IPtable` services.

MacVTap (VEPA mode):

- Better performance than a bridge.
- Interfaces from the same lower device can be shared across the VMs.
- Inter-VM communication using the same lower device is possible only if the upstream or downstream switch supports VEPA mode.

MacVTap (private mode):

- Better performance than a bridge.
- Interfaces from the same lower device can be shared across the VMs.
- Inter-VM communication using the same lower device is not possible.

MacVTap (bridge mode):

- Better as compared to bridge.
- Interfaces out of the same lower device can be shared across the VMs.
- Inter-VM communication using the same lower device is possible, if the lower device link is UP.

MacVTap (Pass-through mode):

- Better as compared to bridge.
- Interfaces out of the same lower device cannot be shared across the VMs.
- Only one VM can use the lower device.

Note:

For best performance by the VPX instance, ensure that the `gro` and `lro` capabilities are switched off on the source interfaces.

Properties of source interfaces

Make sure that you switch off the generic-receive-offload (`gro`) and large-receive-offload (`lro`) capabilities of the source interfaces. To switch off the `gro` and `lro` capabilities, run the following commands at the host Linux shell prompt.

```
ethtool -K eth6 gro off
ethool -K eth6 lro off
```

Example:

```
1 [root@localhost ~]# ethtool -K eth6
2
3         Offload parameters for eth6:
4
5             rx-checksumming: on
6
7             tx-checksumming: on
8
9         scatter-gather: on
10
11        tcp-segmentation-offload: on
12
13        udp-fragmentation-offload: off
14
15        generic-segmentation-offload: on
16
17        generic-receive-offload: off
18
19        large-receive-offload: off
20
21        rx-vlan-offload: on
22
23        tx-vlan-offload: on
24
25        ntuple-filters: off
26
27        receive-hashing: on
28
29 [root@localhost ~]#
```

Example:

If the host Linux bridge is used as a source device, as in the following example, and **lro** capabilities must be switched off on the VNet interfaces, which are the virtual interfaces connecting the host to the guest VMs.

```
1 [root@localhost ~]# brctl show eth6_br
2
3 bridge name      bridge id          STP enabled interfaces
4
5 eth6_br          8000.00e0ed1861ae  no          eth6
6
7                                     vnet0
8
9                                     vnet2
10
11 [root@localhost ~]#
```

In the preceding example, the two virtual interfaces are derived from the eth6_br and are represented as vnet0 and vnet2. Run the following commands to switch off **gro** and **lro** capabilities on these interfaces.

```
1 ethtool -K vnet0 gro off
2 ethtool -K vnet2 gro off
3 ethtool -K vnet0 lro off
4 ethtool -K vnet2 lro off
```

Promiscuous mode

The promiscuous mode must be enabled for the following features to work:

- L2 mode
- Multicast traffic processing
- Broadcast
- IPV6 traffic
- virtual MAC
- Dynamic routing

Use the following command to enable the promiscuous mode.

```
1 [root@localhost ~]# ifconfig eth6 promisc
2 [root@localhost ~]# ifconfig eth6
3 eth6      Link encap:Ethernet  HWaddr 78:2b:cb:51:54:a3
4           inet6 addr: fe80::7a2b:cbff:fe51:54a3/64 Scope:Link
5           UP BROADCAST RUNNING PROMISC MULTICAST  MTU:9000  Metric:1
6           RX packets:142961 errors:0 dropped:0 overruns:0 frame:0
7           TX packets:2895843 errors:0 dropped:0 overruns:0 carrier:0
8           collisions:0 txqueuelen:1000
9           RX bytes:14330008 (14.3 MB)  TX bytes:1019416071 (1.0 GB)
10
11 [root@localhost ~]#
```

Module required

For better network performance, make sure the `vhost_net` module is present in the Linux host. To check the existence of `vhost_net` module, run the following command on the Linux host:

```
1 lsmod | grep "vhost\_net"
```

If `vhost_net` is not yet running, enter the following command to run it:

```
1 modprobe vhost\_net
```

Provision the Citrix ADC VPX instance by using OpenStack

September 6, 2024

You can provision a Citrix ADC VPX instance in an OpenStack environment either by using the **Nova boot** command (OpenStack CLI) or Horizon (OpenStack dashboard) .

Provisioning a VPX instance, optionally involves using data from the config drive. Config drive is a special configuration drive that attaches to the instance as a CD-ROM device when it boots. This configuration drive can be used to pass networking configuration such as management IP address, network mask, default gateway, and to inject customer scripts.

In a Citrix ADC appliance, the default authentication mechanism is password based. Now, the SSH key-pair authentication mechanism is supported for Citrix ADC VPX instances on the OpenStack environment.

The key-pair (public key and private key) is generated before using the Public Key Cryptography mechanism. You can use different mechanisms, such as Horizon, Puttygen.exe for Windows, and [ssh-keygen](#) for the Linux environment, to generate the key pair. Refer to online documentation of respective mechanisms for more information about generating key pair.

Once a key pair is available, copy the private key to a secure location to which authorized persons have access. In OpenStack, public key can be deployed on a VPX instance by using the Horizon or Nova boot command. When a VPX instance is provisioned by using OpenStack, it first detects that the instance is booting in an OpenStack environment by reading a specific BIOS string. This string is “OpenStack Foundation” and for Red Hat Linux distributions it is stored in `/etc/nova/release`. This is a standard mechanism that is available in all OpenStack implementations based on the KVM hypervisor platform. The drive must have a specific OpenStack label.

If the config drive is detected, the instance attempts to read the network configuration, custom scripts, and SSH key pair if provided.

User data file

The Citrix ADC VPX instance uses a customized OVF file, also known as the user data file, to inject network configuration, custom scripts. This file is provided as part of config drive. Here is an example of a customized OVF file.

```
1  `` `
2  <?xml version="1.0" encoding="UTF-8" standalone="no"?>
3  <Environment xmlns:oe="http://schemas.dmtf.org/ovf/environment/1"
4  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
5  oe:id=""
6  xmlns="http://schemas.dmtf.org/ovf/environment/1"
```

```
7 xmlns:cs="http://schemas.citrix.com/openstack">
8 <PlatformSection>
9 <Kind></Kind>
10 <Version>2016.1</Version>
11 <Vendor>VPX</Vendor>
12 <Locale>en</Locale>
13 </PlatformSection>
14 <PropertySection>
15 <Property oe:key="com.citrix.netscaler.ovf.version" oe:value="1.0"/>
16 <Property oe:key="com.citrix.netscaler.platform" oe:value="NSVPX"/>
17 <Property oe:key="com.citrix.netscaler.orch_env" oe:value="openstack-
    orch-env"/>
18 <Property oe:key="com.citrix.netscaler.mgmt.ip" oe:value="10.1.2.22"/>
19 <Property oe:key="com.citrix.netscaler.mgmt.netmask" oe:value="
    255.255.255.0"/>
20 <Property oe:key="com.citrix.netscaler.mgmt.gateway" oe:value="10.1.2.1
    "/>
21 </PropertySection>
22 <cs:ScriptSection>
23 <cs:Version>1.0</cs:Version>
24 <ScriptSettingSection xmlns="http://schemas.citrix.com/openstack"
    xmlns:i="http://www.w3.org/2001/XMLSchema-instance">
25 <Scripts>
26 <Script>
27 <Type>shell</Type>
28 <Parameter>X Y</Parameter>
29 <Parameter>Z</Parameter>
30 <BootScript>before</BootScript>
31 <Text>
32 <pre>#!/bin/bash
33 echo "Hi, how are you" $1 $2 >> /var/sample.txt
34 </pre>
35 </Script>
36 <Script>
37 <Type>python</Type>
38 <BootScript>after</BootScript>
39 <Text>
40 <pre>#!/bin/python
41 print("Hello");
42 </pre>
43 </Script>
44 <Script>
45 <Type>perl</Type>
46 <BootScript>before</BootScript>
47 <Text>
48 <pre>!/usr/bin/perl
49 my $name = "VPX";
50 print "Hello, World $name !\n" ;
51 </pre>
52 </Script>
53 <Script>
54 <Type>nscli</Type>
55 <BootScript>after</BootScript>
```

```

56         <Text>
57             add vlan 33
58     bind vlan 33 -ifnum 1/2
59         </Text>
60     </Script>
61 </Scripts>
62 </ScriptSettingSection>
63 </cs:ScriptSection>
64 </Environment>
65 ```

```

In the OVF file preceding “PropertySection” is used for NetScaler networking configuration while `<cs:ScriptSection>` is used to enclose all scripts. `<Scripts></Scripts>` tags are used to bundle all scripts together. Each script is defined in between `<Script>` `</Script>` tags. Each script tag has following fields/tags:

- a) `<Type>`: Specifies value for script type. Possible values: Shell/Perl/Python/NSLCL (for NetScaler CLI scripts)
- b) `<Parameter>`: Provides parameters to the script. Each script can have multiple `<Parameter>` tags.
- c) `<BootScript>`: Specifies script execution point. Possible values for this tag: before/after. “before” specifies script is run before PE comes up. “after” specifies that the script will be run after PE comes up.
- d) `<Text>`: Pastes content of a script.

Note:

Currently the VPX instance does not take care of sanitization of scripts. As an administrator, you must check the validity of the script.

Not all sections need to be present. Use an empty “PropertySection” to only define scripts to run on first boot or an empty `<cs:ScriptSection>` to only define networking configuration.

After the required sections of the OVF file (user data file) are populated, use that file to provision the VPX instance.

Network configuration

As part of the network configuration, the VPX instance reads:

- Management IP address
- Network mask
- Default gateway

After the parameters are successfully read, they are populated in the NetScaler configuration, to allow managing the instance remotely. If the parameters are not read successfully or the config drive is not available, the instance transitions to the default behavior, which is:

- The instance attempts to retrieve the IP address information from DHCP.
- If DHCP fails or times-out, the instance comes up with the default network configuration (192.168.100.1/16).

Customer script

The VPX instance allows to run a custom script during initial provisioning. The appliance supports script of type Shell, Perl, Python, and Citrix ADC CLI commands.

SSH key pair authentication

The VPX instance copies public key, available within the configuration drive as part of instance meta data, into its “authorized_keys”file. This allows the user to access the instance with private key.

Note:

When an SSH key is provided, the default credentials (nsroot/nsroot) no longer work, if password-based access is needed, log on with the respective SSH private key and manually set a password.

Before you begin

Before you provision a VPX instance on OpenStack environment, extract the `.qcow2` file from the `.tgz` file and build

An OpenStack image from the qcow2 image. Follow these steps:

1. Extract the `.qcow2` file from the `.tgz` file by typing the following command

```
1 tar xvzf <TAR file>
2 tar xvzf <NSVPX-KVM-12.0-26.2_nc.tgz>
3 NSVPX-KVM.xml
4 NSVPX-KVM-12.0-26.2_nc.qcow2
```

2. Build an OpenStack image using the `.qcow2` file extracted in step 1 by typing the following command.

```
1 openstack image create --container-format bare --property
  hw_disk_bus=ide --disk-format qcow2 --file <path to qcow2 file>
  --public <name of the OpenStack image>
2
3 glance image-create --name="NS-VPX-12-0-26-2" --property
  hw_disk_bus=ide --ispublic=
4 true --container-format=bare --disk-format=qcow2< NSVPX-KVM
  -12.0-26.2_nc.qcow2
```


Figure 1: The following illustration provides a sample output for the glance image-create command.

Field	Value
checksum	154ade3fc7dca7d1706b1d03d7d97552
container_format	bare
created_at	2017-03-13T08:52:31Z
disk_format	qcow2
file	/v2/images/322c1e0f-cce8-4b7b-b53e-bd8152c388ed/file
id	322c1e0f-cce8-4b7b-b53e-bd8152c388ed
min_disk	0
min_ram	0
name	VPX-KVM-12.0-26.2
owner	58d17d81df5d4406afbb4fdab3a58d79
properties	hw_disk_bus='ide'
protected	False
schema	/v2/schemas/image
size	784338944
status	active
updated_at	2017-03-13T08:52:43Z
virtual_size	None
visibility	public

Provisioning the VPX instance


You can provision a VPX instance in two ways by using one of the options:

- Horizon (OpenStack dashboard)
- Nova boot command (OpenStack CLI)

Provision a VPX instance by using the OpenStack dashboard

Follow these steps to provision the VPX instance by using Horizon:

1. Log on to the OpenStack dashboard.
2. In the Project panel on the left hand side of the dashboard, select **Instances**.
3. In the Instances panel, click **Launch Instance** to open the Instance Launching wizard.



4. In the Launch Instance wizard, fill in the details, like:

- a) Instance Name
- b) Instance Flavor
- c) Instance Count
- d) Instance Boot Source
- e) Image Name

Launch Instance ✕

Details *
Access & Security *
Networking *
Post-Creation
Advanced Options

Availability Zone:

nova ▼

Instance Name: *

NSVPX_10_1

Flavor: *

m1.medium ▼

Instance Count: *

1

Instance Boot Source: *

Boot from image ▼

Image Name:

NS-VPX-10-1-130-11 (20.0 GB) ▼

Specify the details for launching an instance.

The chart below shows the resources used by this project in relation to the project's quotas.

Flavor Details

Name	m1.medium
VCPUs	2
Root Disk	40 GB
Ephemeral Disk	0 GB
Total Disk	40 GB
RAM	4,096 MB

Project Limits

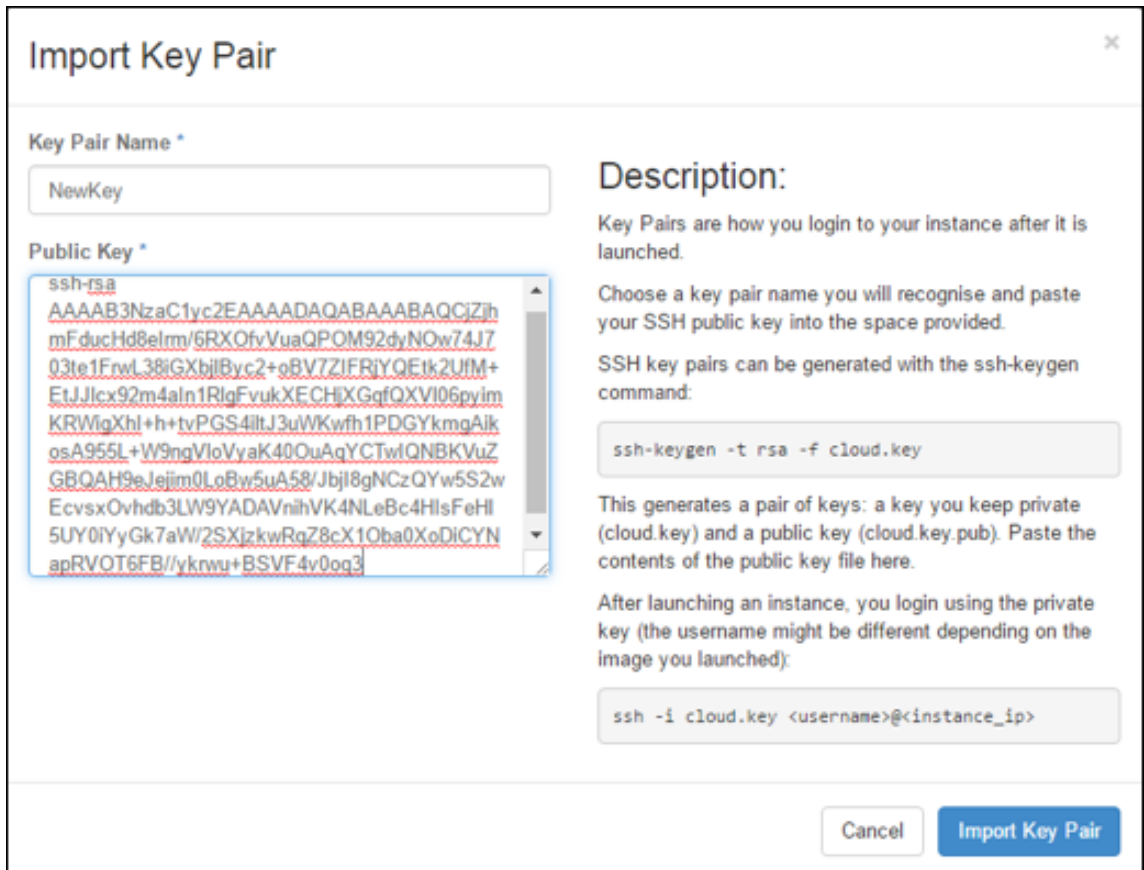
Number of Instances 6 of 10 Used

Number of VCPUs 12 of 20 Used

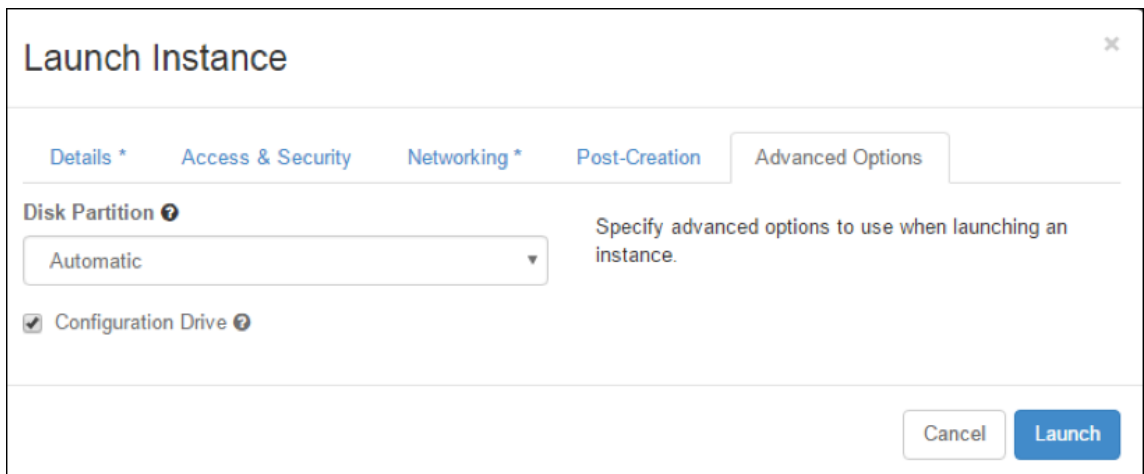
Total RAM 24,576 of 51,200 MB Used

Cancel
Launch

5. Deploy a new key pair or an existing key pair through Horizon by completing the following steps:
 - a) If you don't have an existing key pair, create the key by using any existing mechanisms. If you've an existing key, skip this step.
 - b) Copy the content of public key.
 - c) Go to **Horizon > Instances > Create New Instances**.
 - d) Click **Access & Security**.
 - e) Click the + sign next to the **Key Pair** drop-down menu and provide values for shown parameters.
 - f) Paste public key content in *Public key* box, give a name to the key and click **Import Key Pair**.



6. Click the **Post Creation** tab in the wizard. In Customization Script, add the content of the user data file. The user data file contains the IP address, Netmask and Gateway details, and customer scripts of the VPX instance.
7. After a key pair is selected or imported, check config-drive option and click **Launch**.



Provision the VPX instance by using OpenStack CLI

Follow these steps to provision a VPX instance by using OpenStack CLI.

1. To create an image from qcow2, type the following command:

```
openstack image create --container-format bare --property hw_disk_bus=ide --diskformat qcow2 --file NSVPX-OpenStack.qcow2 --public VPX-ToT-Image
```

2. To select an image for creating an instance, type the following command:

```
openstack image list | more
```

3. To create an instance of a particular flavor, type the following command to choose a flavor ID/Name of from a list:

```
openstack flavor list
```

4. To attach a NIC to a particular network, type the following command to choose a network ID from a network list:

```
openstack network list
```

5. To create an instance, type the following command:

```
1 openstack server create --flavor FLAVOR_ID --image IMAGE_ID --key-name KEY_NAME
2 --user-data USER_DATA_FILE_PATH --config-drive True --nic net-id=net-uuid
3 INSTANCE_NAME
4 openstack server create --image VPX-ToT-Image --flavor m1.medium
  --user-data
5 ovf.xml --config-drive True --nic net-id=2734911b-ee2b-48d0-a1b6-3
  efd44b761b9
6 VPX-ToT
```

Figure 2: The following illustration provides a sample output.

Field	Value
OS-DCF:diskConfig	MANUAL
OS-EXT-AZ:availability_zone	
OS-EXT-SRV-ATTR:host	None
OS-EXT-SRV-ATTR:hypervisor_hostname	None
OS-EXT-SRV-ATTR:instance_name	instance-000001c2
OS-EXT-STS:power_state	0
OS-EXT-STS:task_state	scheduling
OS-EXT-STS:vm_state	building
OS-SRV-USG:launched_at	None
OS-SRV-USG:terminated_at	None
accessIPv4	
accessIPv6	
addresses	
adminPass	pFVvMtq7N8Z6
config_drive	True
created	2017-03-13T10:32:59Z
flavor	m1.medium (3)
hostId	
id	a1fe991e-3604-43a0-9dd6-59fa0f3749df
image	VPX-ToT-Image (f0c2f9d1-08f2-4b2e-9943-2ee6bc2edbc7)
key_name	None
name	VPX-ToT
os-extended-volumes:volumes_attached	[]
progress	0
project_id	58d17d81df5d4406afbb4fdab3a58d79
properties	
security_groups	[{'name': 'u'default'}]
status	BUILD
updated	2017-03-13T10:33:00Z
user_id	a6347b33916b4eb1b1f76360a9c8f935

Provision the Citrix ADC VPX instance by using the Virtual Machine Manager

January 9, 2025

The Virtual Machine Manager is a desktop tool for managing VM guests. It enables you to create new VM guests and various types of storage, and manage virtual networks. You can access the graphical console of VM guests with the built-in VNC viewer and view performance statistics, either locally or remotely.

After installing your preferred Linux distribution, with KVM virtualization enabled, you can proceed with provisioning virtual machines.

While using the Virtual Machine Manager to provision a Citrix ADC VPX instance, you have two options:

- Enter the IP address, gateway, and netmask manually
- Assign the IP address, gateway, and netmask automatically (auto-provisioning)

You can use two kinds of images to provision a Citrix ADC VPX instance:

- RAW
- QCOW2

You can convert a Citrix ADC VPX RAW image to a QCOW2 image and provision the Citrix ADC VPX instance. To convert the RAW image to a QCOW2 image, type the following command:

```
qemu-img convert -O qcow2 original-image.raw image-converted.qcow2
```

For example:

```
qemu-img convert -O qcow2 NSVPX-KVM-11.1-12.5_nc.raw NSVPX-KVM-11.1-12.5_nc.qcow2
```

A typical Citrix ADC VPX deployment on KVM includes the following steps:

- Checking prerequisites for auto-provisioning a NetScaler VPX instance
- Provisioning the NetScaler VPX instance by using a RAW image
- Provisioning the NetScaler VPX instance by using a QCOW2 image
- Adding more interfaces to a VPX instance by using a virtual machine manager

Check the prerequisites for auto-provisioning a Citrix ADC VPX instance

Auto-provisioning is an optional feature, and it involves using data from the CDROM drive. If this feature is enabled, you need not enter the management IP address, network mask, and default gateway of the Citrix ADC VPX instance during initial setup.

You need to complete the following tasks before you can auto-provision a VPX instance:

1. Create a customized Open Virtualization Format (OVF) XML file or user data file.
2. Convert the OVF file into an ISO image by using an online application (for example PowerISO).
3. Mount the ISO image on the KVM host by using any secure copy (SCP)-based tools.

Sample OVF XML file:

Here's is an example of the contents an OVF XML file, which you can use as a sample to create your file.

```
1 <?xml version="1.0" encoding="UTF-8" standalone="no"?>
2
3 <Environment xmlns:oe="`http://schemas.dmtf.org/ovf/environment/1"`
4
5 xmlns:xsi="`http://www.w3.org/2001/XMLSchema-instance"`
6
7 oe:id=""
8
9 xmlns="`http://schemas.dmtf.org/ovf/environment/1"`
10
11 xmlns:cs="`http://schemas.citrix.com/openstack">`
12
13 <PlatformSection>
14
15 <Kind></Kind>
```

```
16
17 <Version>2016.1</Version>
18
19 <Vendor>VPX</Vendor>
20
21 <Locale>en</Locale>
22
23 </PlatformSection>
24
25 <PropertySection>
26
27 <Property oe:key="com.citrix.netscaler.ovf.version" oe:value="1.0"/>
28
29 <Property oe:key="com.citrix.netscaler.platform" oe:value="NSVPX"/>
30
31 <Property oe:key="com.citrix.netscaler.orch\_env" oe:value="KVM"/>
32
33 <Property oe:key="com.citrix.netscaler.mgmt.ip" oe:value="10.1.2.22"/>
34
35 <Property oe:key="com.citrix.netscaler.mgmt.netmask" oe:value="
    255.255.255.0"/>
36
37 <Property oe:key="com.citrix.netscaler.mgmt.gateway" oe:value="10.1.2.1
    "/>
38
39 </PropertySection>
40
41 </Environment>
```

In the OVF XML file preceding, “PropertySection” is used for NetScaler networking configuration. When you create the file, specify values for the parameters that are highlighted at the end of the example:

- Management IP address
- Netmask
- Gateway

Important:

If the OVF file is not properly XML formatted, the VPX instance is assigned the default network configuration, not the values specified in the file.

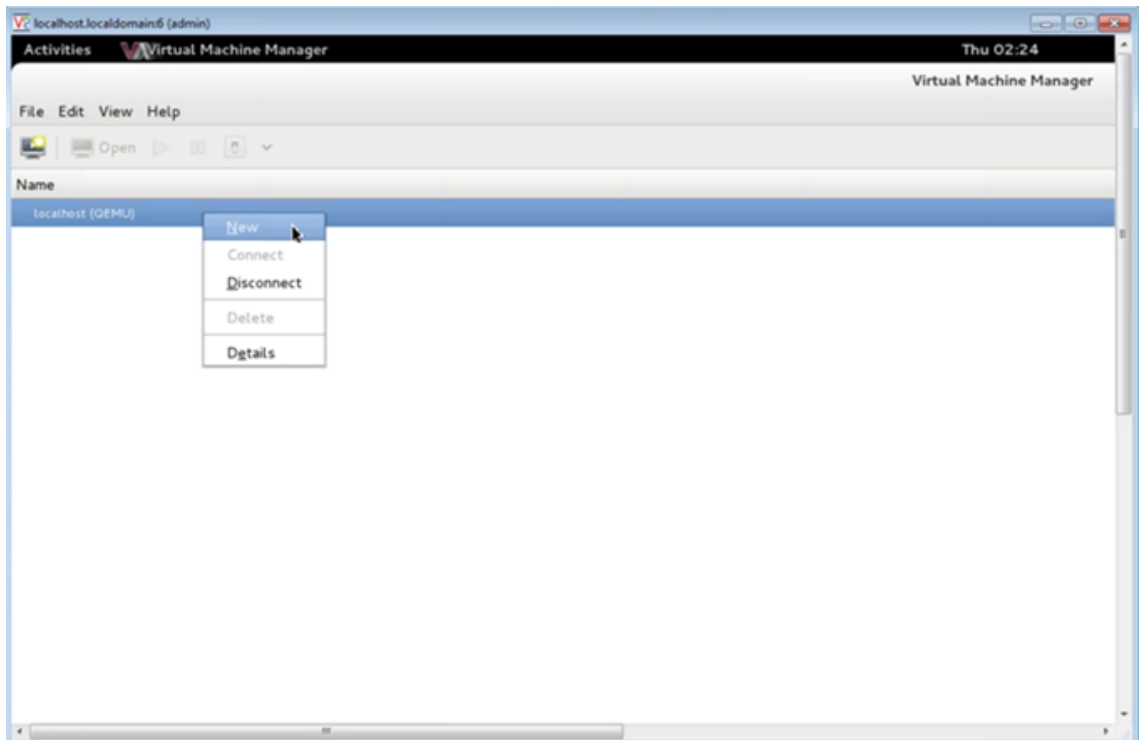
Provision the Citrix ADC VPX instance by using a RAW image

The Virtual Machine Manager enables you to provision a Citrix ADC VPX instance by using a RAW image.


To provision a Citrix ADC VPX instance by using the Virtual Machine Manager, follow these steps:

1. Open the Virtual Machine Manager (**Application > System Tools > Virtual Machine Manager**) and enter the logon credentials in the **Authenticate** window.


2. Click the  icon or right-click **localhost (QEMU)** to create a new Citrix ADC VPX instance.




3. In the **Name** text box, enter a name for the new VM (for example, NetScaler-VPX).
4. In the **New VM** window, under “Choose how you would like to install the operating system,” select **Import existing disk image**, and then and click **Forward**.




5. In the **Provide the existing storage path** field, navigate the path to the image. Choose the OS type as UNIX and Version as FreeBSD 6.x. Then, click **Forward**.




6. Under **Choose Memory and CPU** settings select the following settings, and then click **Forward**:
- Memory (RAM)—2048 MB
 - CPUs—2



7. Select the **Customize configuration before install** checkbox. Optionally, under **Advanced options** you can customize the MAC address. Make sure the **Virt Type** selected is KVM and the Architecture selected is x86_64. Click **Finish**.



8. Select a NIC and provide the following configuration:
- Source device—ethX macvtap or Bridge
 - Device model—virtio
 - Source mode—Bridge



9. Click **Apply**.
10. If you want to auto-provision the VPX instance, see the section **Enabling Auto-Provisioning by Attaching a CDROM Drive** in this document. Otherwise, click **Begin Installation**. After you have provisioned the Citrix ADC VPX on KVM, you can add more interfaces.

Provision the Citrix ADC VPX instance by using a QCOW2 image

Using the Virtual Machine Manager, you can provision the Citrix ADC VPX instance by using a QCOW2 image.


To provision a Citrix ADC VPX instance by using a QCOW2 image, follow these steps:

1. Follow **step 1 to step 8** in [Provision the Citrix ADC VPX instance by using a RAW image](#).

Note:

Ensure that you select the **qcow2** image in **step 5**.


2. Select **Disk 1** and click **Advanced options**.
3. Select **qcow2** from the Storage format drop-down list.



4. Click **Apply**, and then click **Begin Installation**. After you have provisioned the Citrix ADC VPX on KVM, you can add more interfaces.

Enable auto-provisioning by attaching a CDROM drive

1. Click Add **Hardware** > **Storage** > **Device type** > **CDROM device**.
2. Click **Manage** and select the correct ISO file that you mounted in the “Prerequisites for Auto-Provisioning a Citrix ADC VPX Instance” section, and click **Finish**. A new CDROM under Resources on your Citrix ADC VPX instance is created.



3. Power on the VPX instance, and it auto-provisions with the network configuration provided in the OVF file, as shown in the example screen capture.

```

File Virtual Machine View Send Key

Aug 11 10:14:55 <local0.alert> ns restart[25781]: Restart: /netscaler/nsstart.sh
exited normally. Exit code (0)
Aug 11 10:14:55 <local0.alert> ns restart[25781]: Successfully deregistered with
h Pitboss ...

login: nsroot
Password:
Aug 11 10:15:04 <auth.notice> ns login: ROOT LOGIN (nsroot) ON ttyv0
Copyright (c) 1992-2013 The FreeBSD Project.
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
The Regents of the University of California. All rights reserved.

Done
> sh ip
      Ippaddress      Traffic Domain  Type      Mode      Arp      Icmp
      Userver  State
      -----
1)    10.1.2.22      0              NetScaler IP  Active    Enabled  Enab
led NA      Enabled
Done
> Aug 11 10:15:13 <local0.alert> ns restart[25781]: Nsshutdown lock released !

```

4. If auto-provision fails, the instance comes up with the default IP address (192.168.100.1). In that case, you must complete the initial configuration manually. For more information, see [Configure the ADC for the first time](#).


Add more interfaces to the Citrix ADC VPX instance by using the Virtual Machine Manager


After you have provisioned the NetScaler VPX instance on KVM, you can add more interfaces.

To add more interfaces, follow these steps.

1. Shut down the NetScaler VPX instance running on the KVM.
2. Right-click the VPX instance and choose **Open** from the pop-up menu.



3. Click the  icon in the header to view the virtual hardware details.
4. Click **Add Hardware**. In the **Add New Virtual Hardware window**, select **Network** from the navigation menu.




5. In the **Host Device** field, select the physical interface type. The host device type can be either Bridge or MacVTap. In case of MacVTap, four modes possible are VEPA, Bridge, Private, and Pass-through.

a) For Bridge


- i. Host device—Select the “Specify shared device name” option.
- ii. Provide the Bridge name that is configured in the KVM host.

Note:


Make sure that you have configured a Linux bridge in the KVM host, bound the physical interface to the bridge, and put the bridge in the UP state.



- iii. Device model—`virtio`.
 - iv. Click **Finish**.
- b) For MacVTap
- i. Host device—Select the physical interface from the menu.
 - ii. Device model—`virtio`.



iii. Click **Finish**. You can view the newly added NIC in the navigation pane.



- iv. Select the newly added NIC and select the Source mode for this NIC. The available modes are VEPA, Bridge, Private, and Passthrough. For more details on the interface and modes, see Source Interface and Modes.
- v. Click **Apply**.

6. If you want to auto-provision the VPX instance, see the section “Adding a Config Drive to Enable Auto-Provisioning” in this document. Otherwise, power on the VPX instance to complete the initial configuration manually.

Important:

Interface parameter configurations such as speed, duplex, and autonegotiation are not supported.

Configure a Citrix ADC VPX instance to use SR-IOV network interfaces

September 6, 2024

You can configure a Citrix ADC VPX instance running on Linux-KVM platform using single root I/O virtualization (SR-IOV) with the following NICs:

- Intel 82599 10G
- Intel X710 10G
- Intel XL710 40G
- Intel X722 10G

This section describes how to:

- Configure a Citrix ADC VPX Instance to Use SR-IOV Network Interface
- Configure Static LA/LACP on the SR-IOV Interface
- Configure VLAN on the SR-IOV Interface

Limitations

Keep the limitations in mind while using Intel 82599, X710, XL710, and X722 NICs. The following features not supported.

Limitations for Intel 82599 NIC:

- L2 mode switching.
- Admin partitioning (shared VLAN mode).
- High availability (active-active mode).
- Jumbo frames.
- IPv6: You can configure only up to 30 unique IPv6 addresses in a VPX instance if you've at least one SR-IOV interface.
- VLAN configuration on Hypervisor for SRIOV VF interface through `ip link` command is not supported.
- Interface parameter configurations such as speed, duplex, and autonegotiations are not supported.

Limitations for Intel X710 10G, Intel XL710 40G, and Intel X722 10G NICs:

- L2 mode switching.
- Admin partitioning (shared VLAN mode).
- In a cluster, Jumbo frames are not supported when the XL710 NIC is used as a data interface.
- Interface list reorders when interfaces are disconnected and reconnected.
- Interface parameter configurations such as speed, duplex, and auto negotiations are not supported.
- Interface name is 40/X for Intel X710 10G, Intel XL710 40G, and Intel X722 10G NICs
- Up to 16 Intel XL710/X710/X722 SRIOV or PCI passthrough interfaces can be supported on a VPX instance.

Note:

For Intel X710 10G, Intel XL710 40G, and Intel X722 10G NICs to support IPv6, you need to enable trust mode on the Virtual Functions (VFs) by typing the following command on the KVM host:

```
# ip link set <PNIC> <VF> trust on
```

Example:

```
# ip link set ens785f1 vf 0 trust on
```

Prerequisites

Before you configure a Citrix ADC VPX instance to use SR-IOV network interfaces, complete the following prerequisite tasks. See the NIC column for details about how to complete the corresponding tasks.

Task	Intel 82599 NIC	Intel X710, XL710, and X722 NICs
1. Add the NIC to the KVM host.	-	-
1. Download and install the latest Intel driver.	IXGBE driver	I40E driver
1. Block list the driver on the KVM host.	Add the following entry in the /etc/modprobe.d/blacklist.conf file: <code>blacklist ixgbev</code> . Use IXGBE driver version 4.3.15 (recommended).	Add the following entry in the /etc/modprobe.d/blacklist.conf file: <code>blacklist i40evf</code> . Use i40e driver version 2.0.26 (recommended).

Task	Intel 82599 NIC	Intel X710, XL710, and X722 NICs
<p>4.Enable SR-IOV Virtual Functions (VFs) on the KVM host. In both the commands in the next two columns: <code>number_of_VFs</code> = the number of Virtual VFs that you want to create. <code>device_name</code> = the interface name.</p> <p>1. Make the VFs persistent by adding the commands that you used to create VFs, to the rc.local file.</p>	<p>If you are using earlier version of kernel 3.8, then add the following entry to the <code>/etc/modprobe.d/ixgbe</code> file and restart the KVM host:</p> <pre>options ixgbe max_vfs =<number_of_VFs>.</pre> <p>If you are using kernel 3.8 version or later, create VFs using the following command:</p> <pre>echo <number_of_VFs> > /sys/class/net/<device_name>/device/sriov_numvfs.</pre> <p>See example in figure 1.</p> <p>See example in figure 3.</p>	<p>If you are using earlier version of kernel 3.8, then add the following entry to the <code>/etc/modprobe.d/i40e.conf</code> file and restart the KVM host:</p> <pre>options i40e max_vfs =<number_of_VFs>.</pre> <p>If you are using kernel 3.8 version or later, create VFs using the following command:</p> <pre>echo<number_of_VFs> > /sys/class/net/<device_name>/device/sriov_numvfs.</pre> <p>See example in figure 2.</p> <p>See example in figure 3.</p>

Important:

When you create the SR-IOV VFs, ensure that you do not assign MAC addresses to the VFs.

Figure 1: Enable SR-IOV VFs on the KVM host for Intel 82599 10G NIC.

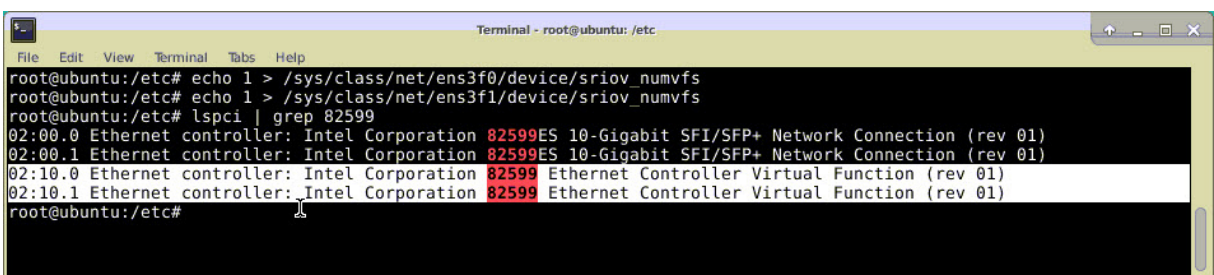


Figure 2: Enable SR-IOV VFs on the KVM host for Intel X710 10G and XL710 40G NICs.

```

root@ubuntu:~# lspci | grep 710
03:00.0 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01)
03:00.1 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01)
03:00.2 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01)
03:00.3 Ethernet controller: Intel Corporation Ethernet Controller X710 for 10GbE SFP+ (rev 01)
03:06.0 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:06.1 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0a.0 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0a.1 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0a.2 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0a.3 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0e.0 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0e.1 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0e.2 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
03:0e.3 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 01)
81:00.0 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 01)
82:00.0 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 02)
82:00.1 Ethernet controller: Intel Corporation Ethernet Controller XL710 for 40GbE QSFP+ (rev 02)
82:02.0 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 02)
82:02.1 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 02)
82:0a.0 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 02)
82:0a.1 Ethernet controller: Intel Corporation XL710/X710 Virtual Function (rev 02)
root@ubuntu:~#

```

Figure 3: Enable SR-IOV VFs on the KVM host for Intel X722 10G NIC.

```

root@ubuntu:~# lspci | grep "37cd"
84:02.0 Ethernet controller: Intel Corporation Device 37cd (rev 04)
84:0a.0 Ethernet controller: Intel Corporation Device 37cd (rev 04)

```

Figure 4: Make the VFs persistent.

```


Terminal - root@ubuntu: /etc
File Edit View Terminal Tabs Help
root@ubuntu:/etc# cat /etc/rc.local
#!/bin/sh -e
#
# rc.local
#
# This script is executed at the end of each multiuser runlevel.
# Make sure that the script will "exit 0" on success or any other
# value on error.
#
# In order to enable or disable this script just change the execution
# bits.
#
# By default this script does nothing.
echo 1 > /sys/class/net/ens3f0/device/sriov_numvfs
echo 1 > /sys/class/net/ens3f1/device/sriov_numvfs
exit 0
root@ubuntu:/etc#

```


Configure a Citrix ADC VPX instance to use SR-IOV network interface

To configure the Citrix ADC VPX instance to use SR-IOV network interface by using Virtual Machine Manager, complete these steps:


1. Power off the Citrix ADC VPX instance.
2. Select the Citrix ADC VPX instance and then select Open.



3. In the <virtual machine on KVM> window, select the **i** icon.



4. Select **Add Hardware**.



5. In the **Add New Virtual Hardware** dialog box, do the following:
 - a) Select PCI Host Device.
 - b) In the Host Device section, select the VF you have created and click Finish.

Figure 4: VF for Intel 82599 10G NIC

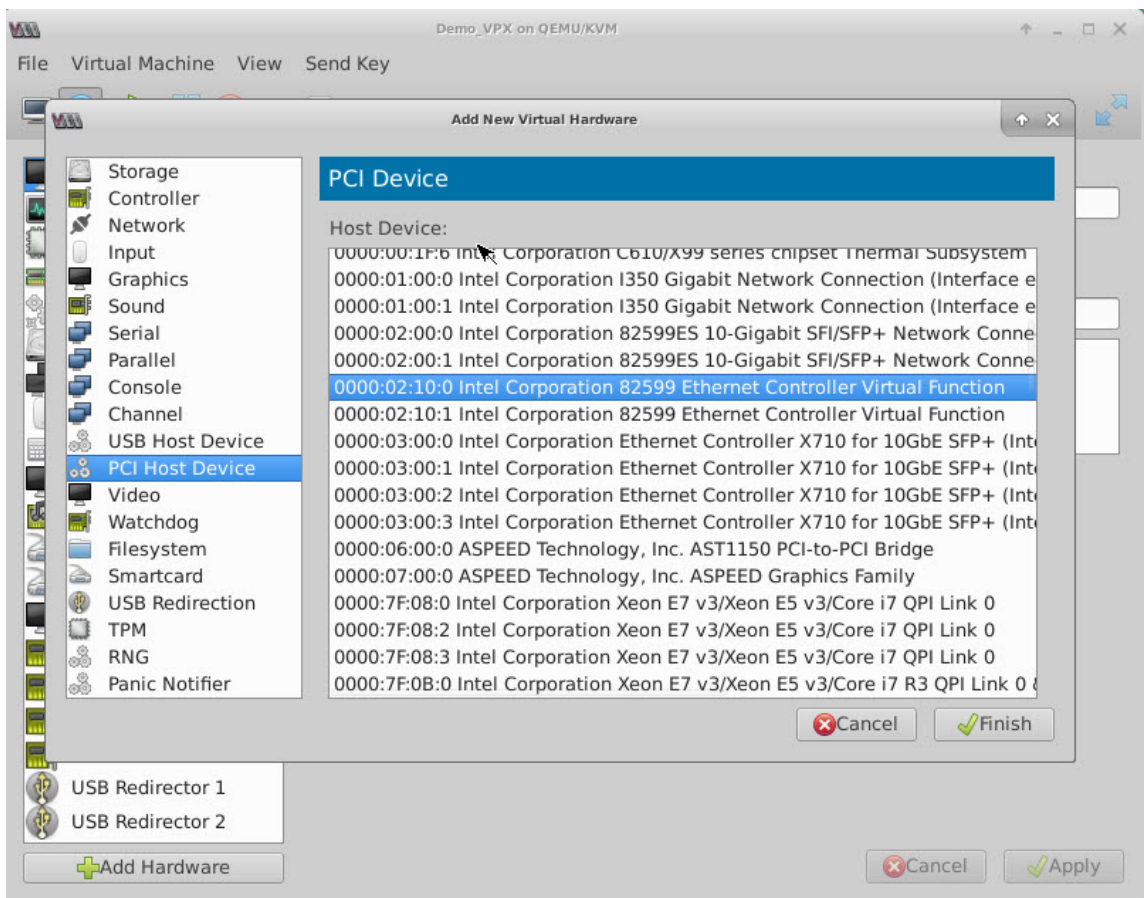


Figure 5: VF for Intel XL710 40G NIC

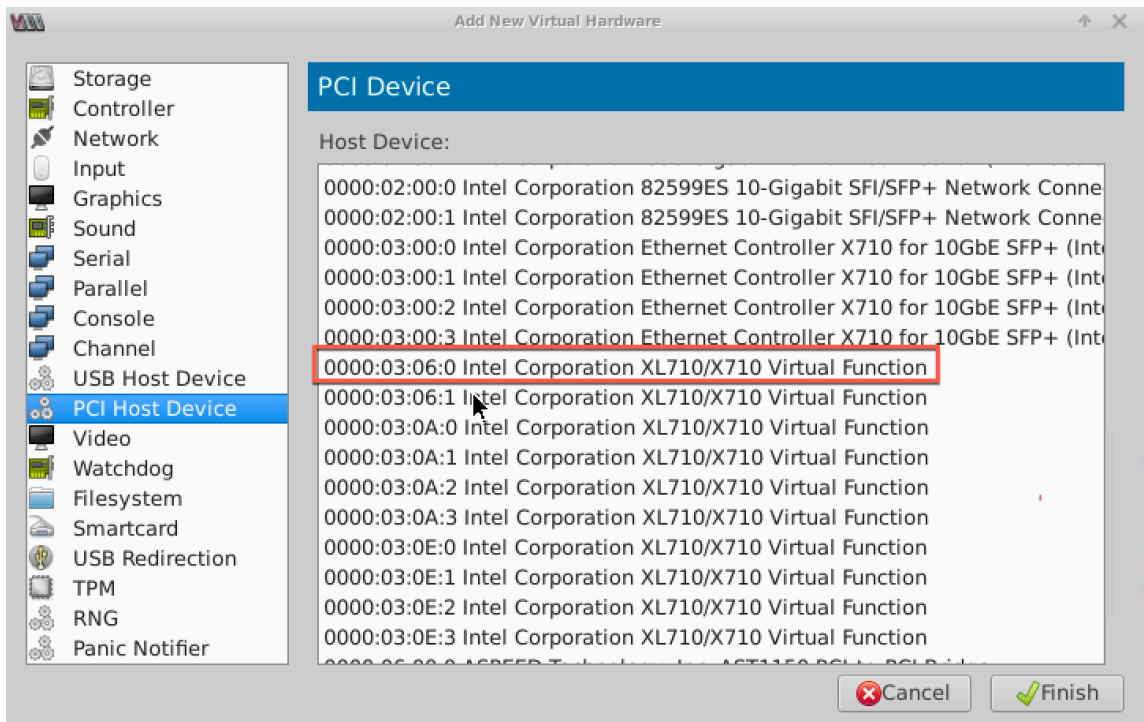



Figure 6: VF for Intel X722 10G NIC



6. Repeat Step 4 and 5 to add the VFs that you have created.
7. Power on the Citrix ADC VPX instance.
8. After the Citrix ADC VPX instance powers on, use the following command to verify the configuration:

```
1 show interface summary
```

The output shows all the interfaces that you configured.

Figure 6: output summary for Intel 82599 NIC.





Figure 7. Output summary for Intel X710 and XL710 NICs.



Configure static LA/LACP on the SR-IOV interface

Important:

When you are creating the SR-IOV VFs, ensure that you do not assign MAC addresses to the VFs.

To use the SR-IOV VFs in link aggregation mode, disable spoof checking for VFs that you have created. On the KVM host, use the following command to disable spoof checking:

```
*ip link set \<interface\_name\> vf \<VF\_id\> spoofchk off*
```

Where:

- Interface_name –is the interface name.
- VF_id –is the Virtual Function id.

Example:

```
Terminal - root@ubuntu: /etc
File Edit View Terminal Tabs Help
root@ubuntu:/etc# ip link show ens3f0
6: ens3f0: <BROADCAST,MULTICAST,UP,LOWER UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether 0c:c4:7a:bd:50:7e brd ff:ff:ff:ff:ff:ff
    vf 0 MAC 8e:e7:e7:06:50:3f, spoof checking on, link-state auto
root@ubuntu:/etc#
root@ubuntu:/etc#
root@ubuntu:/etc# ip link show ens3f1
7: ens3f1: <BROADCAST,MULTICAST,UP,LOWER UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether 0c:c4:7a:bd:50:7f brd ff:ff:ff:ff:ff:ff
    vf 0 MAC 8e:1a:71:cc:a8:3e, spoof checking on, link-state auto
root@ubuntu:/etc#
root@ubuntu:/etc# ip link set ens3f0 vf 0 spoofchk off
root@ubuntu:/etc# ip link show ens3f0
6: ens3f0: <BROADCAST,MULTICAST,UP,LOWER UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether 0c:c4:7a:bd:50:7e brd ff:ff:ff:ff:ff:ff
    vf 0 MAC 8e:e7:e7:06:50:3f, spoof checking off, link-state auto
root@ubuntu:/etc# ip link set ens3f1 vf 0 spoofchk off
root@ubuntu:/etc# ip link show ens3f1
7: ens3f1: <BROADCAST,MULTICAST,UP,LOWER UP> mtu 1500 qdisc mq state UP mode DEFAULT group default qlen 1000
    link/ether 0c:c4:7a:bd:50:7f brd ff:ff:ff:ff:ff:ff
    vf 0 MAC 8e:1a:71:cc:a8:3e, spoof checking off, link-state auto
root@ubuntu:/etc#
```

After you disable spoof checking for all the VFs that you have created. Restart the Citrix ADC VPX instance and configure link aggregation. For detailed instructions, see [Configuring Link Aggregation](#).

Configuring VLAN on the SR-IOV Interface

You can configure VLAN on SR-IOV VFs. For detailed instructions, see [Configuring a VLAN](#).

Important:

Ensure that the KVM host does not contain VLAN settings for the VF interface.

Configure a Citrix ADC VPX instance to use PCI passthrough network interfaces

September 6, 2024

After you have installed and configured a Citrix ADC VPX instance on the Linux-KVM platform, you can use the Virtual Machine Manager to configure the virtual appliance to use PCI passthrough network interfaces.

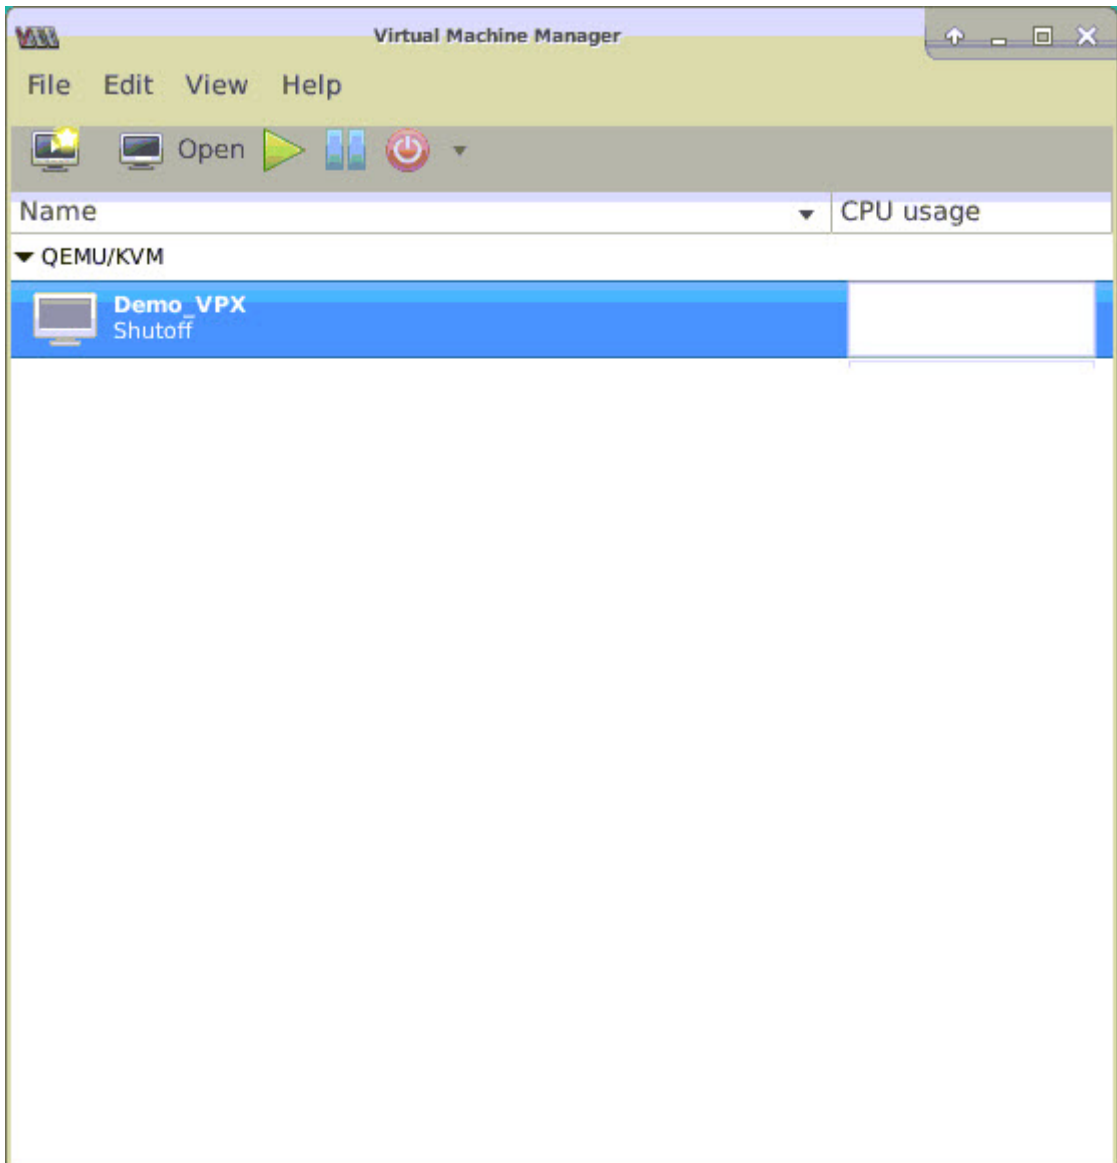
Prerequisites

- The firmware version of the Intel XL710 NIC (NIC) on the KVM Host is 5.04.
- The KVM Host supports input–output memory management unit (IOMMU) and Intel VT-d, and they are enabled in the BIOS of the KVM Host. On the KVM Host, to enable IOMMU, add the following entry to the `/boot/grub2/grub.cfg` file: **intel_iommu=1**


- Run the following command and reboot the KVM Host: **Grub2-mkconfig -o /boot/grub2/grub.cfg**

To configure Citrix ADC VPX instances to use PCI passthrough network interfaces by using the Virtual Machine Manager:

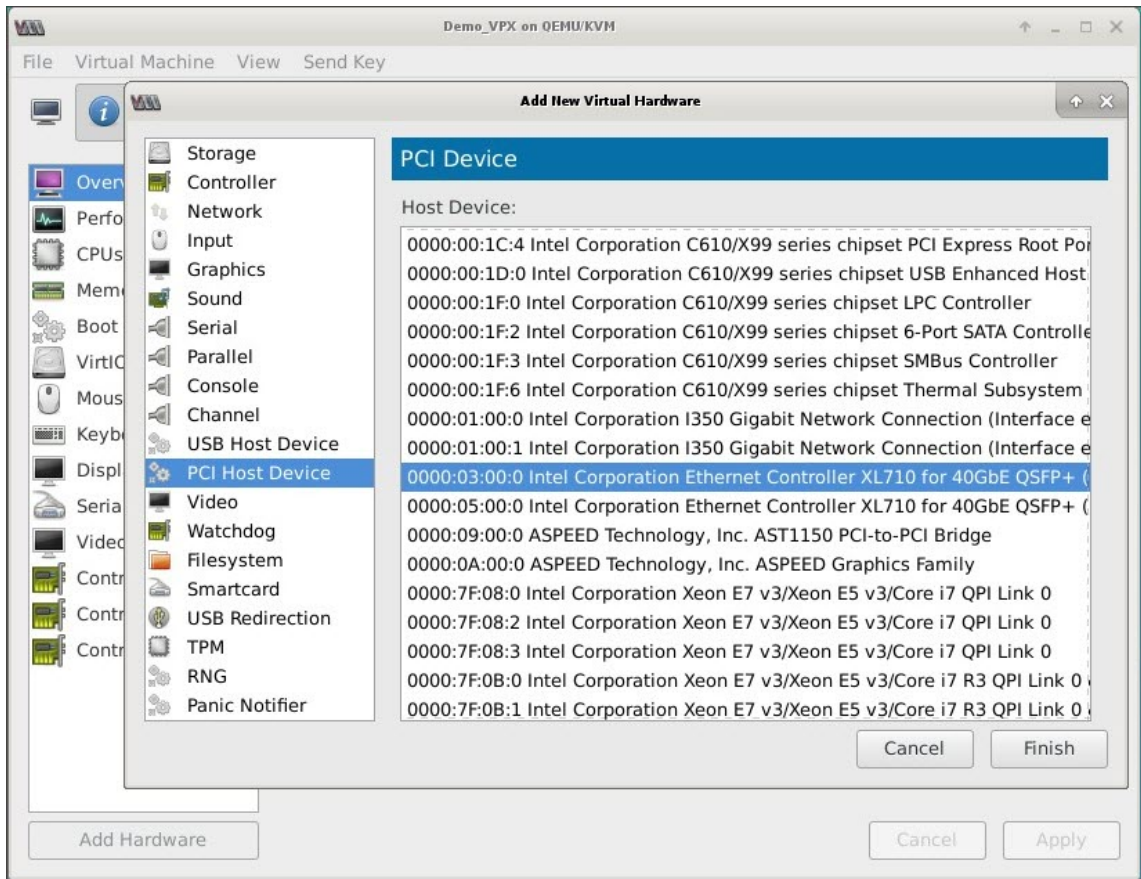
1. Power off the Citrix ADC VPX instance.
2. Select the Citrix ADC VPX instance and click **Open**.



3. In the **virtual_machine on KVM>** window, click the **i** icon.



4. Click **Add Hardware**.
5. In the **Add New Virtual Hardware** dialog box, do the following:
 - a. Select **PCI Host Device**.
 - b. In the **Host Device** section, select the Intel XL710 physical function.
 - c. Click **Finish**.

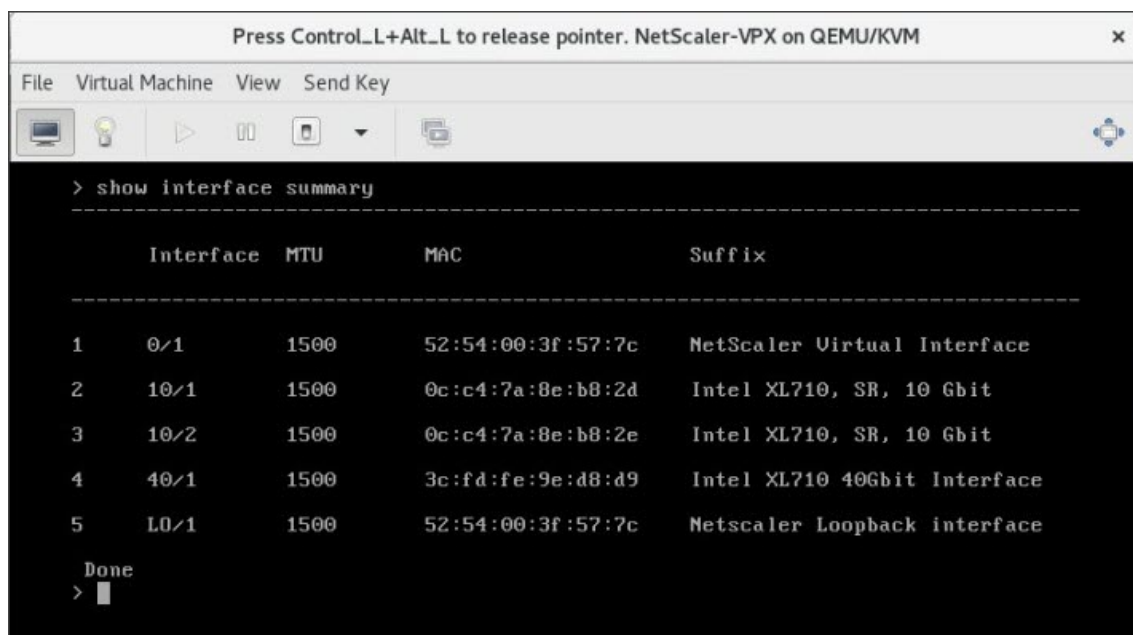


6. Repeat steps 4 and 5 to add any additional Intel XL710 physical functions.
7. Power on the Citrix ADC VPX instance.
8. Once the Citrix ADC VPX instance powers on, you can use the following command to verify the configuration:

```

COMMAND
> show interface summary
    
```

The output must show all the interfaces that you configured:



```

> show interface summary
-----
      Interface  MTU      MAC      Suffix
-----
1      0/1        1500     52:54:00:3f:57:7c  NetScaler Virtual Interface
2      10/1        1500     0c:c4:7a:8e:b8:2d  Intel XL710, SR, 10 Gbit
3      10/2        1500     0c:c4:7a:8e:b8:2e  Intel XL710, SR, 10 Gbit
4      40/1        1500     3c:fd:fe:9e:d8:d9  Intel XL710 40Gbit Interface
5      L0/1        1500     52:54:00:3f:57:7c  Netscaler Loopback interface

Done
> █

```

Provision the Citrix ADC VPX instance by using the virsh program

September 18, 2024

The `virsh` program is a command line tool for managing VM Guests. Its functionality is similar to that of Virtual Machine Manager. It enables you to change a VM Guest's status (start, stop, pause, and so on), to set up new Guests and devices, and to edit existing configurations. The `virsh` program is also useful for scripting VM Guest management operations.

To provision Citrix ADC VPX by using the `virsh` program, follow these steps:

1. Use the `tar` command to untar the Citrix ADC VPX package. The `NSVPX-KVM-*_nc.tgz` package contains the following components:
 - The Domain XML file specifying VPX attributes [`NSVPX-KVM-*_nc.xml`]
 - Check sum of NS-VM Disk Image [`Checksum.txt`]
 - NS-VM Disk Image [`NSVPX-KVM-*_nc.raw`]

Example:

```

1 tar -xvzf NSVPX-KVM-10.1-117_nc.tgz
2 NSVPX-KVM-10.1-117_nc.xml
3 NSVPX-KVM-10.1-117_nc.raw
4 checksum.txt

```

2. Copy the `NSVPX-KVM-\<*\>_nc.xml` XML file to a file named `\<DomainName\>-NSVPX-KVM-\<*\>_nc.xml`. The `<DomainName>` is also the name of the virtual machine. Example:

```
1 cp NSVPX-KVM-10.1-117_nc.xml NetScaler-VPX-NSVPX-KVM-10.1-117_nc.xml
```

3. Edit the `\<DomainName\>-NSVPX-KVM-*_nc.xml` file to specify the following parameters:

- name—Specify the name.
- Mac—Specify the MAC address.

Note:

The domain name and the MAC address have to be unique.

- source file—Specify the absolute disk-image source path. The file path has to be absolute. You can specify the path of the RAW image file or a QCOW2 image file.

If you want to specify a RAW image file, specify the disk image source path as shown in the following example:

Example:

```
1 <name>NetScaler-VPX</name>
2 <mac address='52:54:00:29:74:b3' />
3 <source file='/root/NSVPX-KVM-10.1-117_nc.raw' />
```

Specify the absolute QCOW2 disk-image source path and define the driver type as **qcow2**, as shown in the following example:

Example:

```
1 <name>NetScaler-VPX</name>
2 <mac address='52:54:00:29:74:b3' />
3 <driver name='qemu' type='qcow2' />
4 <source file='/root/NSVPX-KVM-10.1-117_nc.qcow' />*
```

4. Edit the `\<DomainName\>-NSVPX-KVM-*_nc.xml` file to configure the networking details:

- source dev—specify the interface.
- mode—specify the mode. The default interface is **Macvtap Bridge**.

Example: Mode: MacVTap Bridge Set target interface as `ethx` and mode as bridge Model type as `virtio`

```
1 <interface type='direct'>
2   <mac address='52:54:00:29:74:b3' />
3   <source dev='eth0' mode='bridge' />
4   <target dev='macvtap0' />
5   <model type='virtio' />
```

```

6     <alias name='net0' />
7     <address type='pci' domain='0x0000' bus='0x00' slot='0x03'
      function='0x0' />
8     </interface>

```

Here, eth0 is the physical interface attached to the VM.

5. Define the VM attributes in the `\<DomainName\>-NSVPX-KVM-*_nc.xml` file by using the following command:

```
1 virsh define \<DomainName\>-NSVPX-KVM-\*\_nc.xml
```

Example:

```
1 virsh define NS-VPX-NSVPX-KVM-10.1-117_nc.xml
```

6. Start the VM by entering the following command:

```
1 virsh start \[\<DomainName\> | \<DomainUUID\>\]
```

Example:

```
1 virsh start NetScaler-VPX
```

7. Connect the Guest VM through the console

```
1 virsh console \[\<DomainName\> | \<DomainUUID\> | \<DomainID\> \]
```

Example:

```
1 virsh console NetScaler-VPX
```

Add more interfaces to Citrix ADC VPX instance using `virsh` program

After you have provisioned the Citrix ADC VPX on KVM, you can add additional interfaces.

To add more interfaces, follow these steps:

1. Shut down the Citrix ADC VPX instance running on the KVM.
2. Edit the `\<DomainName\>-NSVPX-KVM-*_nc.xml` file using the command:

```
1 virsh edit \[\<DomainName\> | \<DomainUUID\>\]
```

3. In the `\<DomainName\>-NSVPX-KVM-*_nc.xml` file, append the following parameters:

a) For MacVTap

- Interface type—Specify the interface type as ‘direct’.

- MAC address—Specify the MAC address and make sure the MAC address is unique across the interfaces.
- source dev—Specify the interface name.
- mode—Specify the mode. The modes supported are - Bridge, VEPA, Private, and Pass-through
- model type—Specify the model type as `virtio`

Example:

Mode: MacVTap Pass-through

Set target interface as

`ethx`, Mode as

bridge, and model type as

`virtio`

```

1 <interface type='direct'>
2     <mac address='52:54:00:29:74:b3' />
3     <source dev='eth1' mode='passthrough' />
4     <model type='virtio' />
5 </interface>

```

Here `eth1` is the physical interface attached to the VM.

b) For Bridge Mode

Note:

Make sure that you have configured a Linux bridge in the KVM host, bound the physical interface to the bridge, and put the bridge in the UP state.

- Interface type—Specify the interface type as ‘bridge’.
- MAC address—Specify the MAC address and make sure the MAC address is unique across the interfaces.
- source bridge—Specify the bridge name.
- model type—Specify the model type as `virtio`

Example: Bridge Mode

```

1 <interface type='bridge'>
2     <mac address='52:54:00:2d:43:a4' />
3     <source bridge='br0' />
4     <model type='virtio' />
5 </interface>

```

Manage the Citrix ADC VPX guest VMs

September 12, 2024

You can use the Virtual Machine Manager and the `virsh` program to perform management tasks such as starting or stopping a VM Guest, setting up new guests and devices, editing existing configurations, and connecting to the graphical console through Virtual Network Computing (VNC).

Manage the VPX guest VMs by using Virtual Machine Manager

- List the VM guests


The main Window of the Virtual Machine Manager displays a list of all the VM Guests for each VM host server it is connected to. Each VM Guest entry contains the virtual machine's name, along with its status (Running, Paused, or Shutoff) displayed as in the icon.

- Open a graphical console

Opening a Graphical Console to a VM Guest enables you to interact with the machine like you would with a physical host through a VNC connection. To open the graphical console in the Virtual Machine Manager, right-click the VM Guest entry and select the Open option from the pop-up menu.

- Start and shut down a guest

You can start or stop a VM Guest from the Virtual Machine Manager. To change the state of the VM, right-click the VM Guest entry and select Run or one of the Shut Down options from the pop-up menu.



- Reboot a guest

You can reboot a VM Guest from the Virtual Machine Manager. To reboot the VM, right-click the VM Guest entry, and then select Shut Down > Reboot from the pop-up menu.

- Delete a guest

Deleting a VM Guest removes its XML configuration by default. You can also delete a guest's storage files. Doing so completely erases the guest.

1. In the Virtual Machine Manager, right-click the VM Guest entry.
2. Select Delete from the pop-up menu. A confirmation window opens.

Note:

The Delete option is enabled only when the VM Guest is shut down.

3. Click **Delete**.
4. To completely erase the guest, delete the associated .raw file by selecting the Delete Associated Storage Files check box.

Manage the Citrix ADC VPX guest VMs using the `virsh` program

- List the VM Guests and their current states.

To use `virsh` to display information about the Guests

```
virsh list --all
```

The command output displays all domains with their states. Example output:

1	Id	Name	State
2	-----		
3	0	Domain-0	running
4	1	Domain-1	paused
5	2	Domain-2	inactive
6	3	Domain-3	crashed

- Open a `virsh` console.

Connect the Guest VM through the console

```
virsh console [<DomainID> | <DomainName> | <DomainUUID>]
```

Example:

```
virsh console NetScaler-VPX
```

- Start and shut down a guest.

Guests can be started using the DomainName or Domain-UUID.

```
virsh start [<DomainName> | <DomainUUID>]
```

Example:

```
virsh start NetScaler-VPX
```

To shut down a guest:

```
virsh shutdown [<DomainID> | <DomainName> | <DomainUUID>]
```

Example:

```
virsh shutdown NetScaler-VPX
```

- Reboot a guest

```
virsh reboot [<DomainID> | <DomainName> | <DomainUUID>]
```

Example:

```
virsh reboot NetScaler-VPX
```

Delete a guest

To delete a Guest VM you must shut down the Guest and undefine the `<DomainName>-NSVPX-KVM-*_nc.xml` before you run the delete command.

```
1 virsh shutdown [<DomainID> | <DomainName> | <DomainUUID>]
2 virsh undefine [<DomainName> | <DomainUUID>]
```

Example:

```
1 virsh shutdown NetScaler-VPX
2 virsh undefine NetScaler-VPX
```

Note:

The delete command doesn't remove disk image file which must be removed manually.

Provision the Citrix ADC VPX instance with SR-IOV, on OpenStack

September 6, 2024

You can deploy high-performance Citrix ADC VPX instances that use single-root I/O virtualization (SR-IOV) technology, on OpenStack.

You can deploy a Citrix ADC VPX instance that uses SR-IOV technology, on OpenStack, in three steps:

- Enable SR-IOV Virtual Functions (VFs) on the host.
- Configure and make the VFs available to OpenStack.
- Provision the Citrix ADC VPX on OpenStack.

Prerequisites

Ensure that you:

- Add the Intel 82599 NIC (NIC) to the host.
- Download and Install the latest IXGBE driver from Intel.
- Block list the IXGBEVF driver on the host. Add the following entry in the `/etc/modprobe.d/blacklist.conf` file: Block list `ixgbevf`

Note:

The `ixgbe` driver version must be minimum 5.0.4.

Enable SR-IOV VFs on the host

Do one of the following steps to enable SR-IOV VFs:

- If you are using a kernel version earlier than 3.8, add the following entry to the `/etc/modprobe.d/ixgbe` file and restart the host: `options ixgbe max_vfs=<number_of_VFs>`
- If you are using kernel 3.8 version or later, create VFs by using the following command:

```
1 echo <number_of_VFs> > /sys/class/net/<device_name>/device/sriov_numvfs
```

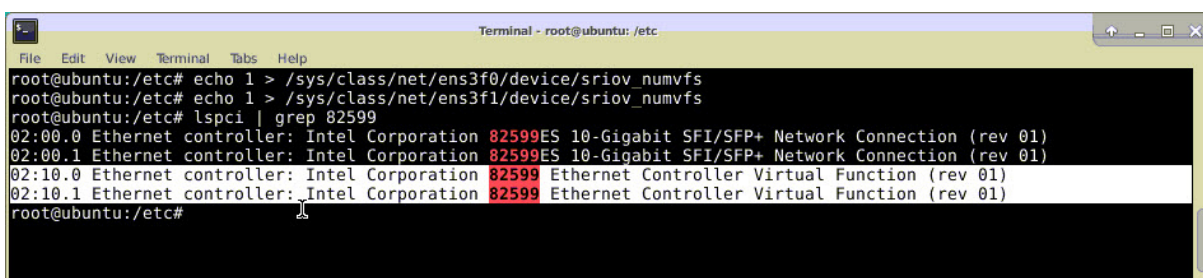
Where:

- `number_of_VFs` is the number of Virtual Functions that you want to create.
- `device_name` is the interface name.

Important:

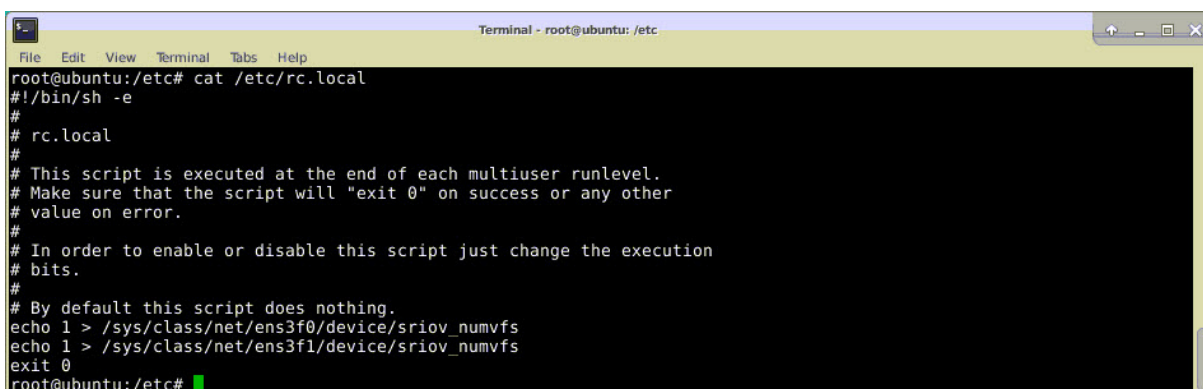
While you are creating the SR-IOV VFs, make sure that you do not assign MAC addresses to the VFs.

Here is an example of four VFs being created.



```
Terminal - root@ubuntu: /etc
root@ubuntu:/etc# echo 1 > /sys/class/net/ens3f0/device/sriov_numvfs
root@ubuntu:/etc# echo 1 > /sys/class/net/ens3f1/device/sriov_numvfs
root@ubuntu:/etc# lspci | grep 82599
02:00.0 Ethernet controller: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
02:00.1 Ethernet controller: Intel Corporation 82599ES 10-Gigabit SFI/SFP+ Network Connection (rev 01)
02:10.0 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
02:10.1 Ethernet controller: Intel Corporation 82599 Ethernet Controller Virtual Function (rev 01)
root@ubuntu:/etc#
```

Make the VFs persistent, add the commands that you used to create VFs to the `rc.local` file. Here is an example showing content of `rc.local` file.



```
Terminal - root@ubuntu: /etc
root@ubuntu:/etc# cat /etc/rc.local
#!/bin/sh -e
#
# rc.local
#
# This script is executed at the end of each multiuser runlevel.
# Make sure that the script will "exit 0" on success or any other
# value on error.
#
# In order to enable or disable this script just change the execution
# bits.
#
# By default this script does nothing.
echo 1 > /sys/class/net/ens3f0/device/sriov_numvfs
echo 1 > /sys/class/net/ens3f1/device/sriov_numvfs
exit 0
root@ubuntu:/etc#
```

For more information, see this [Intel SR-IOV Configuration Guide](#).

Configure and make the VFs available to OpenStack

Follow the steps given at the link below to configure SR-IOV on OpenStack: <https://wiki.openstack.org/wiki/SR-IOV-Passthrough-For-Networking>.

Provision the Citrix ADC VPX instance on OpenStack

You can provision a Citrix ADC VPX instance in an OpenStack environment by using the OpenStack CLI.

Provisioning a VPX instance, optionally involves using data from the config drive. The config drive is a special configuration drive that attaches to the instance when it boots. This configuration drive can be used to pass networking configuration information such as management IP address, network mask, and default gateway and so on to the instance before you configure the network settings for the instance.

When OpenStack provisions a VPX instance, it first detects that the instance is booting in an OpenStack environment, by reading a specific BIOS string (OpenStack Foundation) that indicates OpenStack. For Red Hat Linux distributions, the string is stored in `/etc/nova/release`. This is a standard mechanism that is available in all OpenStack implementations based on KVM hyper-visor platform. The drive must have a specific OpenStack label. If the config drive is detected, the instance attempts to read the following information from the file name specified in the `nova` boot command. In the procedures below, the file is called “`userdata.txt`.”

- Management IP address
- Network mask
- Default gateway

Once the parameters are successfully read, they are populated in the NetScaler stack. This helps in managing the instance remotely. If the parameters are not read successfully or the config drive is not available, the instance transitions to the default behavior, which is:

- The instance attempts to retrieve the IP address information from DHCP.
- If DHCP fails or times-out, the instance comes up with default network configuration (192.168.100.1/16).

Provision the Citrix ADC VPX instance on OpenStack through CLI

You can provision a VPX instance in an OpenStack environment by using the OpenStack CLI. Here’s the summary of the steps to provision a Citrix ADC VPX instance on OpenStack:

1. Extracting the `.qcow2` file from the `.tgz` file
2. Building an OpenStack image from the `qcow2` image
3. Provisioning a VPX instance

To provision a VPX instance in an OpenStack environment, do the following steps.

1. Extract the `qcow2` file from the `.tgz` file by typing the command:

```

1 tar xvzf <TAR file>
2 tar xvzf NSVPX-KVM-12.0-26.2_nc.tgz
3 NSVPX-KVM.xml
4 NSVPX-KVM-12.0-26.2_nc.qcow2

```

2. Build an OpenStack image using the `.qcow2` file extracted in step 1 by typing the following command:

```

1 glance image-create --name="<name of the OpenStack image>" --
  property hw_disk_bus=ide --is-public=true --container-format=
  bare --disk-format=qcow2< <name of the qcow2 file>
2
3 glance image-create --name="NS-VPX-12-0-26-2" --property
  hw_disk_bus=ide --is-public= true --container-format=bare --
  disk-format=qcow2< NSVPX-KVM-12.0-26.2_nc.qcow2

```

The following illustration provides a sample output for the `glance image-create` command.

Property	Value
checksum	735dae4ea6e46e39ed3f0acfba02e755
container_format	bare
created_at	2017-02-16T10:03:29Z
disk_format	qcow2
hw_disk_bus	ide
id	aeaa13e9-b49b-411c-ab54-c61820a8e2f3
min_disk	0
min_ram	0
name	NSVPX-KVM-12.0-26.2
owner	06c41a73b32f4b48af55359fd7d3502c
protected	False
size	717946880
status	active
tags	[]
updated_at	2017-02-16T10:03:38Z
virtual_size	None
visibility	private

3. After an OpenStack image is created, provision the Citrix ADC VPX instance.

```

1 nova boot --image NSVPX-KVM-12.0-26.2 --config-drive=true --
  userdata
2 ./userdata.txt --flavor m1.medium --nic net-id=3b258725-eaae-
3 455e-a5de-371d6d1f349f --nic port-id=218ba819-9f55-4991-adb6-
4 02086a6bdee2 NSVPX-10

```

In the preceding command, `userdata.txt` is the file which contains the details like, IP address, netmask, and default gateway for the VPX instance. The user data file is a user customizable file. `NSVPX-KVM-12.0-26.2` is the name of the virtual appliance that you want to provision. `-NIC port-id=218ba819-9f55-4991-adb6-02086a6bdee2` is the OpenStack VF.

The following illustration gives a sample output of the `nova boot` command.

Property	Value
OS-DCF:diskConfig	MANUAL
OS-EXT-AZ:availability_zone	
OS-EXT-SRV-ATTR:host	-
OS-EXT-SRV-ATTR:hypervisor_hostname	-
OS-EXT-SRV-ATTR:instance_name	instance-0000003c
OS-EXT-STS:power_state	0
OS-EXT-STS:task_state	scheduling
OS-EXT-STS:vm_state	building
OS-SRV-USG:launched_at	-
OS-SRV-USG:terminated_at	-
accessIPv4	
accessIPv6	
adminPass	43EjPdM5shLz
config_drive	True
created	2017-02-20T11:53:37Z
flavor	m1.medium (3)
hostId	
id	6b9f6968-aab9-463c-b619-d58c73db3187
image	NSVPX-KVM-12.0-26.2 (a5478b8a-8435-48d1-b4a0-1494e2c8f8b1)
key_name	-
metadata	{}
name	NSVPX-10
os-extended-volumes:volumes_attached	[]
progress	0
security_groups	default
status	BUILD
tenant_id	06c41a73b32f4b48af55359fd7d3502c
updated	2017-02-20T11:53:38Z
user_id	418524f7101b4f0389ecbb36da9916b5

The following illustration shows a sample of the `userdata.txt` file. The values within the `<PropertySection></PropertySection>` tags are the values which are user configurable and holds the information like, IP address, netmask, and default gateway.

```

1 <?xml version="1.0" encoding="UTF-8" standalone="no"?>
2 <Environment xmlns:oe="http://schemas.dmtf.org/ovf/environment/1"
3 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
4 oe:id=""
5 xmlns="http://schemas.dmtf.org/ovf/environment/1">
6 <PlatformSection>
7 <Kind>NOVA</Kind>
8 <Version>2013.1</Version>
9 <Vendor>Openstack</Vendor>
10 <Locale>en</Locale>
11 </PlatformSection>
12 <PropertySection>
13 <Property oe:key="com.citrix.netscaler.ovf.version" oe:value="1.0"
    />

```

```

14 <Property oe:key="com.citrix.netscaler.platform" oe:value="vpx"/>
15 citrix.com 4
16 <Property oe:key="com.citrix.netscaler.orch_env"
17 oe:value="openstack-orch-env"/>
18 <Property oe:key="com.citrix.netscaler.mgmt.ip"
19 oe:value="10.1.0.100"/>
20 <Property oe:key="com.citrix.netscaler.mgmt.netmask"
21 oe:value="255.255.0.0"/>
22 <Property oe:key="com.citrix.netscaler.mgmt.gateway"
23 oe:value="10.1.0.1"/>
24 </PropertySection>
25 </Environment>

```

Additional supported Configurations: Creating and Deleting VLANs on SR-IOV VFs from the Host

Type the following command to create a VLAN on the SR-IOV VF:

```
ip link show enp8s0f0 vf 6 vlan 10
```

In the preceding command “enp8s0f0” is the name of the physical function.

Example: VLAN 10, created on vf 6

```

4: enp8s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT qlen 1000
   link/ether 00:1b:21:7b:d7:88 brd ff:ff:ff:ff:ff:ff
   vf 0 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off
   vf 1 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off
   vf 2 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off
   vf 3 MAC fa:16:3e:1e:0b:ee, spoof checking on, link-state auto, trust off
   vf 4 MAC fa:16:3e:0d:05:62, spoof checking on, link-state auto, trust off
   vf 5 MAC 5e:46:0d:79:de:f8, spoof checking on, link-state auto, trust off
   vf 6 MAC fa:16:3e:db:ea:b3, vlan 10, spoof checking on, link-state auto, trust off
   vf 7 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off

```

Type the following command to delete a VLAN on the SR-IOV VF:

```
ip link show enp8s0f0 vf 6 vlan 0
```

Example: VLAN 10, removed from vf 6

```

[root@localhost ~]# ip link show enp8s0f0
4: enp8s0f0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP mode DEFAULT qlen 1000
   link/ether 00:1b:21:7b:d7:88 brd ff:ff:ff:ff:ff:ff
   vf 0 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off
   vf 1 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off
   vf 2 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off
   vf 3 MAC fa:16:3e:1e:0b:ee, spoof checking on, link-state auto, trust off
   vf 4 MAC fa:16:3e:0d:05:62, spoof checking on, link-state auto, trust off
   vf 5 MAC 5e:46:0d:79:de:f8, spoof checking on, link-state auto, trust off
   vf 6 MAC fa:16:3e:db:ea:b3, spoof checking on, link-state auto, trust off
   vf 7 MAC 00:00:00:00:00:00, spoof checking on, link-state auto, trust off

```

These steps complete the procedure for deploying a Citrix ADC VPX instance that uses SRIOV technology, on OpenStack.

Configure a Citrix ADC VPX instance on KVM to use OVS DPDK-based host interfaces

June 20, 2024

You can configure a Citrix ADC VPX instance running on KVM (Fedora and RHOS) to use Open vSwitch (OVS) with Data Plane Development Kit (DPDK) for better network performance. This document describes how to configure the Citrix ADC VPX instance to operate on the `vhost-user` ports exposed by OVS-DPDK on the KVM host.

[OVS](#) is a multilayer virtual switch licensed under the open-source Apache 2.0 license. [DPDK](#) is a set of libraries and drivers for fast packet processing.

The following Fedora, RHOS, OVS, and DPDK versions are qualified for configuring a Citrix ADC VPX instance:

Fedora	RHOS
Fedora 25	RHOS 7.4
OVS 2.7.0	OVS 2.6.1
DPDK 16.11.12	DPDK 16.11.12

Prerequisites

Before you install DPDK, make sure the host has 1 GB huge pages.

For more information, see this [DPDK system requirements documentation](#). Here is a summary of the steps required to configure a Citrix ADC VPX instance on KVM to use OVS DPDK-based host interfaces:

- Install DPDK.
- Build and Install OVS.
- Create an OVS bridge.
- Attach a physical interface to the OVS bridge.
- Attach `vhost-user` ports to the OVS data path.
- Provision a KVM-VPX with OVS-DPDK based `vhost-user` ports.

Install DPDK

To install DPDK, follow the instruction given at this [Open vSwitch with DPDK](#) document.

Build and install OVS

Download OVS from the OVS [download page](#). Next, build, and install OVS by using a DPDK datapath. Follow the instructions given in the [Installing Open vSwitch](#) document.

For more detailed information, [DPDK Getting Started Guide for Linux](#).

Create an OVS bridge

Depending on your need, type the Fedora or RHOS command to create an OVS bridge:

Fedora command:

```
1 > $OVS_DIR/utilities/ovs-vsctl add-br ovs-br0 -- set bridge ovs-br0
   datapath_type=netdev
```

RHOS command:

```
1 ovs-vsctl add-br ovs-br0 -- set bridge ovs-br0 datapath_type=netdev
```

Attach the physical interface to the OVS bridge

Bind the ports to DPDK and then attach them to the OVS bridge by typing the following Fedora or RHOS commands:

Fedora command:

```
1 > $OVS_DIR/utilities/ovs-vsctl add-port ovs-br0 dpdk0 -- set Interface
   dpdk0 type=dpgk options:dpdk-devargs=0000:03:00.0
2
3 > $OVS_DIR/utilities/ovs-vsctl add-port ovs-br0 dpdk1 -- set Interface
   dpdk1 type=dpgk options:dpdk-devargs=0000:03:00.1
```

RHOS command:

```
1 ovs-vsctl add-port ovs-br0 dpdk0 -- set Interface dpdk0 type=dpgk
   options:dpdk-devargs=0000:03:00.0
2
3
4 ovs-vsctl add-port ovs-br0 dpdk1 -- set Interface dpdk1 type=dpgk
   options:dpdk-devargs=0000:03:00.1
```

The `dpdk-devargs` shown as part of the options specifies the PCI BDF of the respective physical NIC.

Attach vhost-user ports to the OVS data path

Type the following Fedora or RHOS commands to attach `vhost-user` ports to the OVS data path:

Fedora command:

```
1 > $OVS_DIR/utilities/ovs-vsctl add-port ovs-br0 vhost-user1 -- set
   Interface vhost-user1 type=dpdkvhostuser -- set Interface vhost-
   user1 mtu_request=9000
2
3 > $OVS_DIR/utilities/ovs-vsctl add-port ovs-br0 vhost-user2 -- set
   Interface vhost-user2 type=dpdkvhostuser -- set Interface vhost-
   user2 mtu_request=9000
4
5 chmod g+w /usr/local/var/run/openvswitch/vhost*
```

RHOS command:

```
1 ovs-vsctl add-port ovs-br0 vhost-user1 -- set Interface vhost-user1
   type=dpdkvhostuser -- set Interface vhost-user1 mtu_request=9000
2
3 ovs-vsctl add-port ovs-br0 vhost-user2 -- set Interface vhost-user2
   type=dpdkvhostuser -- set Interface vhost-user2 mtu_request=9000
4
5 chmod g+w /var/run/openvswitch/vhost*
```

Provision a KVM-VPX with OVS-DPDK-based vhost-user ports

You can provision a VPX instance on Fedora KVM with OVS-DPDK-based `vhost-user` ports only from the CLI by using the following QEMU commands:

Fedora command:

```
1 qemu-system-x86_64 -name KVM-VPX -cpu host -enable-kvm -m 4096M \
2
3 -object memory-backend-file,id=mem,size=4096M,mem-path=/dev/hugepages,
   share=on -numa node,memdev=mem \
4
5 -mem-prealloc -smp sockets=1,cores=2 -drive file=<absolute-path-to-disc
   -image-file>,if=none,id=drive-ide0-0-0,format=<disc-image-format> \
6
7 -device ide-drive,bus=ide.0,unit=0,drive=drive-ide0-0-0,id=ide0-0-0,
   bootindex=1 \
8
9 -netdev type=tap,id=hostnet0,script=no,downscript=no,vhost=on \
10
11 -device virtio-net-pci,netdev=hostnet0,id=net0,mac=52:54:00:3c:d1:ae,
   bus=pci.0,addr=0x3 \
12
13 -chardev socket,id=char0,path=</usr/local/var/run/openvswitch/vhost-
   user1> \
```

```
14
15 -netdev type=vhost-user,id=mynet1,chardev=char0,vhostforce -device
    virtio-net-pci,mac=00:00:00:00:00:01,netdev=mynet1,mrg_rxbuf=on \
16
17 -chardev socket,id=char1,path=</usr/local/var/run/openvswitch/vhost-
    user2> \
18
19 -netdev type=vhost-user,id=mynet2,chardev=char1,vhostforce -device
    virtio-net
20
21 pci,mac=00:00:00:00:00:02,netdev=mynet2,mrg_rxbuf=on \
22
23 --nographic
```

For RHOS, use the following sample XML file to provision the Citrix ADC VPX instance, by using `virsh`

```
1 <domain type='kvm'>
2
3   <name>dpgk-vpx1</name>
4
5   <uuid>aedb844b-f6bc-48e6-a4c6-36577f2d68d6</uuid>
6
7   <memory unit='KiB'>16777216</memory>
8
9   <currentMemory unit='KiB'>16777216</currentMemory>
10
11  <memoryBacking>
12
13    <hugepages>
14
15      <page size='1048576' unit='KiB' />
16
17    </hugepages>
18
19  </memoryBacking>
20
21  <vcpu placement='static'>6</vcpu>
22
23  <cputune>
24
25    <shares>4096</shares>
26
27    <vcpupin vcpu='0' cpuset='0' />
28
29    <vcpupin vcpu='1' cpuset='2' />
30
31    <vcpupin vcpu='2' cpuset='4' />
32
33    <vcpupin vcpu='3' cpuset='6' />
34
35    <emulatorpin cpuset='0,2,4,6' />
36
```

```
37 </cputune>
38
39 <numatune>
40
41   <memory mode='strict' nodeset='0' />
42
43 </numatune>
44
45 <resource>
46
47   <partition>/machine</partition>
48
49 </resource>
50
51 <os>
52
53   <type arch='x86\_64' machine='pc-i440fx-rhel7.0.0'>hvm</type>
54
55   <boot dev='hd' />
56
57 </os>
58
59 <features>
60
61   <acpi />
62
63   <apic />
64
65 </features>
66
67 <cpu mode='custom' match='minimum' check='full'>
68
69   <model fallback='allow'>Haswell-noTSX</model>
70
71   <vendor>Intel</vendor>
72
73   <topology sockets='1' cores='6' threads='1' />
74
75   <feature policy='require' name='ss' />
76
77   <feature policy='require' name='pcid' />
78
79   <feature policy='require' name='hypervisor' />
80
81   <feature policy='require' name='arat' />
82
83 <domain type='kvm'>
84
85   <name>dpdk-vpx1</name>
86
87   <uuid>aedb844b-f6bc-48e6-a4c6-36577f2d68d6</uuid>
88
89   <memory unit='KiB'>16777216</memory>
```

```
90
91 <currentMemory unit='KiB'>16777216</currentMemory>
92
93 <memoryBacking>
94   <hugepages>
95     <page size='1048576' unit='KiB' />
96   </hugepages>
97 </memoryBacking>
98
99 <vcpu placement='static'>6</vcpu>
100
101 <cputune>
102   <shares>4096</shares>
103   <vcupin vcpu='0' cpuset='0' />
104   <vcupin vcpu='1' cpuset='2' />
105   <vcupin vcpu='2' cpuset='4' />
106   <vcupin vcpu='3' cpuset='6' />
107   <emulatorpin cpuset='0,2,4,6' />
108 </cputune>
109
110 <numatune>
111   <memory mode='strict' nodeset='0' />
112 </numatune>
113
114 <resource>
115   <partition>/machine</partition>
116 </resource>
117
118 <os>
119   <type arch='x86_64' machine='pc-i440fx-rhel7.0.0'>hvm</type>
120   <boot dev='hd' />
121 </os>
122
123 <features>
```

```
143     <acpi/>
144
145     <apic/>
146
147 </features>
148
149 <cpu mode='custom' match='minimum' check='full'>
150
151     <model fallback='allow'>Haswell-noTSX</model>
152
153     <vendor>Intel</vendor>
154
155     <topology sockets='1' cores='6' threads='1'/>
156
157     <feature policy='require' name='ss'/>
158
159     <feature policy='require' name='pcid'/>
160
161     <feature policy='require' name='hypervisor'/>
162
163     <feature policy='require' name='arat'/>
164
165     <feature policy='require' name='tsc\_adjust'/>
166
167     <feature policy='require' name='xsaveopt'/>
168
169     <feature policy='require' name='pdpe1gb'/>
170
171     <numa>
172
173         <cell id='0' cpus='0-5' memory='16777216' unit='KiB' memAccess='
174             shared'/>
175     </numa>
176
177 </cpu>
178
179 <clock offset='utc'/>
180
181 <on\_poweroff>destroy</on\_poweroff>
182
183 <on\_reboot>restart</on\_reboot>
184
185 <on\_crash>destroy</on\_crash>
186
187 <devices>
188
189     <emulator>/usr/libexec/qemu-kvm</emulator>
190
191     <disk type='file' device='disk'>
192
193         <driver name='qemu' type='qcow2' cache='none'/>
194
```

```
195     <source file='/home/NSVPX-KVM-12.0-52.18\_nc.qcow2' />
196
197     <target dev='vda' bus='virtio' />
198
199     <address type='pci' domain='0x0000' bus='0x00' slot='0x07'
200         function='0x0' />
201 </disk>
202
203 <controller type='ide' index='0'>
204
205     <address type='pci' domain='0x0000' bus='0x00' slot='0x01'
206         function='0x1' />
207 </controller>
208
209 <controller type='usb' index='0' model='piix3-uhci'>
210
211     <address type='pci' domain='0x0000' bus='0x00' slot='0x01'
212         function='0x2' />
213 </controller>
214
215 <controller type='pci' index='0' model='pci-root' />
216
217 <interface type='direct'>
218
219     <mac address='52:54:00:bb:ac:05' />
220
221     <source dev='enp129s0f0' mode='bridge' />
222
223     <model type='virtio' />
224
225     <address type='pci' domain='0x0000' bus='0x00' slot='0x03'
226         function='0x0' />
227 </interface>
228
229 <interface type='vhostuser'>
230
231     <mac address='52:54:00:55:55:56' />
232
233     <source type='unix' path='/var/run/openvswitch/vhost-user1' mode=
234         'client' />
235
236     <model type='virtio' />
237
238     <address type='pci' domain='0x0000' bus='0x00' slot='0x04'
239         function='0x0' />
240 </interface>
241 <interface type='vhostuser'>
```

```
242
243     <mac address='52:54:00:2a:32:64' />
244
245     <source type='unix' path='/var/run/openvswitch/vhost-user2' mode=
      'client' />
246
247     <model type='virtio' />
248
249     <address type='pci' domain='0x0000' bus='0x00' slot='0x05'
      function='0x0' />
250
251 </interface>
252
253 <interface type='vhostuser'>
254
255     <mac address='52:54:00:2a:32:74' />
256
257     <source type='unix' path='/var/run/openvswitch/vhost-user3' mode=
      'client' />
258
259     <model type='virtio' />
260
261     <address type='pci' domain='0x0000' bus='0x00' slot='0x06'
      function='0x0' />
262
263 </interface>
264
265 <interface type='vhostuser'>
266
267     <mac address='52:54:00:2a:32:84' />
268
269     <source type='unix' path='/var/run/openvswitch/vhost-user4' mode=
      'client' />
270
271     <model type='virtio' />
272
273     <address type='pci' domain='0x0000' bus='0x00' slot='0x09'
      function='0x0' />
274
275 </interface>
276
277 <serial type='pty'>
278     <target port='0' />
279
280 </serial>
281
282 <console type='pty'>
283     <target type='serial' port='0' />
284
285 </console>
286
287
288
```



```
289     <input type='mouse' bus='ps2' />
290
291     <input type='keyboard' bus='ps2' />
292
293     <graphics type='vnc' port='-1' autoport='yes'>
294         <listen type='address' />
295     </graphics>
296
297     <video>
298         <model type='cirrus' vram='16384' heads='1' primary='yes' />
299         <address type='pci' domain='0x0000' bus='0x00' slot='0x02'
300             function='0x0' />
301     </video>
302
303     <memballoon model='virtio'>
304         <address type='pci' domain='0x0000' bus='0x00' slot='0x08'
305             function='0x0' />
306     </memballoon>
307
308 </devices>
309
310 </domain
```

Points to note

In the XML file, the `hugepage` size must be 1 GB, as shown in the sample file.

```
1 <memoryBacking>
2
3     <hugepages>
4         <page size='1048576' unit='KiB' />
5     </hugepages>
```

Also, in the sample file `vhost-user1` is the `vhost` user port bound to `ovs-br0`.

```
1 <interface type='vhostuser'>
2
3     <mac address='52:54:00:55:55:56' />
4
5     <source type='unix' path='/var/run/openvswitch/vhost-user1' mode=
6         'client' />
```

```
7     <model type='virtio' />
8
9     <address type='pci' domain='0x0000' bus='0x00' slot='0x04'
      function='0x0' />
10
11    </interface>
```

To bring up the Citrix ADC VPX instance, start using the `virsh` command.

Citrix ADC VPX on AWS

September 19, 2024

You can launch a Citrix ADC VPX instance on Amazon Web Services (AWS). The Citrix ADC VPX appliance is available as an Amazon Machine Image (AMI) in AWS marketplace. A Citrix ADC VPX instance on AWS enables you to use AWS cloud computing capabilities and use Citrix ADC load balancing and traffic management features for their business needs. The VPX instance supports all the traffic management features of a physical Citrix ADC appliance, and it can be deployed as standalone instances or in HA pairs. For more information on VPX features, see the [VPX data sheet](#).

Getting started

Before you get started with your VPX deployment, you must be familiar with the following information:

- [AWS terminology](#)
- [AWS-VPX support matrix](#)
- [Limitations and usage guidelines](#)
- [Prerequisites](#)
- [How a Citrix ADC VPX instance on AWS works](#)

Deploy a Citrix ADC VPX instance on AWS

In AWS, the following deployment types are supported for VPX instances:

- [Standalone](#)
- [High availability \(Active-Passive\)](#)
 - [High availability within same zone](#)
 - [High availability across different zones using Elastic IP](#)
 - [High availability across different zones using Private IP](#)

- [Active-Active GSLB](#)
- [Autoscaling \(Active-Active\) using ADM](#)

Hybrid Deployments

- [Deploy Citrix ADC in AWS Outpost](#)
- [Deploy Citrix ADC in VMC in AWS](#)

Licensing

A Citrix ADC VPX instance on AWS requires a license. The following licensing options are available for Citrix ADC VPX instances running on AWS:

- [Free \(unlimited\)](#)
- [Hourly](#)
- [Annual](#)
- [BYOL](#)
- [Free Trial \(all Citrix ADC VPX-AWS subscription offerings for 21 days free in AWS marketplace.\)](#)

Automation

- [Citrix ADM: Smart Deployment](#)
- [GitHub CFTs: Citrix ADC templates and scripts for AWS deployment](#)
- [GitHub Ansible: Citrix ADC templates and scripts for AWS deployment](#)
- [GitHub Terraform: Citrix ADC templates and scripts for AWS deployment](#)
- [AWS Pattern Library \(PL\): Citrix ADC VPX](#)

Blogs

- [How Citrix ADC on AWS Helps Customers Deliver Applications Securely](#)
- [Application delivery in hybrid cloud with Citrix ADC and AWS](#)
- [Citrix is an AWS Networking Competency Partner](#)
- [Citrix ADC: Always ready for public clouds](#)
- [Scale out or scale in with ease in public clouds through Citrix ADC](#)
- [Citrix expands ADC deployment choice with AWS Outposts](#)

- [Using Citrix ADC with Amazon VPC ingress routing](#)
- [Citrix delivers choice, performance, and simplified deployment in AWS](#)
- [The security of Citrix Web App Firewall –now on the AWS Marketplace](#)
- [How Aria Systems uses Citrix Web App Firewall on AWS](#)

Videos

- [Simplifying public cloud NetScaler deployments through ADM](#)
- [Provisioning and configuring NetScaler VPX in AWS using ready-to-use terraform scripts](#)
- [Deploy NetScaler HA in AWS using CloudFormation Template](#)
- [Deploy NetScaler HA across Availability Zones using AWS QuickStart](#)
- [NetScaler Autoscale using ADM](#)

Customer case studies

- [Technology Solution - Xenit AB](#)
- [A better way to do business with Citrix and AWS cloud –Aria](#)
- [Discover the Citrix ADC and AWS advantage](#)
- [Rain for Rent - Customer story](#)

Solutions

- [Deploy digital advertising platform on AWS with Citrix ADC](#)
- [Enhancing Clickstream analytics in AWS using Citrix ADC](#)

Support

- [Open a Support case](#)
- For Citrix ADC subscription offering, see [Troubleshoot a VPX instance on AWS](#). To file a support case, find your AWS account number and support PIN code, and call Citrix support.
- For Citrix ADC Customer Licensed offering or BYOL, ensure that you have the valid support and maintenance agreement. If you do not have an agreement, contact your Citrix representative.

Additional References

- [AWS On-Demand Webinar - Citrix ADC on AWS](#)
- [Citrix ADC VPX data sheet](#)
- [Citrix ADC in AWS Marketplace](#)
- [Citrix ADC is part of AWS networking partner solutions \(load balancers\)](#)
- [AWS FAQs](#)

AWS terminology

September 6, 2024

This section describes the list of commonly used AWS terms and phrases. For more information, see [AWS Glossary](#).

Term	Definition
Amazon Machine Image (AMI)	A machine image, which provides the information required to launch an instance, which is a virtual server in the cloud.
Elastic Block Store	Provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud.
Simple Storage Service (S3)	Storage for the Internet. It is designed to make web-scale computing easier for developers.
Elastic Compute Cloud (EC2)	A web service that provides secure, resizable compute capacity in the cloud. It is designed to make web-scale cloud computing easier for developers.
Elastic Load Balancing (ELB)	Distributes incoming application traffic across multiple EC2 instances, in multiple Availability Zones. This increases the fault tolerance of your applications.
Elastic network interface (ENI)	A virtual network interface that you can attach to an instance in a Virtual Private Cloud (VPC).

Term	Definition
Elastic IP (EIP) address	A static, public IPv4 address that you have allocated in Amazon EC2 or Amazon VPC and then attached to an instance. Elastic IP addresses are associated with your account, not a specific instance. They are elastic because you can easily allocate, attach, detach, and free them as your needs change.
Instance type	Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications.
Identity and Access Management (IAM)	An AWS identity with permission policies that determine what the identity can and cannot do in AWS. You can use an IAM role to enable applications running on an EC2 instance to securely access your AWS resources. IAM role is required for deploying VPX instances in a high-availability setup.
Internet Gateway	Connects a network to the Internet. You can route traffic for IP addresses outside your VPC to the Internet gateway.
Key pair	A set of security credentials that you use to prove your identity electronically. A key pair consists of a private key and a public key.
Route tables	A set of routing rules that controls the traffic leaving any subnet that is associated with the route table. You can associate multiple subnets with a single route table, but a subnet can be associated with only one route table at a time.
Security groups	A named set of allowed inbound network connections for an instance.

Term	Definition
Subnets	A segment of the IP address range of a VPC that EC2 instances can be attached to. You can create subnets to group instances according to security and operational needs.
Virtual Private Cloud (VPC)	A web service for provisioning a logically isolated section of the AWS cloud where you can launch AWS resources in a virtual network that you define.
Auto Scaling	A web service to launch or terminate Amazon EC2 instances automatically based on user-defined policies, schedules, and health checks.
CloudFormation	A service for writing or changing templates that create and delete related AWS resources together as a unit.

VPX-AWS support matrix

September 9, 2024

The following tables list the supported VPX model and AWS regions, instance types, and services.

Table 1: Supported VPX models on AWS

Supported VPX model
NetScaler VPX Advanced - 200 Mbps
NetScaler VPX Premium - 1 Gbps
NetScaler VPX Premium - 5 Gbps
NetScaler VPX - Customer Licensed
NetScaler VPX FIPS - Customer Licensed
NetScaler VPX FIPS ENA - Customer Licensed

Table 2: Supported AWS regions

Supported AWS regions

US West (Oregon)
US West (N. California)
US East (Ohio)
US East (N. Virginia)
Asia Pacific (Mumbai)
Asia Pacific (Seoul)
Asia Pacific (Singapore)
Asia Pacific (Sydney)
Asia Pacific (Tokyo)
Asia Pacific (Hong Kong)
Asia Pacific (Osaka)
Canada (Central)
EU (Frankfurt)
EU (Ireland)
EU (London)
EU (Paris)
EU (Milan)
South America (São Paulo)
AWS GovCloud (US-East)
AWS GovCloud (US-West)
AWS Top Secret (C2S)
Middle East (Bahrain)
Africa (Cape Town)
C2S

Note:

For AWS Hong Kong region, NetScaler VPX support is available only with BYOL licenses.

Table 3: Supported AWS instance types

Supported AWS instance types

t2.medium, t2.large, t2.xlarge, t2.2xlarge

m3.large, m3.xlarge, m3.2xlarge

c4.large, c4.xlarge, c4.2xlarge, c4.4xlarge, c4.8xlarge

m4.large, m4.xlarge, m4.2xlarge, m4.4xlarge, m4.10xlarge, m4.16xlarge

m5.large, m5.xlarge, m5.2xlarge, m5.4xlarge, m5.12xlarge, m5.24xlarge

c5.large, c5.xlarge, c5.2xlarge, c5.4xlarge, c5.9xlarge, c5.18xlarge, c5.24xlarge

C5n.large, C5n.xlarge, C5n.2xlarge, C5n.4xlarge, C5n.9xlarge, C5n.18xlarge

D2.xlarge, D2.2xlarge, D2.4xlarge, D2.8xlarge

m5a.large, m5a.xlarge, m5a.2xlarge, m5a.8xlarge, m5a.12xlarge, m5a.16xlarge, m5a.24xlarge

t3a.medium, t3a.large, t3a.xlarge, t3a.2xlarge

Table 4: Supported AWS Services

Supported AWS services

EC2: Launches ADC instances.

Lambda: Invokes Citrix ADC VPX NITRO APIs during provisioning of Citrix ADC VPX instances from CFT.

VPC and VPC ingress routing: VPC creates isolated networks in which ADC can be launched. VPC ingress routing

Route53: Distributes traffic across all the ADC VPX nodes in the Citrix ADC Autoscale solution.

ELB: Distributes traffic across all the ADC VPX nodes in the Citrix ADC Autoscale solution.

Cloudwatch: Monitors performance and system parameters for Citrix ADC VPX instance.

AWS Autoscaling: Used for back-end server autoscaling.

Cloud formation: CloudFormation templates are used to deploy Citrix ADC VPX instances.

Simple Queue Service (SQS): Monitors scale up and scale down events in back-end autoscaling.

Simple Notification Service (SNS): Monitors scale up and scale down events in back-end autoscaling.

Identity and Access Management (IAM): Provides access to AWS services and resources.

AWS Outposts: Provisions Citrix ADC VPX instances in AWS Outposts.

Citrix recommends the following AWS instance types:

- M5 and C5n series for marketplace editions or bandwidth-based pool licensing.

- C5n series for vCPU-based pool licensing.

VPX offering in AWS marketplace	AWS instance recommendation
VPX 200	M5.xLarge
VPX 1G, VPX 5G	M5.2xLarge

Citrix recommends the following AWS instance types based on throughput.

VPX with Pooled licensing (Bandwidth licenses)	AWS instance recommendation
VPX 8G	C5n.4xLarge
VPX 10G, VPX 15G, VPX 25G	C5n.9xLarge

Note:

The VPX 25G offering doesn't give the desired 25G throughput in AWS but can give higher SSL transactions rate.

To achieve throughput more than 5G, do the following:

- Choose **Citrix ADC VPX - Customer Licensed (BYOL)** offering in AWS marketplace.
- Select **Pooled Licensing (Bandwidth licenses)** in Citrix ADC GUI or CLI.

To determine your instance based on different metrics such as packets per second, SSL transactions rate, reach out to your Citrix contact for guidance. For vCPU based Pool licensing and sizing guidance, reach out to Citrix support.

Limitations and usage guidelines

September 9, 2024

The following limitations and usage guidelines apply when deploying a Citrix ADC VPX instance on AWS:

- Before you start, read the AWS terminology section in [Deploy a Citrix ADC VPX instance on AWS](#).
- The clustering feature is not supported for VPX.

- For the high availability setup to work effectively, associate a dedicated NAT device to management Interface or associate EIP to NSIP. For more information on NAT, in the AWS documentation, see [NAT Instances](#).
- Data traffic and management traffic must be segregated with ENIs belonging to different subnets.
- Only the NSIP address must be present on the management ENI.
- If a NAT instance is used for security instead of assigning an EIP to the NSIP, appropriate VPC level routing changes are required. For instructions on making VPC level routing changes, in the AWS documentation, see [Scenario 2: VPC with Public and Private Subnets](#).
- A VPX instance can be moved from one EC2 instance type to another (for example, from m3.large to an m3.xlarge).
- For storage options for VPX on AWS, Citrix recommends EBS, because it is durable and the data is available even after it is detached from the instance.
- Dynamic addition of ENIs to VPX is not supported. Restart the VPX instance to apply the update. Citrix recommends you to stop the standalone or HA instance, attach the new ENI, and then restart the instance.
- You can assign multiple IP addresses to an ENI. The maximum number of IP addresses per ENI is determined by the EC2 instance type, see the section “IP Addresses Per Network Interface Per Instance Type” in [Elastic Network Interfaces](#). You must allocate the IP addresses in AWS before you assign them to ENIs. For more information, see [Elastic Network Interfaces](#).
- Citrix recommends that you avoid using the enable and disable interface commands on Citrix ADC VPX interfaces.
- The Citrix ADC `set ha node \<NODE_ID\> -haStatus STAYPRIMARY` and `set ha node \<NODE_ID\> -haStatus STAYSECONDARY` commands are disabled by default.
- IPv6 is not supported for VPX.
- Due to AWS limitations, these features are not supported:
 - Gratuitous ARP(GARP)
 - L2 mode
 - Tagged VLAN
 - Dynamic Routing
 - virtual MAC
- For RNAT to work, ensure **Source/Destination Check** is disabled. For more information, see “Changing the Source/Destination Checking” in [Elastic Network Interfaces](#).

- In a Citrix ADC VPX deployment on AWS, in some AWS regions, the AWS infrastructure might not be able to resolve AWS API calls. This happens if the API calls are issued through a nonmanagement interface on the Citrix ADC VPX instance.

As a workaround, restrict the API calls to the management interface only. To do that, create an NSVLAN on the VPX instance and bind the management interface to the NSVLAN by using the appropriate command.

For example:

```
set ns config -nsvlan <vlan id> -ifnum 1/1 -tagged NO
save config
```

Restart the VPX instance at the prompt. For more information about configuring `nsvlan`, see [Configuring NSVLAN](#).

- In the AWS console, the vCPU usage shown for a VPX instance under the **Monitoring** tab might be high (up to 100 percent), even when the actual usage is much lower. To see the actual vCPU usage, navigate to **View all CloudWatch metrics**. For more information, see [Monitor your instances using Amazon CloudWatch](#).

Prerequisites

September 9, 2024

Before attempting to create a VPX instance in AWS, ensure you have the following:

- **An AWS account:** to launch a Citrix ADC VPX AMI in an AWS Virtual Private Cloud (VPC). You can create an AWS account for free at [Amazon website] (<http://www.aws.amazon.com>).
- **An AWS Identity and Access Management (IAM) user account:** to securely control access to AWS services and resources for your users. For more information about how to create an IAM user account, see [Creating IAM Users \(Console\)](#). An IAM role is mandatory for both standalone and high availability deployments.

The IAM role associated with your AWS account must have the following IAM permissions for various scenarios.

HA pair in the same AWS zone:

```
1  "ec2:DescribeInstances",
2  "ec2:AssignPrivateIpAddresses",
3  "iam:SimulatePrincipalPolicy",
4  "iam:GetRole"
```

HA pair with elastic IP addresses across different AWS zones:

```

1  "ec2:DescribeInstances",
2  "ec2:DescribeAddresses",
3  "ec2:AssociateAddress",
4  "ec2:DisassociateAddress",
5  "iam:SimulatePrincipalPolicy",
6  "iam:GetRole"

```

HA pair with private IP addresses across different AWS zones:

```

1  "ec2:DescribeInstances",
2  "ec2:DescribeRouteTables",
3  "ec2:DeleteRoute",
4  "ec2:CreateRoute",
5  "ec2:ModifyNetworkInterfaceAttribute",
6  "iam:SimulatePrincipalPolicy",
7  "iam:GetRole"

```

HA pair with both private IP addresses and elastic IP addresses across different AWS zones:

```

1  "ec2:DescribeInstances",
2  "ec2:DescribeAddresses",
3  "ec2:AssociateAddress",
4  "ec2:DisassociateAddress",
5  "ec2:DescribeRouteTables",
6  "ec2:DeleteRoute",
7  "ec2:CreateRoute",
8  "ec2:ModifyNetworkInterfaceAttribute",
9  "iam:SimulatePrincipalPolicy",
10 "iam:GetRole"

```

AWS backend autoscaling:

```

1  "ec2:DescribeInstances",
2  "autoscaling:*",
3  "sns:CreateTopic",
4  "sns:DeleteTopic",
5  "sns:ListTopics",
6  "sns:Subscribe",
7  "sqs:CreateQueue",
8  "sqs:ListQueues",
9  "sqs:DeleteMessage",
10 "sqs:GetQueueAttributes",
11 "sqs:SetQueueAttributes",
12 "iam:SimulatePrincipalPolicy",
13 "iam:GetRole",

```

Note:

- If you use any combination of the preceding features, use the combination of IAM permissions for each of the features.
- If you use the Citrix CloudFormation template, the IAM role is automatically created.

The template does not allow selecting an already created IAM role.

- When you log on to the VPX instance through the GUI, a prompt to configure the required privileges for the IAM role appears. Ignore the prompt if you've already configured the privileges.

- **AWS CLI:** To use all the functionality provided by the AWS Management Console from your terminal program. For more information, see the [AWS CLI user guide](#). You also need the AWS CLI to change the network interface type to SR-IOV.
- **Elastic Network Adapter (ENA):** For ENA driver-enabled instance type, for example M5, C5 instances, the firmware version must be 13.0 and later.

How a Citrix ADC VPX instance on AWS works

September 12, 2024

The Citrix ADC VPX instance is available as an AMI in AWS marketplace, and it can be launched as an EC2 instance within an AWS VPC. The Citrix ADC VPX AMI instance requires a minimum of 2 virtual CPUs and 2 GB of memory. An EC2 instance launched within an AWS VPC can also provide the multiple interfaces, multiple IP addresses per interface, and public and private IP addresses needed for VPX configuration. Each VPX instance requires at least three IP subnets:

- A management subnet
- A client-facing subnet (VIP)
- A back-end facing subnet (SNIP, MIP, and so on)

Citrix recommends three network interfaces for a standard VPX instance on AWS installation.

AWS currently makes multi-IP functionality available only to instances running within an AWS VPC. A VPX instance in a VPC can be used to load balance servers running in EC2 instances. An Amazon VPC allows you to create and control a virtual networking environment, including your own IP address range, subnets, route tables, and network gateways.

Note:

By default, you can create up to 5 VPC instances per AWS region for each AWS account. You can request higher VPC limits by submitting Amazon's request form <http://aws.amazon.com/contact-us/vpc-request>.

Figure 1. A Sample Citrix ADC VPX Instance Deployment on AWS Architecture




Figure 1 shows a simple topology of an AWS VPC with a Citrix ADC VPX deployment. The AWS VPC has:

1. A single Internet gateway to route traffic in and out of the VPC.
2. Network connectivity between the Internet gateway and the Internet.
3. Three subnets, one each for management, client, and server.
4. Network connectivity between the Internet gateway and the two subnets (management and client).
5. A standalone Citrix ADC VPX instance deployed within the VPC. The VPX instance has three ENIs, one attached to each subnet.

Deploy a Citrix ADC VPX standalone instance on AWS

September 12, 2024

You can deploy a Citrix ADC VPX standalone instance on AWS by using the following options:

- AWS web console
- Citrix-authored CloudFormation template
- AWS CLI

This topic describes the procedure for deploying a Citrix ADC VPX instance on AWS.

Before you start your deployment, read the following topics:

- [Prerequisites](#)
- [Limitation and usage guidelines](#)

Deploy a Citrix ADC VPX instance on AWS by using the AWS web console

You can deploy a Citrix ADC VPX instance on AWS through the AWS web console. The deployment process includes the following steps:

1. Create a Key Pair
2. Create a Virtual Private Cloud (VPC)
3. Add more subnets
4. Create security groups and security rules
5. Add route tables
6. Create an internet gateway
7. Create a Citrix ADC VPX instance
8. Create and attach more network interfaces
9. Attach elastic IPs to the management NIC
10. Connect to the VPX instance

Step 1: Create a key pair.

Amazon EC2 uses a key pair to encrypt and decrypt logon information. To log on to your instance, you must create a key pair, specify the name of the key pair when you launch the instance, and provide the private key when you connect to the instance.

When you review and launch an instance by using the AWS Launch Instance wizard, you are prompted to use an existing key pair or create a new key pair. For more information about how to create a key pair, see [Amazon EC2 Key Pairs](#).

Step 2: Create a VPC.

A Citrix ADC VPC instance is deployed inside an AWS VPC. A VPC allows you to define the virtual network dedicated to your AWS account. For more information about AWS VPC, see [Getting Started With Amazon VPC](#).

While creating a VPC for your Citrix ADC VPX instance, keep the following points in mind.

- Use the VPC with a Single Public Subnet Only option to create an AWS VPC in an AWS availability zone.
- Citrix recommends that you create at least **three subnets**, of the following types:
 - One subnet for management traffic. You place the management IP(NSIP) on this subnet. By default elastic network interface (ENI) eth0 is used for management IP.
 - One or more subnets for client-access (user-to-Citrix ADC VPX) traffic, through which clients connect to one or more virtual IP (VIP) addresses assigned to Citrix ADC load balancing virtual servers.
 - One or more subnets for the server-access (VPX-to-server) traffic, through which your servers connect to VPX-owned subnet IP (SNIP) addresses. For more information about Citrix ADC load balancing and virtual servers, virtual IP addresses (VIPs), and subnet IP addresses (SNIPs), see:
 - All subnets must be in the same availability zone.

Step 3: Add subnets.

When you used the VPC wizard, only one subnet was created. Depending on your requirement, you might want to create more subnets. For more information about how to create more subnets, see [Adding a Subnet to Your VPC](#).

Step 4: Create security groups and security rules.

To control inbound and outbound traffic, create security groups and add rules to the groups. For more information how to create groups and add rules, see [Security Groups for Your VPC](#).

For Citrix ADC VPX instances, the EC2 wizard gives default security groups, which are generated by AWS Marketplace and is based on recommended settings by Citrix. However, you can create more security groups based on your requirements.

Note:

Port 22, 80, 443 to be opened on the Security group for SSH, HTTP, and HTTPS access respectively.

Step 5: Add route tables.

Route table contains a set of rules, called routes, that are used to determine where network traffic is directed. Each subnet in your VPC must be associated with a route table. For more information about how to create a route table, see [Route Tables](#).

Step 6: Create an internet gateway.

An internet gateway serves two purposes: to provide a target in your VPC route tables for internet-routable traffic, and to perform network address translation (NAT) for instances that have been assigned public IPv4 addresses.


Create an internet gateway for internet traffic. For more information about how to create an Internet Gateway, see the section [Attaching an Internet Gateway](#).

Step 7: Create a Citrix ADC VPX instance by using the AWS EC2 service.

To create a Citrix ADC VPX instance by using the AWS EC2 service, complete the following steps.

1. From the AWS dashboard, go to **Compute > EC2 > Launch Instance > AWS Marketplace**.


Before you click **Launch Instance**, ensure your region is correct by checking the note that appears under **Launch Instance**.



2. In the Search AWS Marketplace bar, search with the keyword Citrix ADC VPX.
3. Select the version you want to deploy and then click **Select**. For the Citrix ADC VPX version, you have the following options:
 - A licensed version
 - Citrix ADC VPX Express appliance (This is a free virtual appliance, which is available from Citrix ADC 12.0 56.20.)
 - Bring your own device

The Launch Instance wizard starts. Follow the wizard to create an instance. The wizard prompts you to:

- Choose Instance Type
- Configure Instance
- Add Storage
- Add Tags
- Configure Security Group
- Review



Step 8: Create and attach more network interfaces.

Create two more network interfaces for VIP and SNIP. For more information about how to create more network interfaces, see the [Creating a Network Interface](#) section.

After you’ve created the network interfaces, you must attach them to the VPX instance. Before attaching the interface, shut down the VPX instance, attach the interface, and power on the instance.

For more information about how to attach network interfaces, see the [Attaching a Network Interface When Launching an Instance](#) section.

Step 9: Allocate and associate elastic IPs.

If you assign a public IP address to an EC2 instance, it remains assigned only until the instance is stopped. After that, the address is released back to the pool. When you restart the instance, a new public IP address is assigned.

In contrast, an elastic IP (EIP) address remains assigned until the address is disassociated from an instance.

Allocate and associate an elastic IP for the management NIC. For more information about how to allocate and associate elastic IP addresses, see these topics:

- [Allocating an Elastic IP Address](#)
- [Associating an Elastic IP Address with a Running Instance](#)

These steps complete the procedure to create a Citrix ADC VPX instance on AWS. It can take a few minutes for the instance to be ready. Check that your instance has passed its status checks. You can view this information in the **Status Checks** column on the Instances page.

Step 10: Connect to the VPX instance.

After you've created the VPX instance, you connect the instance by using the GUI and an SSH client.

- GUI

The following are the default administrator credentials to access a Citrix ADC VPX instance

User name: `nsroot`

Password: The default password for the ns root account is set to the AWS instance-ID of the Citrix ADC VPX instance. On your first logon, you are prompted to change the password for security reasons. After changing the password, you must save the configuration. If the configuration is not saved and the instance restarts, you must log on with the default password. Change the password again at the prompt.

- SSH client

From the AWS management console, select the Citrix ADC VPX instance and click **Connect**. Follow the instructions given on the **Connect to Your Instance** page.

For more information about how to deploy a Citrix ADC VPX standalone instance on AWS by using the AWS web console, see [Scenario: standalone instance](#)

Configure a Citrix ADC VPX instance by using the Citrix CloudFormation template

You can use the Citrix-provided CloudFormation template to automate VPX instance launch. The template provides functionality to launch a single Citrix ADC VPX instance, or to create a high availability environment with a pair of Citrix ADC VPX instances.

You can launch the template from AWS Marketplace or GitHub.

The CloudFormation template requires an existing VPC environment, and it launches a VPX instance with three elastic network interfaces (ENIs). Before you start the CloudFormation template, ensure that you complete the following requirements:

- An AWS virtual private cloud (VPC)
- Three subnets within the VPC: one for management, one for client traffic, and one for back-end servers
- An EC2 key pair to enable SSH access to the instance
- A security group with UDP 3003, TCP 3009–3010, HTTP, SSH ports open

See the “Deploy a Citrix ADC VPX Instance on AWS by Using the AWS Web Console” section or AWS documentation for more information about how to complete the prerequisites.

Further, you configure and launch a Citrix ADC VPX Express standalone instance by using the Citrix CloudFormation template available in GitHub:

<https://github.com/citrix/citrix-adc-aws-cloudformation/tree/master/templates/standalone/>

An IAM role is not mandatory for a standalone deployment. However, Citrix recommends that you create and attach an IAM role with the required privileges to the instance, for future need. The IAM role ensures that the standalone instance is easily converted to a high availability node with SR-IOV, when required.

For more information about the required privileges, see [Configuring Citrix ADC VPX instances to Use SR-IOV Network Interface](#).

Note:

If you deploy a Citrix ADC VPX instance on AWS by using the AWS web console, the CloudWatch service is enabled by default. If you deploy a Citrix ADC VPX instance by using the Citrix CloudFormation template, the default option is “Yes.” If you want to disable the CloudWatch service, select “No.” For more information, see [Monitor your instances using Amazon CloudWatch](#)

Configure a Citrix ADC VPX instance by using the AWS CLI

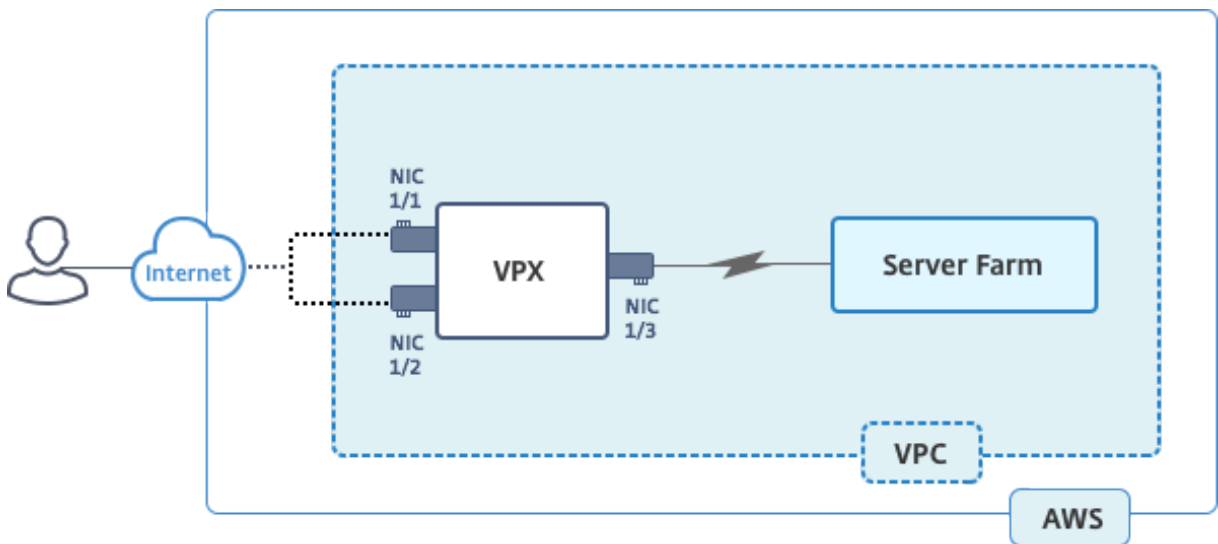
You can use the AWS CLI to launch instances. For more information, see the [AWS Command Line Interface Documentation](#).

Scenario: standalone instance

September 9, 2024

This scenario illustrates how to deploy a Citrix ADC VPX standalone EC2 instance in AWS by using the AWS GUI. Create a standalone VPX instance with three NICs. The instance, which is configured as a load balancing virtual server, communicates with back-end servers (the server farm). For this configuration, set up the required communication routes between the instance and the back-end servers, and between the instance and the external hosts on the public internet.

For more details about the procedure for deploying a VPX instance, see [Deploy a Citrix ADC VPX standalone instance on AWS](#).



Create three NICs. Each NIC can be configured with a pair of IP addresses (public and private). The NICs serve the following purposes.

NIC	Purpose	Associated with
eth0	Serves management traffic (NSIP)	A public IP address and a private IP address
eth1	Serves client-side traffic (VIP)	A public IP address and a private IP address
eth2	Communicates with back-end servers (SNIP)	A public IP address (Private IP address not mandatory)

Step 1: Create a VPC.

1. Log on to the AWS web console and navigate to **Networking & Content Delivery > VPC**. Click **Start VPC Wizard**.

2. Select **VPC with a Single Public Subnet** and click **Select**.
3. Set the IP CIDR Block to 10.0.0.0/16, for this scenario.
4. Give a name for the VPC.
5. Set the public subnet to 10.0.0.0/24. (This is the management network).
6. Select an availability zone.
7. Give a name for the subnet.
8. Click Create **VPC**.

The screenshot shows the AWS VPC console configuration page for 'Step 2: VPC with a Single Public Subnet'. The form includes the following fields and options:

- IPv4 CIDR block:** 10.0.0.0/16 (65531 IP addresses available)
- IPv6 CIDR block:** No IPv6 CIDR Block, Amazon provided IPv6 CIDR block
- VPC name:** NSDoc
- Public subnet's IPv4 CIDR:** 10.0.0.0/24 (251 IP addresses available)
- Availability Zone:** ap-south-1a
- Subnet name:** NSDoc-MGMT
- Service endpoints:** Add Endpoint button
- Enable DNS hostnames:** Yes, No
- Hardware tenancy:** Default

At the bottom right, there are three buttons: 'Cancel and Exit', 'Back', and 'Create VPC' (highlighted with a red border).

Step 2: Create extra subnets.

1. Open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. In the navigation pane, choose Subnets, Create Subnet after you enter the following details.
 - Name tag: Provide a name for your subnet.
 - VPC: Choose the VPC for which you're creating the subnet.
 - Availability Zone: Choose the availability zone in which you created the VPC in step 1.
 - IPv4 CIDR block: Specify an IPv4 CIDR block for your subnet. For this scenario, choose 10.0.1.0/24.

Create Subnet ✕

Use the CIDR format to specify your subnet's IP address block (e.g., 10.0.0.0/24). Note that block sizes must be between a /16 netmask and /28 netmask. Also, note that a subnet can be the same size as your VPC. An IPv6 CIDR block must be a /64 CIDR block.

Name tag ⓘ

VPC ⓘ

VPC CIDRs

CIDR	Status	Status Reason
10.0.0.0/16	associated	

Availability Zone ⓘ

IPv4 CIDR block ⓘ

Cancel
Yes, Create

3. Repeat the steps to create one more subnet for back-end servers.

Create Subnet ✕

Use the CIDR format to specify your subnet's IP address block (e.g., 10.0.0.0/24). Note that block sizes must be between a /16 netmask and /28 netmask. Also, note that a subnet can be the same size as your VPC. An IPv6 CIDR block must be a /64 CIDR block.

Name tag ⓘ

VPC ⓘ

VPC CIDRs

CIDR	Status	Status Reason
10.0.0.0/16	associated	

Availability Zone ⓘ

IPv4 CIDR block ⓘ

Cancel
Yes, Create

Step 3: Create a route table.

1. Open the Amazon VPC console at <https://console.aws.amazon.com/vpc/>.
2. In the navigation pane, choose **Route Tables > Create Route Table**.
3. In the Create Route Table window, add a name and select the VPC that you created in step 1.
4. Click **Yes, Create**.

Create Route Table ✕

A route table specifies how packets are forwarded between the subnets within your VPC, the Internet, and your VPN connection.

Name tag ⓘ

VPC ⓘ

Cancel
Yes, Create

The route table is assigned to all the subnets that you created for this VPC, so that routing of traffic from an instance in one subnet can reach an instance in another subnet.

5. Click Subnet Associations, and then click Edit.
6. Click the management and client subnet and click Save. This creates a route table for internet traffic only.

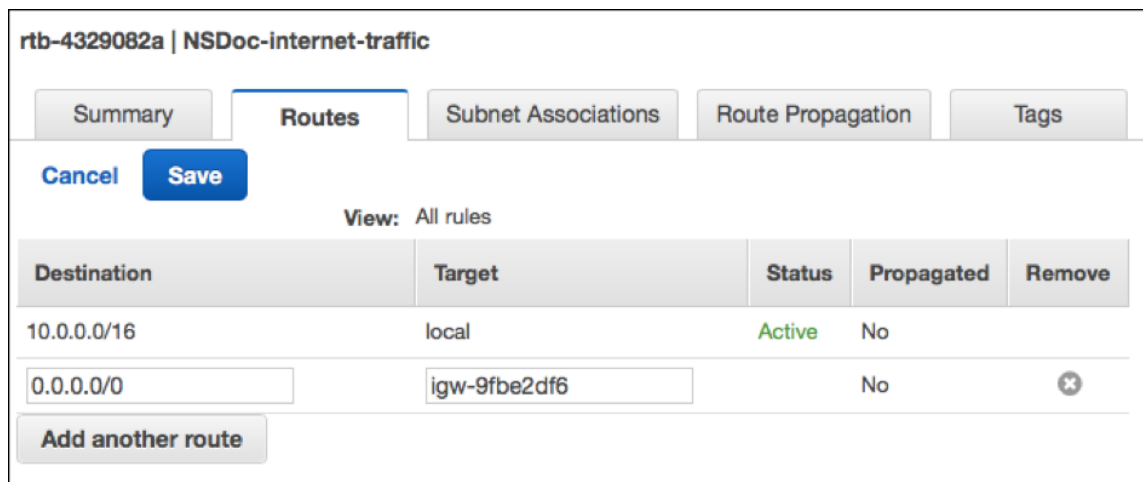
rtb-4329082a | NSDoc-internet-traffic

Summary
Routes
Subnet Associations
Route Propagation
Tags

Cancel
Save

Associate	Subnet	IPv4 CIDR	IPv6 CIDR	Current Route Table
<input checked="" type="checkbox"/>	subnet-c4ce9aad NSDoc-MGMT	10.0.0.0/24	-	rtb-735a7b1a
<input checked="" type="checkbox"/>	subnet-31ce9a58 NSDoc-client	10.0.1.0/24	-	Main
<input type="checkbox"/>	subnet-d0cd99b9 NSDoc-server	10.0.2.0/24	-	Main


7. Click **Routes > Edit > Add another route**.
8. In the Destination field add 0.0.0.0/0, and click the Target field to select igw-<xxxx> the Internet Gateway that the VPC Wizard created automatically.
9. Click Save.



10. Follow the steps to create a route table for server-side traffic.

Step 4: Create a Citrix ADC VPX instance.

1. Log on the AWS management console and click **EC2** under **Compute**.
2. Click AWS Marketplace. In the Search AWS Marketplace bar, type Citrix ADC VPX and press Enter. The available Citrix ADC VPX editions are displayed.
3. Click **Select** to choose the desired Citrix ADC VPX edition. The EC2 instance wizard starts.
4. In the **Choose Instance Type** page, select **m4. Xlarge** (recommended) and click **Next: Configure Instance Details**.
5. In the Configure Instance Details page, select the following, and then click Next: Add Storage.
 - Number of instances: 1
 - Network: the VPC that created in Step 1
 - Subnet: the management subnet
 - Auto-assign Public IP: Enable



6. In the Add Storage page, select the default option, and click Next: Add Tags.
7. In the Add Tags page, add a name for the instance, and click Next: Configure Security Group.
8. In the Configure Security Group page, select the default option (which is generated by AWS Marketplace and is based on recommended settings by Citrix Systems) and then click **Review and Launch > Launch**.
9. You are prompted to select an existing key pair or create a new key pair. From the Select a key pair drop-down list, select the key pair that you created as a prerequisite (See the Prerequisite section.)
10. Check the box to acknowledge the key pair and click Launch Instances.

Select an existing key pair or create a new key pair ✕

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

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

Choose an existing key pair ⌵

Select a key pair

NSDOCKeypair ⌵

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

Cancel Launch Instances

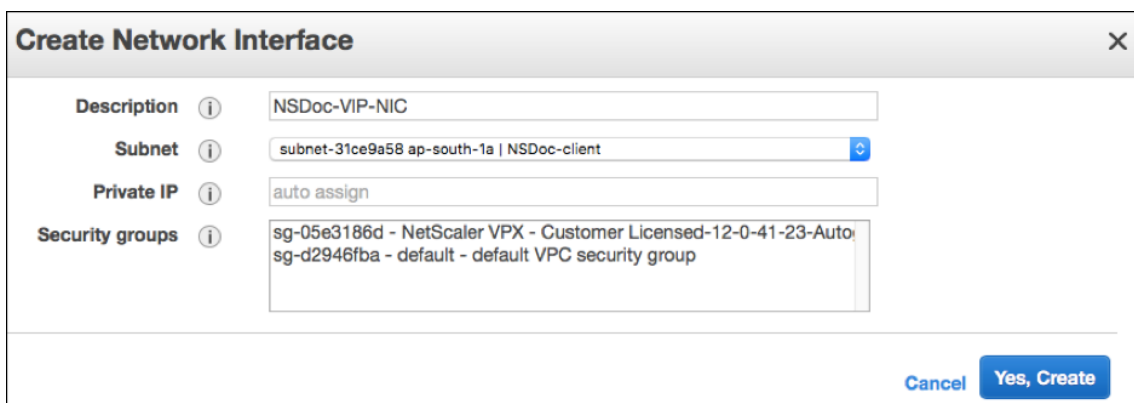
Launch Instance Wizard displays the Launch Status, and the instance appears in the list of instances when it is fully launched.

The check instance, go the AWS console click EC2 > Running Instances. Select the instance and add a name. Make sure the Instance State is running and Status Checks is complete.

Step 5: Create and attach more network interfaces.

When you created the VPC, only one network interface associated with it. Now add two more network interfaces to the VPC, for the VIP and SNIP.

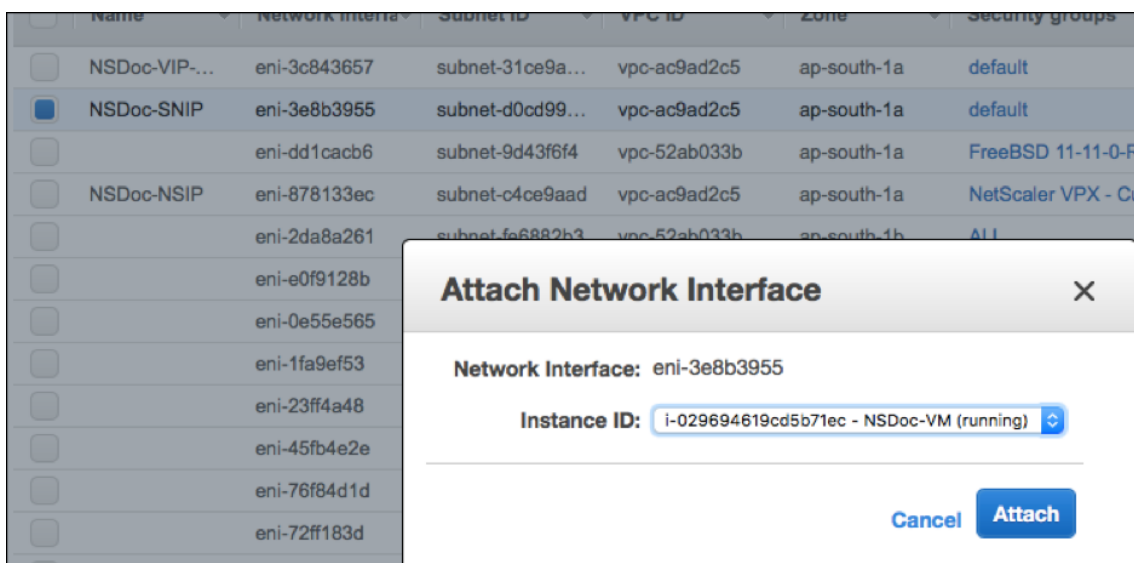
1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose Network Interfaces.
3. Choose Create Network Interface.
4. For Description, enter a descriptive name.
5. For Subnet, select the subnet that you created previously for the VIP.
6. For Private IP, leave the default option.
7. For Security groups, select the group.
8. Click **Yes, Create**.



9. After the network interface is created, add a name to the interface.
10. Repeat the steps to create a network interface for server-side traffic.

Attach the network interfaces:


1. In the navigation pane, choose Network Interfaces.
2. Select the network interface and choose Attach.
3. In the Attach Network Interface dialog box, select the instance and choose Attach.



Step 6: Attach an elastic IP to the NSIP.

1. From the AWS management console, go to **NETWORK & SECURITY > Elastic IPs**.
2. Check for available free EIP to attach. If none, click **Allocate new address**.
3. Select the newly allocated IP address and choose **Actions > Associate address**.
4. Click the **Network interface** radio button.
5. From the Network interface drop-down list, select the management NIC.

6. From the **Private IP** drop-down menu, select the AWS-generated IP address.
7. Select the **Reassociation** check box.
8. Click **Associate**.



Access the VPX instance:

After you’ve configured a standalone Citrix ADC VPX instance with three NICs, log on to the VPX instance to complete the Citrix ADC-side configuration. Use of the following options:

- GUI: Type the public IP of the management NIC in the browser. Log on by using `nsroot` as the user name and the instance ID (`i-0c1ffe1d987817522`) as the password.

Note:

On your first logon, you are prompted to change the password for security reasons. After changing the password, you must save the configuration. If the configuration is not saved and the instance restarts, you must log on with the default password. Change the password again at the prompt and save the configuration.

- SSH: Open an SSH client and type:

```
ssh -i \<location of your private key\> ns root@\<public DNS of the instance\>
```

To find the public DNS, click the instance, and click **Connect**.

Related information:

- To configure the Citrix ADC-owned IP addresses (NSIP, VIP, and SNIP), see [Configuring Citrix ADC-Owned IP Addresses](#).
- You’ve configured a BYOL version of the Citrix ADC VPX appliance, for more information see the VPX Licensing Guide at <http://support.citrix.com/article/CTX122426>

Download a Citrix ADC VPX license

September 12, 2024

After the launch of Citrix ADC VPX-customer licensed instance from the AWS marketplace, a license is required. For more information on VPX licensing, see [Licensing overview](#).

You have to:

1. Use the licensing portal within the Citrix website to generate a valid license.
2. Upload the license to the instance.

If this is a **paid** marketplace instance, then you do not need to install a license. The correct feature set and performance activate automatically.

If you use a Citrix ADC VPX instance with a model number higher than VPX 5000, the network throughput might not be the same as specified by the instance's license. However, other features, such as SSL throughput and SSL transactions per second, might improve.

5 Gbps network bandwidth is observed in the `c4.8xlarge` instance type.

How to migrate the AWS subscription to BYOL


This section describes the procedure to migrate from AWS subscription to Bring your own license (BYOL), and conversely.

Do the following steps to migrate an AWS subscription to BYOL:

Note:

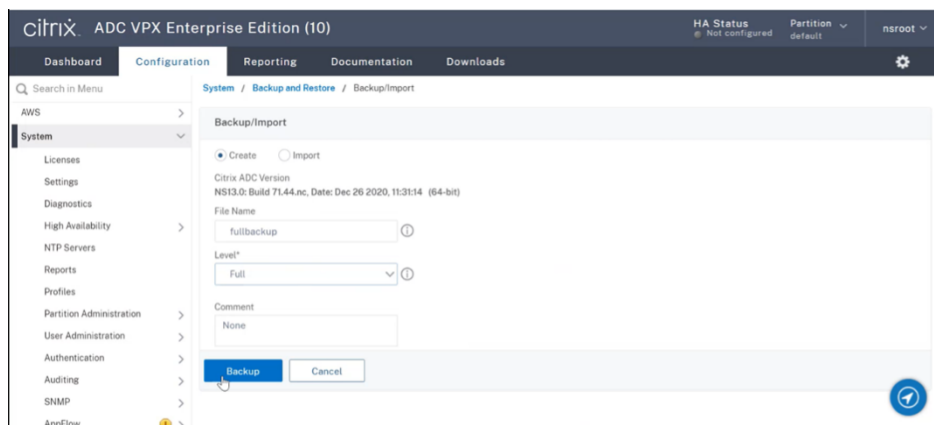
The **Step 2** and **Step 3** are done on the Citrix ADC VPX instance, and all other steps are done on the AWS portal.

1. Create a BYOL EC2 instance using [Citrix ADC VPX - Customer Licensed](#) in the same availability zone as the old EC2 instance that has the same security group, IAM role, and subnet. The new EC2 instance must have only one ENI interface.
2. To back up the data on the old EC2 instance using the Citrix ADC GUI, follow these steps.
 - a) Navigate to **System > Backup and Restore**.
 - b) In the **Welcome** page, click **Backup/Import** to start the process.




c) In the **Backup/Import** page, fill in the following details:

- **Name** –Name of the backup file.
- **Level** –Select the backup level as **Full**.
- **Comment** –Provide a brief description of the backup.




d) Click **Backup**. Once the backup is complete, you can select the file and download it to your local machine.




3. To restore the data on the new EC2 instance using the Citrix ADC GUI, follow these steps:

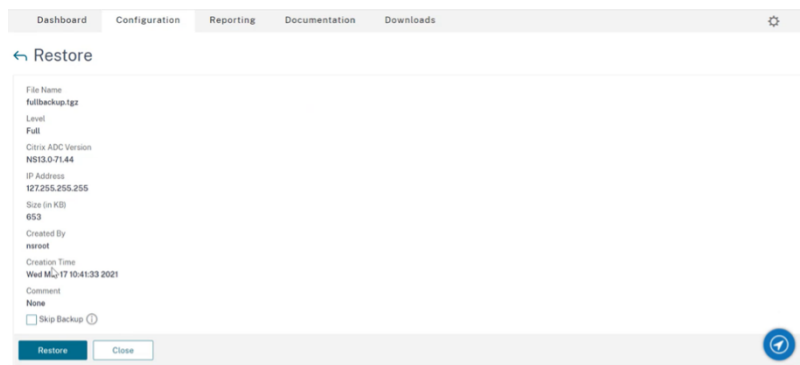
- a) Navigate to **System > Backup and Restore**.
- b) Click **Backup/Import** to start the process.
- c) Select the **Import** option and upload the backup file.



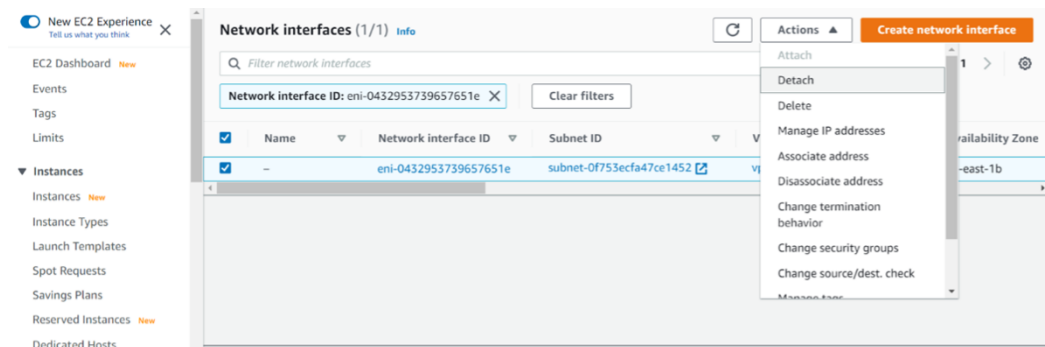
- d) Select the file.
- e) From the **Select Action** drop-down menu, select **Restore**.




- f) On the **Restore** page, verify the file details, and click **Restore**.



- g) After the restore, reboot the EC2 instance.
4. Move all interfaces (except the management interface to which the NSIP address is bound) from the old EC2 instance to the new EC2 instance. To move a network interface from one EC2 instance to another, follow these steps:
- a) In the **AWS Portal**, stop both the old and new EC2 instances.
 - b) Navigate to **Network Interfaces**, and select the network interface attached to the old EC2 instance.
 - c) Detach the EC2 instance by clicking **Actions > Detach**.




- d) Attach the network interface to the new EC2 instance by clicking **Actions > Attach**. Enter the EC2 instance name to which the network interface must be attached.



- e) Do the **Step 1 to Step 4** for all other interfaces that are attached. Make sure to follow the

sequence and maintain the interface order. That is, first detach interface 2 and attach it, and then detach interface 3 and attach it, and so on.

5. You can't detach the management interface from an old EC2 instance. So, move all the secondary IP addresses (if any) on the management interface (primary network interface) of the old EC2 instance to the new EC2 instance. To move an IP address from one interface to another, follow these steps:
 - a) In the **AWS Portal**, make sure that both the old and new EC2 instances are in **Stop** state.
 - b) Navigate to **Network Interfaces**, and select the management network interface attached to the old EC2 instance.
 - c) Click **Actions > Manage IP Address**, and make note of all the secondary IP addresses assigned (if any).
 - d) Navigate to the management network interface or primary interface of the new EC2 instance.
 - e) Click **Actions > Manage IP Addresses**.
 - f) Under **IPv4 Addresses**, click **Assign new IP address**.
 - g) Enter the IP addresses, which are noted in the **Step 3**.
 - h) Select **Allow secondary private IP addresses to be reassigned** check box.
 - i) Click **Save**.



6. Start the new EC2 instance and verify the configuration. After all the configuration is moved, you can delete or keep the old EC2 instance as per your requirement.
7. If any EIP address is attached to the NSIP address of the old EC2 instance, move the old instance NSIP address to the new instance NSIP address.

8. If you want to revert to the old instance, then follow the same steps in the opposite way between the old and new instance.
9. After you move from subscription instance to BYOL instance, a license is required. To install a license follow these steps:
 - Use the licensing portal in the Citrix website to generate a valid license.
 - Upload the license to the instance.

Note:

When you move BYOL instance to subscription instance (paid marketplace instance), you need not install the license. The correct feature set and performance is automatically activated.

Limitations

The management interface can't be moved to the new EC2 instance. So Citrix recommends you manually configure the management interface. For more information, see **Step 5** in the preceding procedure. A new EC2 instance is created with the exact replica of the old EC2 instance but only the NSIP address has a new IP address.

Load balancing servers in different availability zones

September 9, 2024

A VPX instance can be used to load balance servers running in the same availability zone, or in:

- A different availability zone (AZ) in the same AWS VPC
- A different AWS region
- AWS EC2 in a VPC

To enable a VPX instance to load balance servers running outside the AWS VPC that the VPX instance is in, configure the instance to use EIPs to route traffic through the Internet gateway, as follows:

1. Configure a SNIP on the Citrix ADC VPX instance by using the Citrix ADC CLI or the GUI.
2. Enable traffic to be routed out of the AZ, by creating a public facing subnet for the server-side traffic.
3. Add an Internet gateway route to the routing table, using the AWS GUI console.
4. Associate the routing table you updated with the server-side subnet.
5. Associate an EIP with the server-side private IP address that is mapped to a Citrix ADC SNIP address.

How high availability on AWS works

September 9, 2024

You can configure two Citrix ADC VPX instances on AWS as a high availability (HA) active-passive pair. When you configure one instance as the primary node and the other as the secondary node, the primary node accepts connections and manages servers. The secondary node monitors the primary. If for any reason, the primary node is unable to accept connections, the secondary node takes over.

In AWS, the following deployment types are supported for VPX instances:

- High availability within same zone
- High availability across different zones

Note:

For high availability to work, ensure both the Citrix ADC VPX instances are attached with IAM roles and assigned with the Elastic IP (EIP) address to the NSIP. You need not assign an EIP on NSIP if the NSIP can reach internet through the NAT instance.

High availability within the same zones

In a high-availability deployment within the same zones, both VPX instances must have similar networking configurations.

Follow these two rules:

Rule 1. Any NIC on one VPX instance must be in the same subnet as the corresponding NIC in the other VPX. Both instances must have:

- Management interface on the same subnet (referred as management subnet)
- Client interface on the same subnet (referred as client subnet)
- Server interface on the same subnet (referred as server subnet)

Rule 2. Sequence of mgmt NIC, client NIC, and server NIC on both instances must be the same.

For example, the following scenario is not supported.

VPX instance 1

NIC 0: management

NIC 1: client

NIC 2: Server

VPX instance 2

NIC 0: management

NIC 1: server

NIC 2: client

In this scenario, NIC 1 of instance 1 is in client subnet while NIC 1 of instance 2 is in server subnet. For HA to work, NIC 1 of both the instances must be either in the client subnet or in the server subnet.

From 13.0 41.xx, high availability can be achieved by migrating secondary private IP addresses attached to the NICs (client and server-side NICs) of the primary HA node to the secondary HA node after failover. In this deployment:

- Both the VPX instances have the same number of NICs and subnet mapping according to NIC enumeration.
- Each VPX NIC has one extra private IP address, except the first NIC - which corresponds to the management IP address. The extra private IP address appears as the primary private IP address in the AWS web console. In our document, we refer to this extra IP address as the dummy IP address).
- The dummy IP addresses must be not configured on the Citrix ADC instance as VIP and SNIP.
- Other secondary private IP addresses must be created, as required, and configured as VIP and SNIP.
- On failover, the new primary node looks for configured SNIPs and VIPs and moves them from NICs attached to the previous primary to corresponding NICs on the new primary.
- Citrix ADC instances require IAM permissions for HA to work. Add the following IAM privileges to the IAM policy added to each instance.

```
"iam:GetRole"  
"ec2:DescribeInstances"  
"ec2:DescribeNetworkInterfaces"  
"ec2:AssignPrivateIpAddresses"
```

Note:

`unassignPrivateIpAddress` is not required.

This method is faster than the legacy method. In the older method, HA depends on the migration of AWS elastic network interfaces of the primary node to the secondary node.

For a legacy method, the following policies are required:

```
"iam:GetRole"  
"ec2:DescribeInstances"  
"ec2:DescribeAddresses"  
"ec2:AssociateAddress"  
"ec2:DisassociateAddress"
```

For more information, see [Deploy a high availability pair on AWS](#).

High availability across different zones

You can configure two Citrix ADC VPX instances on two different subnets or two different AWS availability zones, as a high availability active-passive pair in Independent Network Configuration (INC) mode. Upon failover, the EIP (Elastic IP) of the VIP of the primary instance migrates to the secondary, which takes over as the new primary. In the failover process, the AWS API:

- Checks the virtual servers that have [IPSets](#) attached to them.
- Finds the IP address that has an associated public IP, from the two IP addresses the virtual server is listening on. One that is directly attached to the virtual server, and one that is attached through the IP set.
- Reassociates the public IP (EIP) to the private IP belonging to the new primary VIP.

For HA across different zones, the following policies are required:

```
"iam:GetRole"  
"ec2:DescribeInstances"  
"ec2:DescribeAddresses"  
"ec2:AssociateAddress"  
"ec2:DisassociateAddress"
```

For more information, see [High availability across AWS availability zones](#).

Before you start your deployment

Before you start any HA deployment on AWS, read the following document:

- [Prerequisites](#)
- [Limitations and usage guidelines](#)
- [Deploy a Citrix ADC VPX instance on AWS](#)
- [High Availability](#)

Troubleshooting

To troubleshoot any failure during a HA failover of Citrix ADC VPX instance on AWS cloud, do the following:

- For release 13.0 build 65.3 and later, check the `cloud-ha-daemon.log` file stored in the `/var/log/` location.
- For releases earlier than 13.0 build 65.3, check the `ns.log` file stored in the `/var/log/` location.

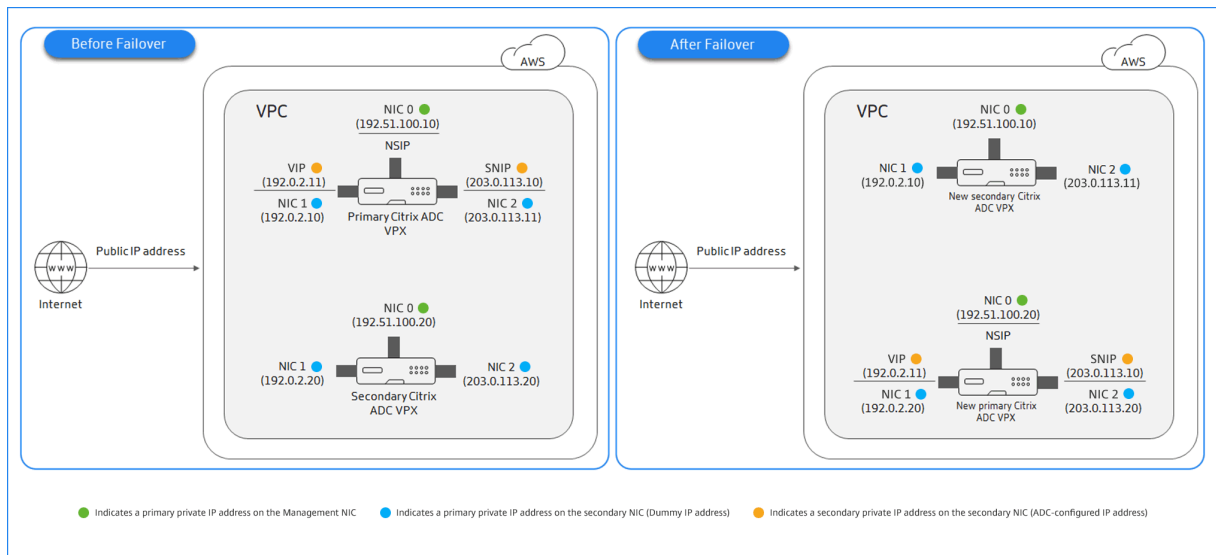
Deploy a VPX HA pair in the same AWS availability zone

September 9, 2024

You can configure two Citrix ADC VPX instances on AWS as a high-availability (HA) pair, in the same AWS zone where both VPX instances are on the same subnet. HA is achieved by migrating secondary private IP addresses attached to the NICs (client and server-side NICs) of the primary HA node to the secondary HA node after failover. All the Elastic IP addresses associated with the secondary private IP addresses are also migrated.

The following illustration depicts an HA failover scenario by migrating secondary private IP addresses.

Figure 1. A Citrix ADC VPX HA Pair on AWS, using private IP migration



Before you start your document, read the following docs:

- [Prerequisites](#)
- [Limitations and usage guidelines](#)
- [Deploy a Citrix ADC VPX instance on AWS](#)
- [High Availability](#)

How to deploy a VPX HA pair in the same zone

Here is the summary of the steps to deploy a VPX HA pair in the same zone:

1. Create two VPX instances on AWS, each with three NICs
2. Assign AWS secondary private IP address to VIP and SNIP of primary node

3. Configure VIP and SNIP on primary node using AWS secondary private IP addresses
4. Configure HA on both nodes

Step 1. Create two VPX instances (primary and secondary nodes) by using the same VPC, each with three NICs (Ethernet 0, Ethernet 1, Ethernet 2)

Follow the steps given in [Deploy a Citrix ADC VPX instance on AWS by using the AWS web console](#).

Step 2. On the primary node, assign secondary private IP addresses for Ethernet 1 (client IP or VIP) and Ethernet 2 (back-end server IP or SNIP)

The AWS console automatically assigns primary private IP addresses to the configured NICs. Assign more private IP addresses to VIP and SNIP, known as secondary private IP addresses.

To assign a secondary private IPv4 address to a network interface, follow these steps:

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose Network Interfaces, and then select the network interface attached to the instance.
3. Choose Actions, Manage IP Addresses.
4. Under IPv4 Addresses, choose Assign new IP.
5. Enter a specific IPv4 address that's within the subnet range for the instance, or leave the field blank to let Amazon select an IP address for you.
6. (Optional) Choose Allow reassignment to allow the secondary private IP address to be reassigned if it is already assigned to another network interface.
7. Choose Yes, Update.

Under the instance description, the assigned secondary private IP addresses appear.

Step 3. Configure VIP and SNIP on the primary node, using secondary private IP addresses

Access the primary node using SSH. Open an ssh client and type:

```
1 ssh -i <location of your private key> nsroot@<public DNS of the instance>
```

Next, configure VIP and SNIP.

For VIP, type:

```
1 add ns ip <IPAddress> <netmask> -type <type>
```

For SNIP, type:


```
1 add ns ip <IPAddress> <netmask> -type SNIP
```

Type `save config` to save.

To see the configured IP addresses, type the following command:

```
1 show ns ip
```

For more information, see the following topics:

- [Configuring and Managing Virtual IP \(VIP\) Addresses](#)
- [Configuring the NSIP address](#)

Step 4: Configure HA on both instances

On the primary node, open a Shell client and type the following command:

```
1 add ha node <id> <private IP address of the management NIC of the secondary node>
```

On the secondary node, type the following command:

```
1 add ha node <id> < private IP address of the management NIC of the primary node >
```

Type `save config` to save the configuration.

To see the configured HA nodes, type `show ha node`.

Upon failover, the secondary private IP addresses configured as VIP and SNIP on the previous primary node are migrated to the new primary node.

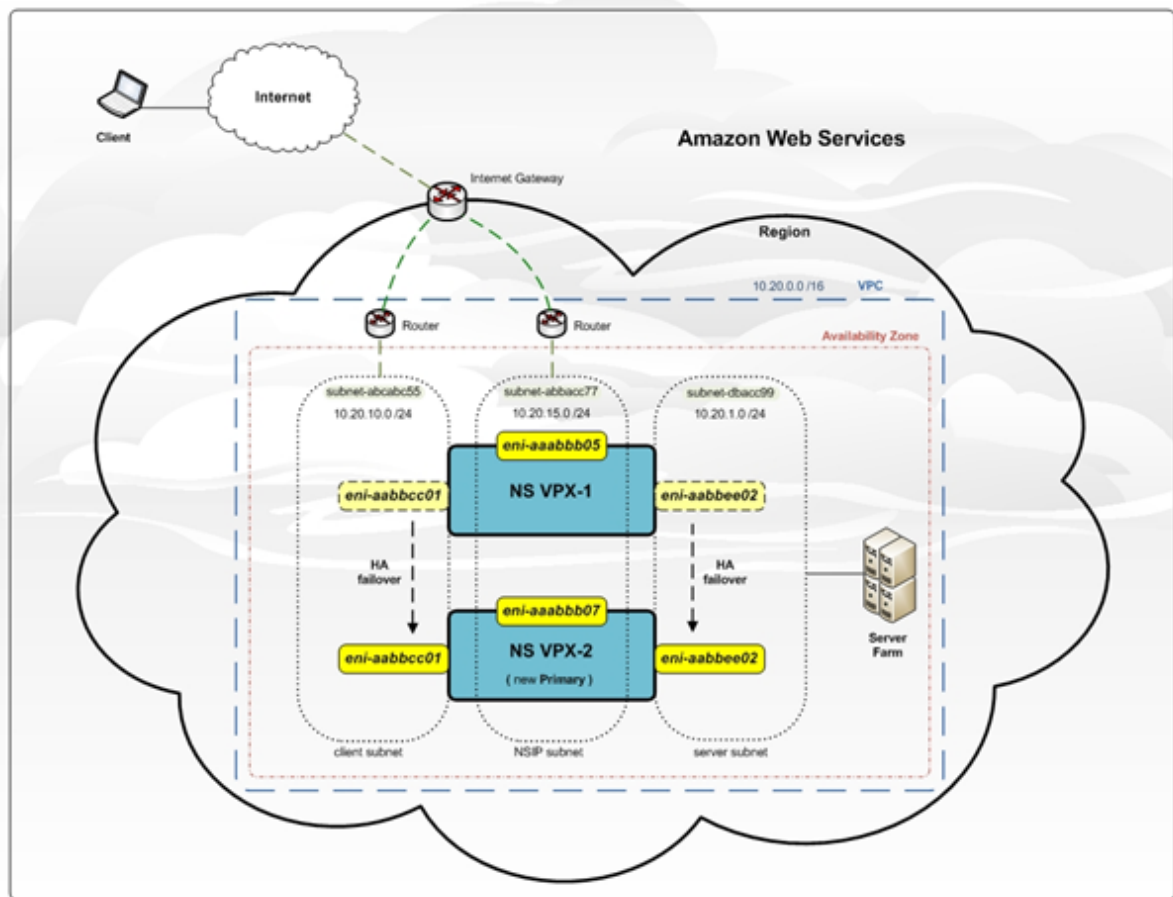
To force a failover on a node, type `force HAfailover`.

Legacy method for deploying a VPX HA pair

Before 13.0 41.x release, HA within the same zone was achieved through AWS elastic network interface (ENI) migration. However, this method is slowly deprecated.

The following figure shows an example of the HA deployment architecture for Citrix ADC VPX instances on AWS.

Figure 1. A Citrix ADC VPX HA Pair on AWS, using ENI migration



You can deploy two VPX instances on AWS as an HA pair by using one of the following options:

- Create the instances with IAM Role manually by using the AWS Management Console and then configure HA on them.
- Or automate the high availability deployment by using the Citrix CloudFormation template.

The CloudFormation template significantly decreases the number of steps involved for creating an HA pair, and it automatically creates an IAM Role. This section shows how to deploy a Citrix ADC VPX HA (active-passive) pair by using the Citrix CloudFormation template.

Keep the following points in mind while deploying two Citrix ADC VPX instances as an HA pair.

Points to note

- HA on AWS requires the primary node to have at least two ENIs (one for management and the other for data traffic), and the secondary node to have one management ENI. However, for security purposes, create three ENIs on the primary node, because this setup allows you to segregate the private and public network (recommended).

- The secondary node always has one ENI interface (for management) and the primary node can have up to four ENIs.
- The NSIP addresses for each VPX instance in a high availability pair must be configured on the default ENI of the instance.
- Amazon does not allow any broadcast/multicast packets in AWS. As a result, in a HA setup, data-plane ENIs are migrated from the primary to the secondary VPX instance when the primary VPX instance fails.
- Because the default (management) ENI cannot be moved to another VPX instance, do not use the default ENI for client and server traffic (data-plane traffic).
- The message `AWSCONFIG_IOCTL_NSAPI_HOTPLUG_INTF success output 0` in the `/var/log/ns.log` indicates that the two data ENIs have successfully attached to the secondary instance (the new primary).
- Failover might take up to 20 seconds due to the AWS detach/attach ENI mechanism.
- Upon failover, the failed instance always restarts.
- The heartbeat packets are received only on the management interface.
- The configuration file of the primary and secondary VPX instances is synchronized, including the `nsroot` password. The `nsroot` password of the secondary node is set to that of the primary node after the HA configuration synchronization.
- To have access to the AWS API servers, either the VPX instance must have a public IP address assigned or routing must be set up correctly at VPC subnet level pointing to the internet gateway of the VPC.
- Nameservers/DNS servers are configured at VPC level using DHCP options.
- The Citrix CloudFormation template does not create an HA setup between different availability zones.
- The Citrix CloudFormation template does not create an INC mode.
- The AWS debug messages are available in the log file, `/var/log/ns.log`, on the VPX instance.

Deploy a high availability pair by using the Citrix CloudFormation template

Before starting the CloudFormation template, ensure that you complete the following requirements:

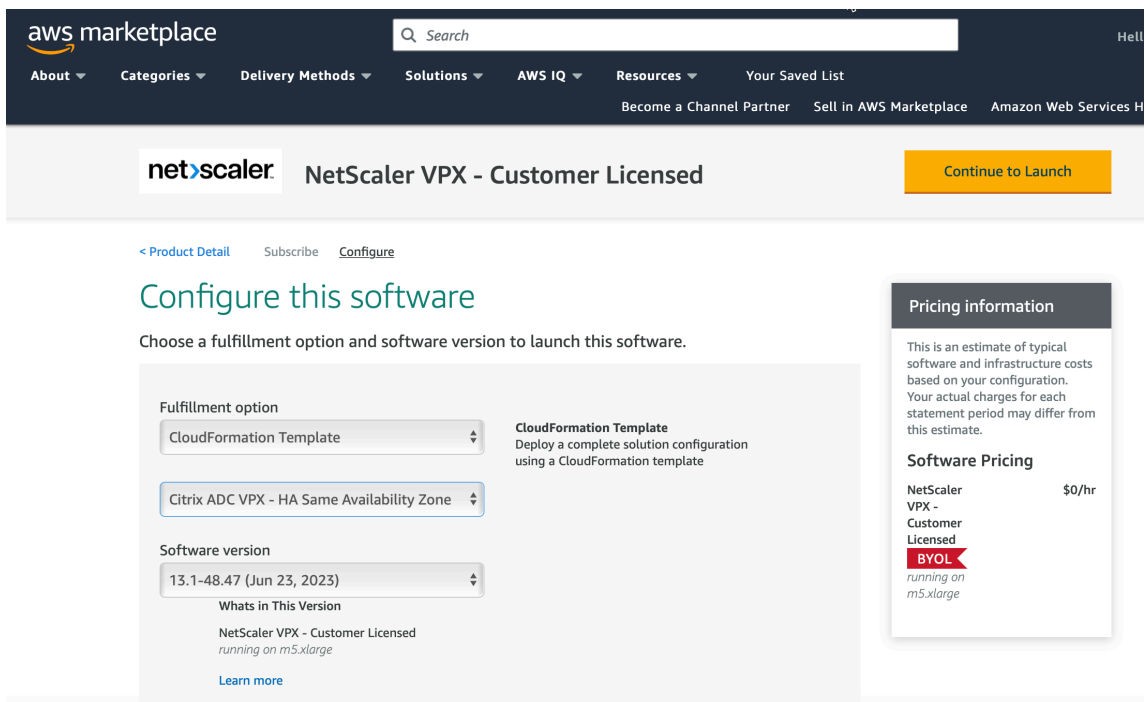
- A VPC
- Three subnets within the VPC
- A security group with UDP 3003, TCP 3009–3010, HTTP, SSH ports open
- A key pair
- Create an internet gateway
- Edit route tables for client and management networks to point to the internet gateway

Note:


The Citrix CloudFormation template automatically creates an IAM Role. Existing IAM Roles do not appear in the template.

To launch the Citrix CloudFormation template:

1. Log on to the [AWS marketplace](#) by using your AWS credentials.
2. In the search field, type **Citrix ADC VPX** to search for the Citrix ADC AMI, and click **Go**.
3. On the search result page, click the desired Citrix ADC VPX offering.
4. Click the **Pricing** tab, to go to **Pricing Information**.
5. Select the region and **Fulfillment Option** as **Citrix ADC VPX –Customer Licensed**.
6. Click **Continue to Subscribe**.
7. Check the details in the **Subscribe** page and click **Continue to Configuration**.
8. Select **Delivery Method** as **CloudFormation Template**.
9. Select the required CloudFormation template.
10. Select **Software Version** and **Region**, and click **Continue to Launch**.




11. Under **Choose Action**, select **Launch CloudFormation**, and click **Launch**. The **Create stack** page appears.
12. Click **Next**.




13. The **Specify stack details** page appears. Enter the following details.

- Type a **Stack name**. The name must be within 25 characters.
- Under **Network Configuration**, perform the following:
 - Select **Management Subnetwork**, **Client Subnetwork**, and **Server Subnetwork**. Ensure that you select the correct subnetworks you created within the VPC that you selected under VPC ID.
 - Add **Primary Management IP**, **Secondary Management IP**, **Client IP**, and **Server IP**. The IP addresses must belong to the same subnets of the respective subnetworks. Alternatively, you can let the template assign the IP addresses automatically.
 - Select **default** for **VPCTenancy**.
- Under **Citrix ADC Configuration**, perform the following:
 - Select **m5.xlarge** for **Instance type**.
 - Select the key pair that you've already created from the menu for **Key Pair**.
 - By default, the **Publish custom metrics to CloudWatch?** option is set to **Yes**. If you want to disable this option, select **No**.
For more information about CloudWatch metrics, see [Monitor your instances using Amazon CloudWatch] (#monitor-your-instances-using-amazon-cloudWatch).
- Under **Optional Configuration**, perform the following:
 - By default, the **Should publicIP(EIP) be assigned to management interfaces?** option is set to **No**.
 - By default, the **Should publicIP(EIP) be assigned to client interface?** option is set to **No**.




14. Click **Next**.


15. The **Configure stack options** page appears. This is an optional page.



16. Click **Next**.
17. The **Options** page appears. (This is an optional page.). Click **Next**.
18. The **Review** page appears. Take a moment to review the settings and make any changes, if necessary.
19. Select the **I acknowledge that AWS CloudFormation might create IAM resources.** check box, and then click **Create stack**.
20. The **CREATE-IN-PROGRESS** status appears. Wait until the status is **CREATE-COMplete**. If the status does not change to **COMPLETE**, check the **Events** tab for the reason of failure, and recreate the instance with proper configurations.



- After an IAM resource is created, navigate to **EC2 Management Console > Instances**. You find two VPX instances created with IAM role. The primary and secondary nodes are created each with three private IP addresses and three network interfaces.
- Log on to the primary node with user name `nsroot` and the instance ID as the password. From the GUI, navigate to **System > High Availability > Nodes**. The Citrix ADC VPX is already configured in HA pair by the CloudFormation template.
- The Citrix ADC VPX HA pair appears.



Monitor your instances using Amazon CloudWatch

You can use the Amazon CloudWatch service to monitor a set of Citrix ADC VPX metrics such as CPU and memory utilization, and throughput. CloudWatch monitors resources and applications that run on AWS, in real time. You can access the Amazon CloudWatch dashboard by using the AWS Management console. For more information, see [Amazon CloudWatch](#).

Points to note


- If you deploy a Citrix ADC VPX instance on AWS by using the AWS web console, the CloudWatch service is enabled by default.

- If you deploy a Citrix ADC VPX instance by using the Citrix CloudFormation template, the default option is “Yes.” If you want to disable the CloudWatch service, select “No.”
- Metrics are available for CPU (management and packet CPU usage), memory, and throughput (inbound and outbound).


How to view CloudWatch metrics

To view CloudWatch metrics for your instance, follow these steps:

1. Log on to **AWS Management console > EC2 > Instances**.
2. Select the instance.
3. Click **Monitoring**.
4. Click **View all CloudWatch metrics**.

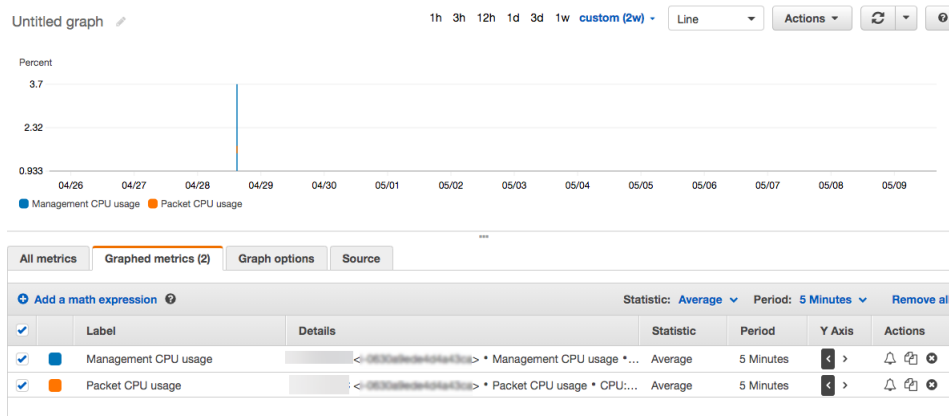


5. Under All metrics, click your instance ID.



6. Click the metrics you want to view, set the duration (by minutes, hours, days, weeks, months).
7. Click **Graphed metrics** to view the statistics of usage. Use the **Graph options** to customize your graph.

Figure. Graphed metrics for CPU usage



Configuring SR-IOV on a high availability setup

Support for SR-IOV interfaces in a high availability setup is available from Citrix ADC release 12.0 57.19 onwards. For more information about how to configure SR-IOV, see [Configuring Citrix ADC VPX instances to Use SR-IOV Network Interface](#).

Related resources

[How high availability on AWS works](#)

High availability across different AWS availability zones

September 9, 2024

You can configure two Citrix ADC VPX instances on two different subnets or two different AWS availability zones, as a high availability active-passive pair in Independent Network Configuration (INC) mode. If for any reason, the primary node is unable to accept connections, the secondary node takes over.

For more information about high availability, see [High availability](#). For more information about INC, see [Configuring high availability nodes in different subnets](#).

Points to note

- Read the following documents before you start your deployment:
 - [AWS terminology](#)
 - [Prerequisites](#)
 - [Limitations and usage guidelines](#)

- The VPX high availability pair can either reside in the same availability zone in a different subnet or in two different AWS availability zones.
- Citrix recommends that you use different subnets for management (NSIP), client traffic (VIP), and back-end server (SNIP).
- High availability must be set in Independent Network Configuration (INC) mode for a failover to work.
- The two instances must have port 3003 open for UDP traffic as that is used for heartbeats.
- The management subnets of both the nodes must have access to internet or to AWS API server through internal NAT so that the rest APIs are functional.
- IAM role must have E2 permission for the public IP or elastic IP (EIP) migration and EC2 Route Table permissions for the private IP migration.

You can deploy high availability across AWS availability zones in the following ways:

- [Using elastic IP addresses](#)
- [Using private IP addresses](#)

Deploy a VPX high-availability pair with elastic IP addresses across different AWS zones

September 9, 2024

You can configure two Citrix ADC VPX instances on two different subnets or two different AWS availability zones using elastic IP addresses in the INC mode.

For more information about high availability, see [High availability](#). For more information about INC, see [Configuring high availability nodes in different subnets](#).

How HA with EIP addresses across different AWS zones work

Upon failover, the EIP of the VIP of the primary instance migrates to the secondary, which takes over as the new primary. In the failover process, AWS API:

1. Checks the virtual servers that have [IPSets](#) attached to them.
2. Finds the IP address that has an associated public IP, from the two IP addresses the virtual server is listening on. One that is directly attached to the virtual server, and one that is attached through the IP set.
3. Reassociates the public IP (EIP) to the private IP belonging to the new primary VIP.

Note:

To protect your network from attacks such as denial-of-service (DoS), when using an EIP, you can create security groups in AWS to restrict the IP access. For high availability, you can switch from EIP to a private IP movement solution as per your deployments.

How to deploy a VPX high-availability pair with elastic IP addresses across different AWS zones

The following is the summary of steps for deploying a VPX pair on two different subnets or two different AWS availability zones.

1. Create an Amazon virtual private cloud.
2. Deploy two VPX instances in two different availability zones or in the same zone but in different subnets.
3. Configure high availability
 - a) Set up high availability in INC mode in both the instances.
 - b) Add an [IP set](#) in both the instances.
 - c) Bind the IP set in both the instances to the VIP.
 - d) Add a virtual server in the primary instance.

For steps 1 and 2, use the AWS console. For steps 3, use the Citrix ADC VPX GUI or the CLI.

Step 1. Create an Amazon virtual private cloud (VPC).

Step 2. Deploy two VPX instance in two different availability zones or in the same zone but in different subnets. Attach an EIP to the VIP of the primary VPX.

For more information about how to create a VPC and deploy a VPX instance on AWS, see [Deploy a Citrix ADC VPX standalone instance on AWS](#) and [Scenario: standalone instance](#)

Step 3. Configure high availability. You can use the Citrix ADC VPX CLI or the GUI to set up high availability.

Configure high availability by using the CLI

1. Set up high availability in INC mode in both the instances.

On the primary node:

```
add ha node 1 <sec_ip> -inc ENABLED
```

On the secondary node:

```
add ha node 1 <prim_ip> -inc ENABLED
```

<sec_ip> refers to the private IP address of the management NIC of the secondary node

<prim_ip> refers to the private IP address of the management NIC of the primary node

2. Add the IP set in both the instances.

Type the following command on both the instances.

```
add ipset <ipsetname>
```

3. Bind the IP set to the VIP set on both the instances.

Type the following command on both the instances:

```
add ns ip <secondary vip> <subnet> -type VIP
```

```
bind ipset <ipsetname> <secondary VIP>
```

Note:

You can bind the IP set to the primary VIP or to the secondary VIP. However, if you bind the IP set to the primary VIP, use the secondary VIP to add to the virtual server, and conversely.

4. Add a virtual server on the primary instance.

Type the following command:

```
add <server_type> vserver <vserver_name> <protocol> <primary_vip>  
<port> -ipset \<ipset_name>
```

Configure high availability by using the GUI

1. Set up high availability in INC mode on both the instances
2. Log on to the primary node with user name `nsroot` and instance ID as password.
3. From the GUI, go to **Configuration > System > High Availability**. Click **Add**.
4. At the **Remote Node IP address** field, add the private IP address of the management NIC of the secondary node.
5. Select **Turn on NIC (Independent Network Configuration)** mode on self-node.
6. Under **Remote System Login Credential**, add the user name and password for the secondary node and click **Create**.
7. Repeat the steps in the secondary node.
8. Add IP set and bind IP set to the VIP set on both the instances.
9. From the GUI, navigate to **System > Network > IPs > Add**.
10. Add the required values for IP Address, Netmask, IP Type (virtual IP) and click **Create**.

11. Navigate to **System > Network > IP Sets > Add**. Add an IP set name and click **Insert**.
12. From the IPv4s page, select the virtual IP and click **Insert**. Click **Create** to create the IP set.
13. Add a virtual server in the primary instance

From the GUI, go to **Configuration > Traffic Management > Virtual Servers > Add**.


Load Balancing Virtual Server | [Export as a Template](#)

Basic Settings			
Name	vserver1	Listen Priority	-
Protocol	HTTP	Listen Policy Expression	NONE
State	● DOWN	Redirection Mode	IP
IP Address	192.168.2.129	Range	1
Port	80	IPset	ipset123
Traffic Domain	0	RHI State	PASSIVE
		AppFlow Logging	ENABLED
		Retain Connections on Cluster	NO

Scenario

In this scenario, a single VPC is created. In that VPC, two VPX instances are created in two availability zones. Each instance has three subnets - one for management, one for client, and one for back-end server. An EIP is attached to the VIP of the primary node.

Diagram: This diagram illustrates the Citrix ADC VPX high availability setup in INC mode, on AWS



For this scenario, use CLI to configure high availability.

1. Set up high availability in INC mode on both the instances.

Type the following commands on the primary and the secondary nodes.

On primary:

```
add ha node 1 192.168.6.82 -inc enabled
```

Here, 192.168.6.82 refers to the private IP address of the management NIC of the secondary node.

On secondary:

```
add ha node 1 192.168.1.108 -inc enabled
```

Here, 192.168.1.108 refers to the private IP address of the management NIC of the primary node.

2. Add an IP set and bind the IP set to the VIP on both the instances

On primary:

```
add ipset ipset123
add ns ip 192.168.7.68 255.255.255.0 -type VIP
bindipset ipset123 192.168.7.68
```

On secondary:

```
add ipset ipset123
add ns ip 192.168.7.68 255.255.255.0 -type VIP
bind ipset ipset123 192.168.7.68
```

3. Add a virtual server on the primary instance.

The following command:

```
add lbvserver vserver1 http 192.168.2.129 80 -ipset ipset123
```

4. Save the configuration.

<input type="checkbox"/>	ID	IP Address	Host Name	Master State	Node State	INC	Synchronization State
<input type="checkbox"/>	0	192.168.1.108		Primary	● UP	ENABLED	ENABLED
<input type="checkbox"/>	1	192.168.6.82		Secondary	● UP	ENABLED	SUCCESS

5. After a forced failover, the secondary becomes the new primary.

<input type="checkbox"/>	ID	IP Address	Host Name	Master State	Node State	INC	Synchronization State
<input type="checkbox"/>	0	192.168.1.108		Secondary	● UP	ENABLED	SUCCESS
<input type="checkbox"/>	1	192.168.6.82		Primary	● UP	ENABLED	ENABLED

Deploy a VPX high-availability pair with private IP addresses across different AWS zones

September 9, 2024

You can configure two Citrix ADC VPX instances on two different subnets or two different AWS availability zones using private IP addresses in the INC mode. This solution can be easily integrated with the existing multizone [VPX high-availability pair with elastic IP addresses](#). Therefore, you can use both the solutions together.

For more information about high availability, see [High availability](#). For more information about INC, see [Configuring high availability nodes in different subnets](#).

Note:

This deployment is supported from Citrix ADC release 13.0 build 67.39 onwards. This deployment is compatible with AWS Transit Gateway and VPC peering.

Prerequisites

Ensure that the IAM role associated with your AWS account has the following IAM permissions:

```
1 {
2
3     "Version": "2012-10-17",
4     "Statement": [
5         {
6
7             "Action": [
8                 "ec2:DescribeInstances",
9                 "ec2:DescribeAddresses",
10                "ec2:AssociateAddress",
11                "ec2:DisassociateAddress",
12                "ec2:DescribeRouteTables",
13                "ec2>DeleteRoute",
14                "ec2:CreateRoute",
15                "ec2:ModifyNetworkInterfaceAttribute",
16                "iam:SimulatePrincipalPolicy",
17                "iam:GetRole"
18            ],
19            "Resource": "*",
20            "Effect": "Allow"
21        }
22    ]
23 }
24 }
```


How to deploy a VPX high-availability pair with private IP addresses across different AWS zones

The following is the summary of steps for deploying a VPX pair on two different subnets or two different AWS availability zones using private IP addresses.

1. Create an Amazon virtual private cloud.
2. Deploy two VPX instances in two different availability zones.
3. Configure high availability
 - a) Set up high availability in INC mode in both the instances.
 - b) Add the respective route tables in the VPC that points to the client interface.
 - c) Add a virtual server in the primary instance.

For steps 1 and 2, use the AWS console. For step 3, use the Citrix ADC VPX GUI or the CLI.

Step 1. Create an Amazon virtual private cloud (VPC).

Step 2. Deploy two VPX instance in two different availability zones with the same number of ENI (Network Interface).

For more information about how to create a VPC and deploy a VPX instance on AWS, see [Deploy a Citrix ADC VPX standalone instance on AWS](#) and [Scenario: standalone instance](#)

Step 3. Configure the ADC VIP addresses by choosing a subnet that does not overlap with the Amazon VPC subnets. If your VPC is 192.168.0.0/16, then to configure ADC VIP addresses, you can choose any subnet from these IP address ranges:

- 0.0.0.0 - 192.167.0.0
- 192.169.0.0 - 254.255.255.0

In this example, the chosen 10.10.10.0/24 subnet and created VIPs in this subnet. You can choose any subnet other than the VPC subnet (192.168.0.0/16).

Step 4. Add a route that points to the client interface (VIP) of the primary node from the VPC route table.

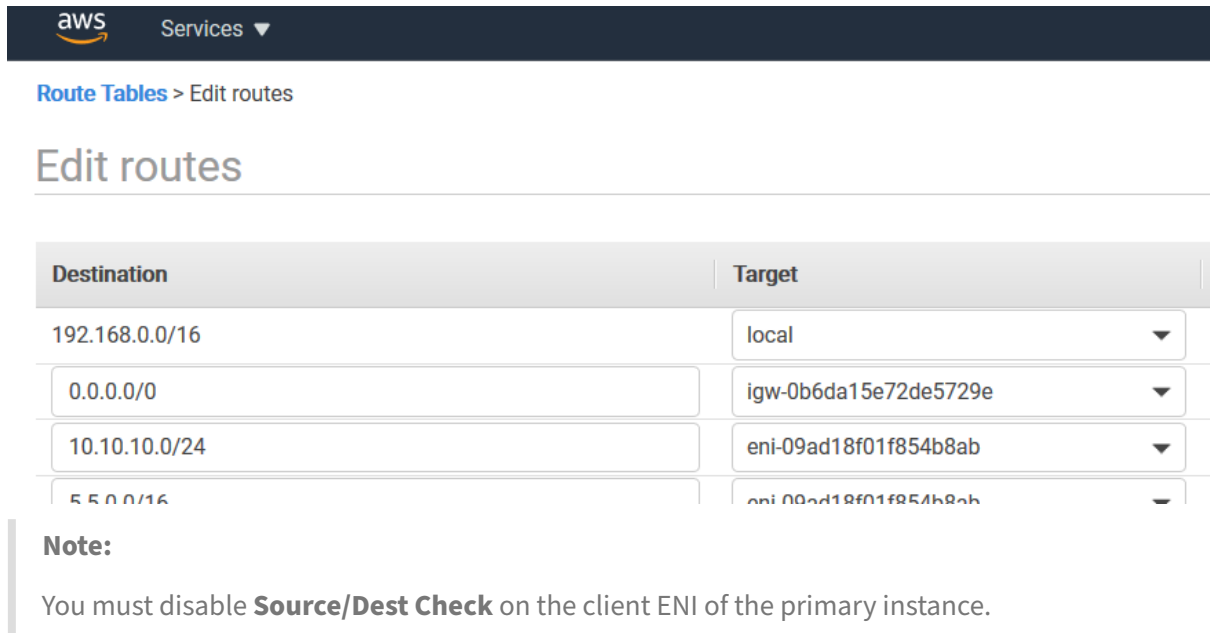
From the AWS CLI, type the following command:

```
1 aws ec2 create-route --route-table-id rtb-2272532 --destination-cidr-block 10.10.10.0/24 --gateway-id <eni-client-primary>
```

From the AWS GUI, perform the following steps to add a route:

1. Open the [Amazon EC2 console](#).
2. In the navigation pane, choose **Route Tables**, and select the route table.
3. Choose **Actions**, and click **Edit routes**.

- To add a route, choose **Add route**. For **Destination**, enter the destination CIDR block, a single IP address, or the ID of a prefix list. For gateway ID, select the ENI of a client interface of the primary node.



aws Services ▾

Route Tables > Edit routes

Edit routes

Destination	Target
192.168.0.0/16	local
0.0.0.0/0	igw-0b6da15e72de5729e
10.10.10.0/24	eni-09ad18f01f854b8ab
5.5.0.0/16	eni-09ad18f01f854b8ab

Note:
 You must disable **Source/Dest Check** on the client ENI of the primary instance.

To disable the source/destination checking for a network interface using the console, perform the following steps:

- Open the [Amazon EC2 console](#).
- In the navigation pane, choose **Network Interfaces**.
- Select the network interface of a primary client interface, and choose **Actions**, and click **Change Source/Dest. Check**.
- In the dialog box, choose **Disabled**, click **Save**.

Change Source/Dest. Check

×

Network Interface eni-0047841c06c3e9012

Source/dest. check Enabled
 Disabled

Cancel
Save

Step 5. Configure high availability. You can use the Citrix ADC VPX CLI or the GUI to set up high availability.

Configure high availability by using the CLI

1. Set up high availability in INC mode in both the instances.

On the primary node:

```
1 add ha node 1 \<sec\_ip\> -inc ENABLED
```

On the secondary node:

```
1 add ha node 1 \<prim\_ip\> -inc ENABLED
```

<sec_ip> refers to the private IP address of the management NIC of the secondary node.

<prim_ip> refers to the private IP address of the management NIC of the primary node.

2. Add a virtual server on the primary instance. You must add it from the chosen subnet, for example, 10.10.10.0/24.

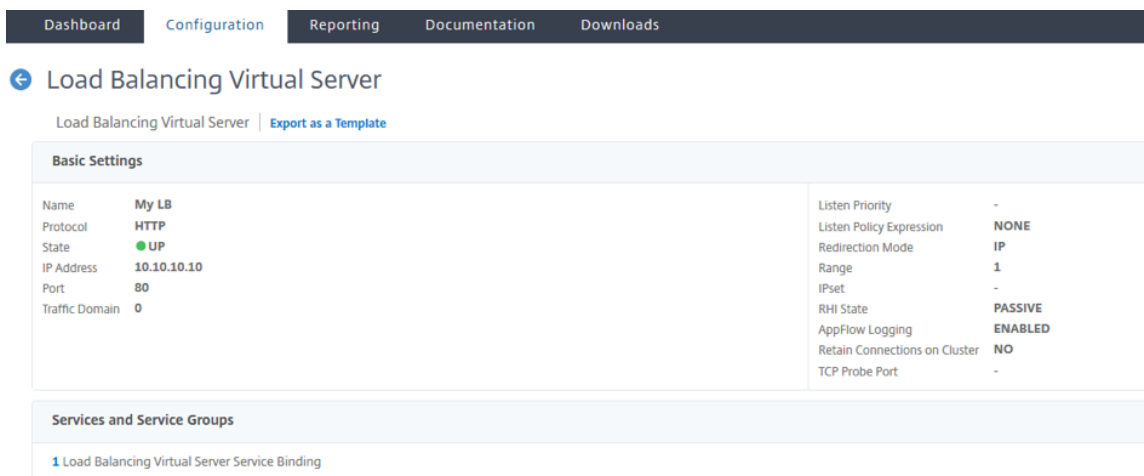
Type the following command:

```
1 add \<server\_type\> vserver \<vserver\_name\> \<protocol\> \<primary\_vip\> \<port\>
```

Configure high availability by using the GUI

1. Set up high availability in INC mode on both the instances
2. Log on to the primary node with user name `nsroot` and instance ID as password.
3. Navigate to **Configuration > System > High Availability**, and click **Add**.
4. At the **Remote Node IP address** field, add the private IP address of the management NIC of the secondary node.
5. Select **Turn on NIC (Independent Network Configuration)** mode on self-node.
6. Under **Remote System Login Credential**, add the user name and password for the secondary node and click **Create**.
7. Repeat the steps in the secondary node.
8. Add a virtual server in the primary instance


Navigate to **Configuration > Traffic Management > Virtual Servers > Add**.




Scenario

In this scenario, a single VPC is created. In that VPC, two VPX instances are created in two availability zones. Each instance has three subnets - one for management, one for client, and one for back-end server.

The following diagrams illustrate the Citrix ADC VPX high availability setup in INC mode, on AWS. The custom subnet 10.10.10.10, which is not part of the VPC is used as VIP. Therefore, the 10.10.10.10 subnet can be used across availability zones.





For this scenario, use CLI to configure high availability.

1. Set up high availability in INC mode on both the instances.

Type the following commands on the primary and the secondary nodes.

On the primary node:

```
1 add ha node 1 192.168.4.10 -inc enabled
```

Here, 192.168.4.10 refers to the private IP address of the management NIC of the secondary node.

On the secondary node:

```
1 add ha node 1 192.168.1.10 -inc enabled
```

Here, 192.168.1.10 refers to the private IP address of the management NIC of the primary node.

2. Add a virtual server on the primary instance.

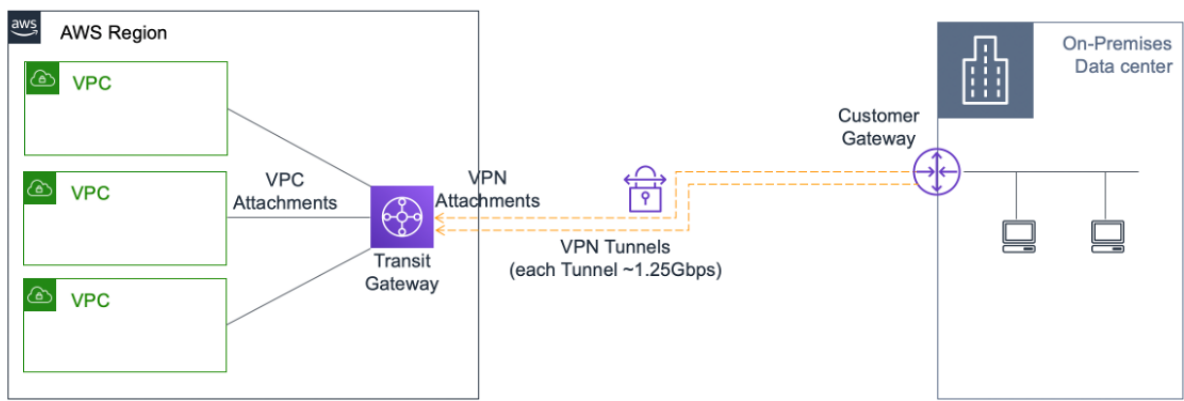
Type the following command:

```
1 add lbvserver vserver1 http 10.10.10.10 80
```

3. Save the configuration.
4. After a forced failover:
 - The secondary instance becomes the new primary instance.
 - The VPC route pointing to the primary ENI migrates to the secondary client ENI.
 - Client traffic resumes to the new primary instance.


AWS Transit Gateway configuration for HA private IP solution

You need AWS Transit Gateway to make the private VIP subnet routable within the internal network, across AWS VPCs, regions, and On-premises networks. The VPC must connect to AWS Transit Gateway. A static route for the VIP subnet or IP pool inside the AWS Transit Gateway route table is created and pointed towards the VPC.



To configure AWS Transit Gateway, follow these steps:

1. Open the [Amazon VPC console](#).
2. On the navigation pane, choose **Transit Gateway Route Tables**.
3. Choose the **Routes** tab, and click **Create static route**.



4. Create a static route where CIDR points to your private VIPS subnet and attachment points to the VPC having ADC VPX.

Transit Gateway Route Tables > Create static route

Create static route

Add a static route to your Transit Gateway route table.

Transit Gateway ID `lgw-0b3e99191e03c16ed`

Transit Gateway route table ID `lgw-rtb-09f12ca61473654a7`

CIDR* ⓘ

Blackhole ⓘ

Choose attachment ↕

* Required

Cancel **Create static route**

5. Click **Create static route**, then choose **Close**.

Troubleshooting

If you face any issues while configuring HA private IP solution across multizone HA, check the following key points for troubleshooting:

- Both primary and secondary nodes have the same set of IAM permissions.
- INC mode is enabled on both the primary and secondary nodes.
- Both primary and secondary nodes have the same number of interfaces.
- While creating an instance, follow the same order of attaching interfaces on both the nodes. On a primary node, if the client interface is attached first and the server interface is attached second. Then, follow the same order on the secondary node as well. If there is any mismatch, detach and reattach the interfaces in the correct order.
- If traffic does not flow, make sure the “Source/dest. Check” is disabled on the client interface of the primary node for the first time.
- Make sure the cloudhadaemon command (`ps -aux | grep cloudha`) is running in Shell.
- Make sure that the Citrix ADC firmware version is 13.0 build 70.x or later.
- For issues with the failover process, check the log file available at: `/var/log/cloud-ha-daemon.log`

Deploy a Citrix ADC VPX instance on AWS Outposts

September 9, 2024

AWS Outposts is a pool of AWS compute and storage capacity deployed at your site. Outposts provides AWS infrastructure and services in your on-premises location. AWS operates, monitors, and manages

this capacity as part of an AWS Region. You can use the same Citrix ADC VPX instances, AWS APIs, tools, and infrastructure across on-premises and the AWS cloud for a consistent hybrid experience.

You can create subnets on your Outposts and specify them when you create AWS resources such as EC2 instances, EBS volumes, ECS clusters, and RDS instances. Instances in the Outposts subnets communicate with other instances in the AWS Region using private IP addresses, all within the same Amazon Virtual Private Cloud (VPC).

For more information, see the [AWS Outposts user guide](#).

How AWS Outposts works

AWS Outposts is designed to operate with a constant and consistent connection between your Outposts and an AWS Region. To achieve this connection to the Region, and to the local workloads in your on-premises environment, you must connect your Outpost to your on-premises network. Your on-premises network must provide WAN access back to the Region and to the internet. The internet must also provide LAN or WAN access to the local network where your on-premises workloads or applications reside.

Prerequisite


- You must install an AWS Outposts at your site.
- The AWS Outposts' compute and storage capacity must be available for use.

For more information on how to place an order for AWS Outposts, see the following AWS documentation:

<https://aws.amazon.com/blogs/aws/aws-outposts-now-available-order-your-racks-today/>

Deploy a Citrix ADC VPX instance on AWS Outposts by using the AWS web console

The following figure depicts a simple deployment of Citrix ADC VPX instances on the Outposts. The Citrix ADC AMI present in the AWS Marketplace is also deployed in the Outposts.



Log in to the AWS web console and complete the following steps to deploy ADC VPX EC2 instances on your AWS Outposts.

1. Create a key pair.
2. Create a Virtual Private Cloud (VPC).
3. Add more subnets.
4. Create security groups and security rules.
5. Add route tables.
6. Create an internet gateway.
7. Create an ADC VPX instance by using the AWS EC2 service.
From the AWS dashboard, navigate to **Compute > EC2 > Launch Instance > AWS Marketplace**.
8. Create and attach more network interfaces.
9. Attach elastic IPs to the management NIC.
10. Connect to the VPX instance.

For detailed instructions on each of the steps, see [Deploy a Citrix ADC VPX instance on AWS by using the AWS web console](#).

For high availability within same availability zone deployment, see [Deploy a high availability pair on AWS](#).

Add back-end AWS Autoscaling service

September 9, 2024

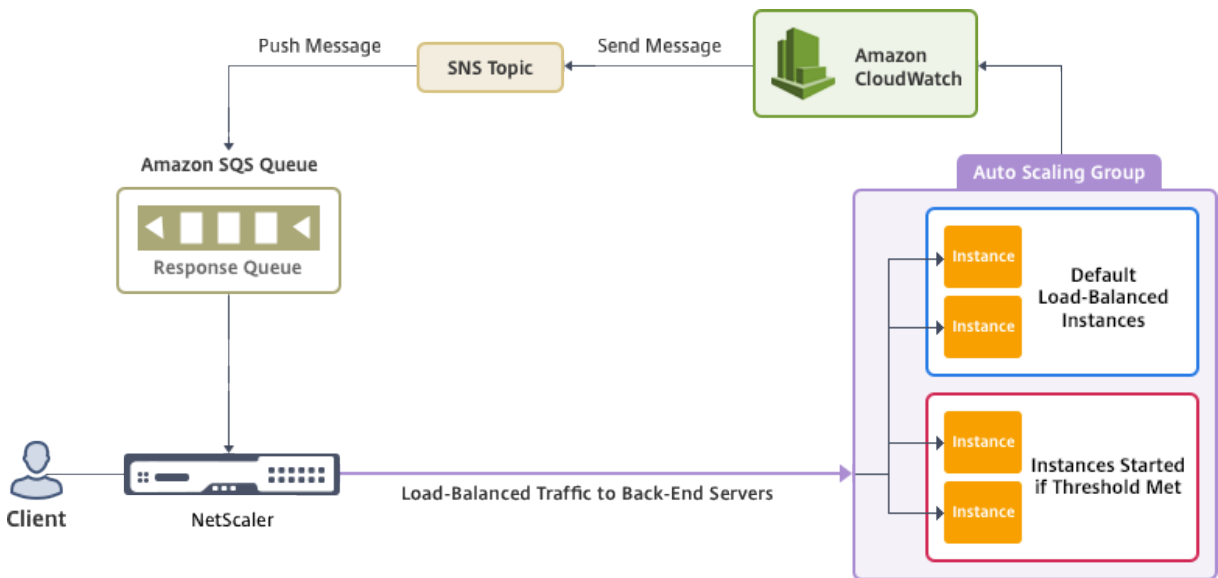
Efficient hosting of applications in a cloud involves easy and cost-effective management of resources depending on the application demand. To meet increasing demand, you have to scale network resources upward. Whether demand subsides, you need to scale down to avoid the unnecessary cost of

idle resources. To minimize the cost of running the application by deploying only as many instances as are necessary during any given time, you constantly have to monitor traffic, memory and CPU use, and so on. However, monitoring traffic manually is cumbersome. For the application environment to scale up or down dynamically, you must automate the processes of monitoring traffic and of scaling resources up and down whenever necessary.

Integrated with the AWS Auto Scaling service, the Citrix ADC VPX instance provides the following advantages:

- **Load balance and management:** Auto configures servers to scale up and scale down, depending on demand. The VPX instance auto detects Autoscale groups in the back-end subnet and allows a user to select the Autoscale groups to balance the load. All of this is done by auto configuring the virtual and subnet IP addresses on the VPX instance.
- **High availability:** Detects Autoscale groups that span multiple availability zones and load-balance servers.
- **Better network availability:** The VPX instance supports:
 - Back-end servers on different VPCs, by using VPC peering
 - Back-end servers on same placement groups
 - Back-end servers on different availability zones
- **Graceful connection termination:** Removes Autoscale servers gracefully, avoiding loss of client connections when scale-down activity occurs, by using the Graceful Timeout feature.

Diagram: AWS Autoscaling service with a Citrix ADC VPX Instance



This diagram illustrates how the AWS Autoscaling service is compatible with a Citrix ADC VPX instance (Load balancing virtual server). For more information, see the following AWS topics.

- [Autoscaling groups](#)

- [CloudWatch](#)
- [Simple Notification Service \(SNS\)](#)
- [Simple Queue Service \(Amazon SQS\)](#)

Before you begin

Before you start using Autoscaling with your Citrix ADC VPX instance, you must complete the following tasks.

1. Read the following topics:
 - [Prerequisites](#)
 - [Limitation and usage guidelines](#)
2. Create a Citrix ADC VPX instance on AWS according to your requirement.
 - For more information about how to create a Citrix ADC VPX standalone instance, see [Deploy a Citrix ADC VPX standalone instance on AWS](#) and [Scenario: standalone instance](#)
 - For more information about how to deploy VPX instances in HA mode, see [Deploy a high availability pair on AWS](#).

Note:

Citrix recommends the CloudFormation template for creating Citrix ADC VPX instances on AWS.

Citrix recommends you create three interfaces: one for management (NSIP), one for client-facing LB virtual server (VIP), and one for subnet IP (NSIP).

3. Create an AWS Autoscale group. If you don't have an existing Autoscaling configuration, you must:
 - a) Create a Launch Configuration
 - b) Create an Autoscaling Group
 - c) Verify the Autoscaling GroupFor more information, see <http://docs.aws.amazon.com/autoscaling/latest/userguide/GettingStartedTutorial.html>.
4. In the AWS Autoscale group, you must specify at least one scale-down policy. The Citrix ADC VPX instance supports only the Step scaling policy. The Simple scaling policy and Target tracking scaling policy are not supported for Autoscale group.

Add the AWS Autoscaling service to a Citrix ADC VPX instance

You can add the Autoscaling service to a VPX instance with a single click by using the GUI. Complete these steps to add the Autoscaling service to the VPX instance:

1. Log on to the VPX instance by using your credentials for `nsroot`.
2. When you log on to the Citrix ADC VPX instance for the first time, you see the default Cloud Profile page. Select the AWS Autoscaling group from the drop-down menu and click **Create** to create a cloud profile. Click **Skip** if you want to create the cloud profile later.

Points to keep in mind while creating a Cloud Profile: By default the CloudFormation Template creates and attaches the below IAM Role.

```
1  {
2
3
4  "Version": "2012-10-17",
5
6  "Statement": \[
7
8    {
9
10
11      "Action": \[
12
13        "ec2:DescribeInstances",
14
15        "ec2:DescribeNetworkInterfaces",
16
17        "ec2:DetachNetworkInterface",
18
19        "ec2:AttachNetworkInterface",
20
21        "ec2:StartInstances",
22
23        "ec2:StopInstances",
24
25        "ec2:RebootInstances",
26
27        "autoscaling:\*",
28
29        "sns:\*",
30
31        "sqs:\*"
32
33        " iam: SimulatePrincipalPolicy "
34
35        " iam: GetRole "
36
37      \],
38
```

```
39         "Resource": "\*",
40
41         "Effect": "Allow"
42     }
43 }
44
45 \]
46
47 }
48 }
```

Ensure the IAM Role of an instance has proper permissions.

- The virtual server IP address is autopopulated from the free IP address available to the VPX instance. <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/MultipleIP.html#ManageMultipleIP>
- Autoscale group is prepopulated from the Autoscale group configured on your AWS account. <http://docs.aws.amazon.com/autoscaling/latest/userguide/AutoScalingGroup.html>.
- While selecting the Autoscaling Group protocol and port, ensure your servers listen on those protocol and ports, and you bind the correct monitor in the service group. By default, the TCP monitor is used.
- For SSL Protocol type Autoscaling, after you create the Cloud Profile the load balance virtual server or service group is down because of a missing certificate. You can bind the certificate to the virtual server or service group manually.
- Select the Graceful Timeout option to remove Autoscale servers gracefully. If this option is not selected the server is the Autoscale group is removed immediately after the load goes down, which might cause service interruption for the existing connected clients. Selecting Graceful and giving a timeout means in the event of scale down. The VPX instance does not remove the server immediately but marks one of the servers for graceful deletion. During this period, the instance does not allow new connections to this server. Existing connections are served until the timeout occurs, and after a timeout, the VPX instance removes the server.

Figure: Default Cloud Profile page

Name
CloudProfile

Virtual Server IP Address*

Load Balancing Server Protocol*
HTTP

Load Balancing Server Port*
80

Auto Scale Group*
SharePoint

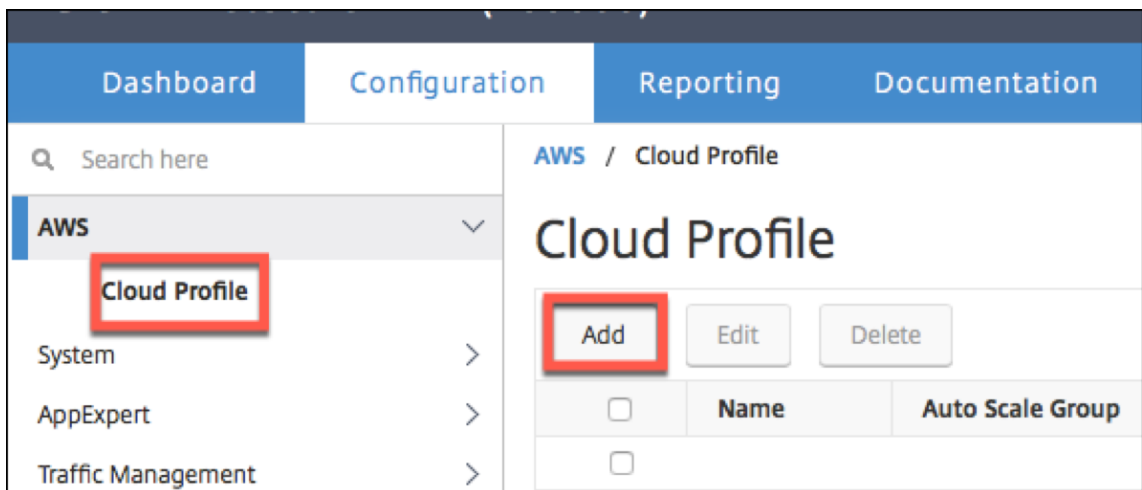
Auto Scale Group Protocol
HTTP

Auto Scale Group Port*
80

Select this option to drain the connections gracefully. Else the connections will be dropped
 Graceful

Create Skip

3. After the first time login if you want to create Cloud Profile, on the GUI go to **System > AWS > Cloud Profile** and click **Add**.



The **Create Cloud Profile** configuration page appears.

Citrix NetScaler VPX (3000)

Dashboard Configuration Reporting Documentation Downloads

← Create Cloud Profile

Name

Virtual Server IP Address*

Load Balancing Server Protocol

Load Balancing Server Port

Auto Scale Group*

Auto Scale Group Protocol

Auto Scale Group Port

Select this option to drain the connections gracefully. Else the connections will be dropped in the event of scale down.

Graceful

Delay (Seconds)

Create **Close**

Cloud Profile creates a Citrix ADC load-balancing virtual server and a service group with members as the servers of the Autoscaling group. Your back-end servers must be reachable through the SNIP configured on the VPX instance.

Citrix NetScaler VPX (3000)

Dashboard Configuration Reporting Documentation Downloads

HA Status Not configured Partition default nsroot

Search here

AWS / Cloud Profile

Cloud Profile

Name	Auto Scale Group	Load Balancing Virtual Server	Auto Scale Group Protocol	Graceful	Delay (Seconds)
<input type="checkbox"/> SharePoint_CloudProfile	SharePoint	_CP_SharePoint_CloudProfile_21.0.2.29_1B_	HTTP	YES	60

Note:

To view Autoscale-related information in the AWS console, go to **EC2 > Dashboard > Auto Scaling > Auto Scaling Group**.

Configure a Citrix ADC VPX instance to use SR-IOV network interface

September 9, 2024

Note:

Support for SR-IOV interfaces in a high availability setup is available from Citrix ADC release 12.0 57.19 onwards.

After you have created a Citrix ADC VPX instance on AWS, you can configure the virtual appliance to use SR-IOV network interfaces, by using the AWS CLI.

In all Citrix ADC VPX models, except Citrix ADC VPX AWS Marketplace Editions of 3G and 5G, SR-IOV is not enabled in the default configuration of a network interface.

Before you start the configuration, read the following topics:

- [Prerequisites](#)
- [Limitations and Usage Guidelines](#)

This section includes the following topics:

- [Change the Interface Type to SR-IOV](#)
- [Configure SR-IOV on a High Availability Setup](#)

Change the interface type to SR-IOV

You can run the show interface summary command to check the default configuration of a network interface.

Example 1: The following CLI screen capture shows the configuration of a network interface where SR-IOV is enabled by default on Citrix ADC VPX AWS Marketplace Editions of 3G and 5G.

```
> show interface summary
-----
Interface  MTU      MAC                Suffix
-----
1    1/1      1500              0a:1e:2e:17:a2:37  Intel 82599 10G VF Interface
2    L0/1     1500              0a:1e:2e:17:a2:37  Netscaler Loopback interface
Done
```

Example 2: The following CLI screen capture shows the default configuration of a network interface where SR-IOV is not enabled.

```

Done
[> sh int s
-----
Interface  MTU      MAC          Suffix
-----
1  1/1      1500      12:fc:04:c5:d0:12  NetScaler Virtual Interface
2  LO/1     1500      12:fc:04:c5:d0:12  Netscaler Loopback interface
Done
>

```

For more information about changing the interface type to SR-IOV, see <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/sriov-networking.html>

To change the interface type to SR-IOV

1. Shut down the Citrix ADC VPX instance running on AWS.
2. To enable SR-IOV on the network interface, type the following command in the AWS CLI.


```
$ aws ec2 modify-instance-attribute --instance-id <instance\_id
\> --sriov-net-support simple
```
3. To check if SR-IOV has been enabled, type the following command in the AWS CLI.


```
$ aws ec2 describe-instance-attribute --instance-id <instance\_
\_id\> --attribute sriovNetSupport
```

Example 3: Network interface type changed to SR-IOV, by using the AWS CLI.

```

aws ec2 modify-instance-attribute --instance-id i-008c1230aaf303bee --sriov-net-support simple
aws ec2 describe-instance-attribute --instance-id i-008c1230aaf303bee --attribute sriovNetSupport
{
  "InstanceId": "i-008c1230aaf303bee",
  "SriovNetSupport": {
    "Value": "simple"
  }
}

```

If SR-IOV is not enabled, value for SriovNetSupport is absent.

Example 4: In the following example, SR-IOV support is not enabled.

```

{
  "InstanceId": "i-0c3e84cfa65b04cc8",
  "SriovNetSupport": {}
}

```

- Power on the VPX instance. To see the changed status of the network interface, type “show interface summary” in the CLI.

Example 5: The following screen capture shows the network interfaces with SR-IOV enabled. The interfaces 10/1, 10/2, 10/3 are SR-IOV enabled.

```
> show interface summary
-----
Interface  MTU      MAC              Suffix
-----
1   10/1      1500            0a:1e:2e:17:a2:37  Intel 82599 10G VF Interface
2   10/2      1500            0a:df:17:0a:fe:83  Intel 82599 10G VF Interface
3   10/3      1500            0a:de:5d:31:bf:c3  Intel 82599 10G VF Interface
4   LO/1      1500            0a:1e:2e:17:a2:37  Netscaler Loopback interface
Done
```

These steps complete the procedure to configure VPX instances to use SR-IOV network interfaces.

Configure SR-IOV on a high availability setup

High availability is supported with SR-IOV interfaces from Citrix ADC release 12.0 build 57.19 onwards.

If the high availability setup was deployed manually or by using the Citrix CloudFormation template for Citrix ADC version 12.0 56.20 and lower, the IAM role attached to the high availability setup must have the following privileges:

- ec2:DescribeInstances
- ec2:DescribeNetworkInterfaces
- ec2:DetachNetworkInterface
- ec2:AttachNetworkInterface
- ec2:StartInstances
- ec2:StopInstances
- ec2:RebootInstances
- autoscaling:*
- sns:*
- sqs:*
- IAM:SimulatePrincipalPolicy
- IAM:GetRole

By default, the Citrix CloudFormation template for Citrix ADC version 12.0 57.19 automatically adds the required privileges to the IAM role.

Note:

A high availability setup with SR-IOV Interfaces takes around 100 seconds of downtime.

Related resources:

For more information about IAM roles, see [AWS documentation](#).

Configure a Citrix ADC VPX instance to use Enhanced Networking with AWS ENA

September 9, 2024

After you have created a Citrix ADC VPX instance on AWS, you can configure the virtual appliance to use [Enhanced Networking](#) with [AWS Elastic Network Adapter \(ENA\)](#), by using AWS CLI.

Coupled with AWS ENA, enhanced networking provides higher bandwidth, higher packet-per-second (PPS) performance, and consistently lower inter-instance latencies.

Before you start the configuration, read the following topics:

- [Prerequisites](#)
- [Limitations and Usage Guidelines](#)

The following HA configurations are supported for ENA-enabled instances:

- Private IP addresses can be moved within the same availability zone.
- Elastic IP addresses can be moved across availability zones.

Upgrade a Citrix ADC VPX instance on AWS

September 9, 2024

You can upgrade the EC2 instance type, throughput, software edition, and the system software of a Citrix ADC VPX running on AWS. For certain types of upgrades, Citrix recommends using the High Availability Configuration method to minimize downtime.

Note:

- Citrix ADC software release 10.1.e-124.1308.e or later for a Citrix ADC VPX AMI (including both utility license and customer license) does not support the M1 and M2 instance families.

- Because of changes in VPX instance support, downgrading from 10.1.e-124 or a later release to 10.1.123.x or an earlier release is not supported.
- Most of the upgrades do not require the launch of a new AMI, and the upgrade can be done on the current Citrix ADC AMI instance. If you do want to upgrade to a new Citrix ADC AMI instance, use the high availability configuration method.

Change the EC2 instance type of a Citrix ADC VPX instance on AWS

If your Citrix ADC VPX instances are running release 10.1.e-124.1308.e or later, you can change the EC2 instance type from the AWS console as follows:

1. Stop the VPX instance.
2. Change the EC2 instance type from the AWS console.
3. Start the instance.

You can also use the above procedure to change the EC2 instance type for a release, earlier than 10.1.e-124.1308.e, unless you want to change the instance type to M3. In that case, you must first follow the standard Citrix ADC upgrade procedure, at, to upgrade the Citrix ADC software to 10.1.e-124 or a later release, and then follow the above steps.

Upgrade the throughput or software edition of a Citrix ADC VPX instance on AWS

To upgrade the software edition (for example, to upgrade from Standard to Premium edition) or throughput (for example, to upgrade from 200 Mbps to 1000mbps), the method depends on the instance's license.

Using a customer license (Bring-Your-Own-License)

If you are using a customer license, you can purchase and download the new license from the Citrix website, and then install the license on the VPX instance. For more information about downloading and installing a license from the Citrix website, see the VPX Licensing Guide.

Using a utility license (Utility license with hourly fee)

AWS does not support direct upgrades for fee-based instances. To upgrade the software edition or throughput of a fee based Citrix ADC VPX instance, launch a new AMI with the desired license and capacity and migrate the older instance configuration to the new instance. This can be achieved by using a Citrix ADC high availability configuration as described in Upgrade to a new Citrix ADC AMI instance by using a Citrix ADC high availability configuration subsection in this page.

Upgrade the system software of a Citrix ADC VPX instance on AWS

If you need to upgrade a VPX instance running 10.1.e-124.1308.e or a later release, follow the standard Citrix ADC upgrade procedure at [Upgrade and downgrade a Citrix ADC appliance](#).

If you need to upgrade a VPX instance running a release older than 10.1.e-124.1308.e to 10.1.e-124.1308.e or a later release, first upgrade the system software, and then change the instance type to M3 as follows:

1. Stop the VPX instance.
2. Change the EC2 instance type from the AWS console.
3. Start the instance.

Upgrade to a new Citrix ADC AMI instance by using a Citrix ADC high availability configuration

To use the high availability method of upgrading to a new Citrix ADC AMI instance, perform the following tasks:

- Create a new instance with the desired EC2 instance type, software edition, throughput, or software release from the AWS marketplace.
- Configure high availability between the old instance (to be upgraded) and the new instance. After high availability is configured between the old and the new instance, configuration from the old instance is synchronized to the new instance.
- Force an HA failover from the old instance to the new instance. As a result, the new instance becomes primary and starts receiving traffic.
- Stop, and reconfigure or remove the old instance from AWS.

Prerequisites and points to consider

- Ensure you understand how high availability works between two Citrix ADC VPX instances on AWS. For more information about high availability configuration between two Citrix ADC VPX instances on AWS, see [Deploy a high availability pair on AWS](#).
- You must create the new instance in the same availability zone as the old instance, having the exact same security group and subnet.
- High availability setup requires access and secret keys associated with the user's AWS Identity and Access Management (IAM) account for both instances. If the correct key information is not used when creating VPX instances, the HA setup fails. For more information about creating an IAM account for a VPX instance, see [Prerequisites](#).
 - You must use the EC2 console to create the new instance. You cannot use the AWS 1-click launch, because it does not accept the access and secret keys as the input.

- The new instance must have only one ENI interface.

To upgrade a Citrix ADC VPX Instance by using a high availability configuration, follow these steps:

1. Configure high availability between the old and the new instance. To configure high availability between two Citrix ADC VPX instances, at the command prompt of each instance, type:
 - `add ha node <nodeID> <IPaddress of the node to be added>`
 - `save config`

Example:

At the command prompt of the old instance, type:

```
1 add ha node 30 192.0.2.30
2 Done
```

At the command prompt of the new instance, type:

```
1 add ha node 10 192.0.2.10
2 Done
```

Note the following:

- In the HA setup, the old instance is the primary node and the new instance is the secondary node.
- The NSIP IP address is not copied from the old instance to the new instance. Therefore, after the upgrade, your new instance has a different management IP address from the previous one.
- The `nsroot` account password of the new instance is set to that of the old instance after HA synchronization.

For more information about high availability configuration between two Citrix ADC VPX instances on AWS, see [Deploy a high availability pair on AWS](#).

2. Force an HA failover. To force a failover in a high availability configuration, at the command prompt of either of the instances, type:

```
1 force HA failover
```

As the result of forcing a failover, the ENIs of the old instance are migrated to the new instance and traffic flows through the new instance (the new primary node). The old instance (the new secondary node) restarts.

If the following warning message appears, type N to abort the operation:

```
1 [WARNING]:Force Failover may cause configuration loss, peer health
   not optimum. Reason(s):
2 HA version mismatch
```

```

3 HA heartbeats not seen on some interfaces
4 Please confirm whether you want force-failover (Y/N)?

```

The warning message appears because the system software of the two VPX instances is not HA compatible. As a result, the configuration of the old instance cannot be automatically synced to the new instance during a forced failover.

Following is the workaround for this issue:

- a) At the Citrix ADC shell prompt of the old instance, type the following command to create a backup of the configuration file (ns.conf):

```
copy /nsconfig/ns.conf to /nsconfig/ns.conf.bkp
```

- b) Remove the following line from the backup configuration file (ns.conf.bkp):

- `set ns config -IPAddress <IP> -netmask <MASK>`

Forexample, `set ns config -IPAddress 192.0.2.10 -netmask 255.255.255.0`

- c) Copy the old instance's backup configuration file (ns.conf.bkp) to the /nsconfig directory of the new instance.

- d) At the Citrix ADC shell prompt of the new instance, type the following command to load the old instance's configuration file (ns.conf.bkp) on the new instance:

- `batch -f /nsconfig/ns.conf.bkp`

- e) Save the configuration on the new instance.

- `save conifg`

- f) At the command prompt of either of the nodes, type the following command to force a failover, and then type Y for the warning message to confirm the force failover operation:

- `force ha failover`

Example:

```

1 > force ha failover
2
3 [WARNING]:Force Failover may cause configuration loss, peer health
  not optimum.
4 Reason(s):
5 HA version mismatch
6 HA heartbeats not seen on some interfaces
7 Please confirm whether you want force-failover (Y/N)? Y

```

3. Remove the HA configuration, so that the two instances are no longer in an HA configuration. First remove the HA configuration from the secondary node and then remove the HA configuration from the primary node.

To remove an HA configuration between two Citrix ADC VPX instances, at the command prompt of each instance, type:

```
1 > remove ha node \<nodeID\>
2 > save config
```

For more information about high availability configuration between two VPX instances on AWS, see [Deploy a high availability pair on AWS](#).

Example:

At the command prompt of the old instance (new secondary node), type:

```
1 > remove ha node 30
2 Done
3 > save config
4 Done
```

At the command prompt of the new instance (new primary node), type:

```
1 > remove ha node 10
2 Done
3 > save config
4 Done
```


Troubleshoot a VPX instance on AWS

September 6, 2024

Amazon does not provide console access to a Citrix ADC VPX instance. To troubleshoot, you have to use the AWS GUI to view the activity log. You can debug only if the network is connected. To view an instance's System Log, right-click the instance and select System Log.

Citrix provides support for AWS Marketplace-licensed Citrix ADC VPX instances (utility license with hourly fee) on AWS. To file a support case, find your AWS account number and support PIN code, and call Citrix support. You will also be asked for your name and email address. To find the support PIN, log on to the VPX GUI and navigate to the System page.

Here is an example of a system page showing the support PIN.



AWS FAQs

September 9, 2024

- **Does a Citrix ADC VPX instance support the encrypted volumes in AWS?**

Encryption and decryption happen at the hypervisor level, and hence it works seamlessly with any instance. For more information about the encrypted volumes see the following AWS document:

<https://docs.aws.amazon.com/kms/latest/developerguide/services-efs.html>

- **What is the best way to provision Citrix ADC VPX instance on AWS?**

You can provision a Citrix ADC VPX instance on AWS by any of the following ways:

- AWS CloudFormation Template (CFT) in AWS marketplace
- Citrix ADM
- AWS Quick Starts
- Citrix AWS CFTs in GitHub
- Citrix Terraform Scripts in GitHub
- Citrix Ansible Playbooks in GitHub
- AWS EC2 launch workflow

You can choose any of the listed options based on the automation tool that you use.

For more details about the options, see [Citrix ADC VPX on AWS](#).

- **How to upgrade Citrix ADC VPX instance in AWS?**

To upgrade the Citrix ADC VPX instance in AWS, you can upgrade the system software or upgrade to a new Citrix ADC VPX Amazon Machine Image (AMI) by following the procedure at [Upgrade a Citrix ADC VPX instance on AWS](#).

The recommended way to upgrade a Citrix ADC VPX instance is using the ADM service by following the procedure at [Use jobs to upgrade Citrix ADC instances](#).

- **What is the HA failover time for Citrix ADC VPX in AWS?**

- HA failover of Citrix ADC VPX within the AWS availability zone takes around 3 seconds.
- HA failover of Citrix ADC VPX across AWS availability zones takes around 5 seconds.

- **What level of support is provided for Citrix ADC VPX marketplace subscription customers who provide the technical support PIN?**

By default, the “Select for Software” service is provided to customers who provide the technical support PIN.

- **In High availability across different zones using Elastic IP deployment, do we need to create Multiple IP Sets for each application?**

Yes. If there are multiple applications with multiple VIPs mapped to multiple EIPs then multiple IP Sets are required. Therefore during HA failover, all the primary VIP mappings of EIPs are changed to secondary (new primary) VIPs.

- **Why is INC mode enabled in high availability across different zone deployments?**

HA pairs across availability zones are in different networks. For HA synchronization, network configuration must not be synchronized. This is achieved by enabling INC mode on HA pair.

- **Can HA node in one availability zone communicate with back-end servers in another availability zone, provided those availability zones are in same VPC?**

Yes, subnets in different availability zones of the same VPC are reachable by adding an extra route pointing to the backend-server subnet via SNIP. For example, if the SNIP subnet of ADC in AZ1 is 192.168.3.0/24 and the backend-server subnet in AZ2 is 192.168.6.0/24, then a route must be added in the Citrix ADC appliance present in AZ1 as 192.168.6.0 255.255.255.0 192.168.3.1.

- **Can High availability across different zones using Elastic IP and High availability across different zones using Private IP deployments work together?**

Yes, both the configurations can be applied on the same HA Pair.

- **In High availability across different zones using Private IP deployment, if there are multiple subnets with multiple route tables in a VPC, how does a secondary node in HA pair know about the route table to be checked during HA failover?**

Secondary node is aware of the primary NICs and searches across all the route tables in a VPC.

• **What is the size of the /var partition when using the default image for VPX on AWS? How to increase the disk space?**

The size of the root disk is limited to 20 GB to keep the disk image small.

If you want to increase the /var/core/ or the /var/crash/ directory space, attach an extra disk. To increase the /var size, currently, you must attach an extra disk and create a symbolic link to /var, after copying the critical contents to the new disk.

• **How many packet engines are activated and allocated to vCPUs?**

The packet engines (PEs) are limited by the number of licensed vCPUs. The Citrix ADC daemons are not pinned to any particular vCPU and might run on any of the non-PE vCPUs. According to AWS, the C5.9xlarge is a 36vCPU instance with 72 GB memory. With pooled licensing, the Citrix ADC VPX instance deploys with the maximum number of PEs. In this case, 19 PEs run on cores 1–19. However, ADC management processes run from CPUs 20–31.

• **How to decide the right AWS instance for ADC?**

1. Understand your use case and requirements like throughput, PPS, SSL requirement, and average packet size.
2. Choose the right ADC offering and licensing that meets your requirements, such as VPX bandwidth offerings or vCPU based licensing.
3. Based on the chosen offering, decide on the AWS instance.

Example:

A 5 Gbps license enables 5 data packet engines. Hence, the vCPU requirement is 6 (5+1 for management). But 6 vCPU instance is not available. So an 8 vCPU is good enough to reach that throughput provided you choose a network that supports 5 Gbps bandwidth. For example, you must choose m5.2xlarge for a 5 Gbps bandwidth license to enable max PE allocation for 5 Gbps license. But if you use vCPU license that is not limited by throughput, you might get 5 Gbps throughput using the m5.xlarge instance itself.

Instance Size	vCPU	Memory (GiB)	Instance Storage (GiB)	Network Bandwidth (Gbps)	EBS Bandwidth (Mbps)
m5.large	2	8	EBS-Only	Up to 10	Up to 4,750
m5.xlarge	4	16	EBS-Only	Up to 10	Up to 4,750
m5.2xlarge	8	32	EBS-Only	Up to 10	Up to 4,750
m5.4xlarge	16	64	EBS-Only	Up to 10	4,750

• **Is three NICs-three subnets deployment mandatory for ADC in AWS?**

Three NICs-three subnets is the recommended deployment, where each one for management, client and server network. This deployment gives better traffic isolation and VPX performance. Two NICs-two subnets, and one NIC-one subnet are the other available options. It is not recommended to have multiple NICs sharing a subnet in AWS, such as a two NICs—one subnet deployment. This scenario can cause networking issues like asymmetric routing. For more information, see [Best practices for configuring network interfaces in AWS](#).

- **Why does an ENA driver on AWS always indicate a 1Gbps (1/1) link speed, irrespective of the instance's network capabilities?**

The reported speed of an AWS Elastic Network Adapter (ENA) is often displayed as 1Gbps (1/1) regardless of the selected instance type. This is because the indicated speed does not directly reflect the actual network performance. Unlike traditional network interfaces, ENA speeds can dynamically scale based on the instance's requirements and workload. The true network performance is primarily determined by the instance type and size. Therefore, the actual network throughput can vary significantly depending on the specific instance type and the current network load.

Deploy a Citrix ADC VPX instance on Microsoft Azure

September 6, 2024

When you deploy a Citrix ADC VPX instance on Microsoft Azure Resource Manager (ARM), you can use both of the following feature sets to achieve your business needs:

- Azure cloud computing capabilities
- Citrix ADC load balancing and traffic management features

You can deploy Citrix ADC VPX instances on ARM either as standalone instances or as high availability pairs in active-standby modes.

You can deploy a Citrix ADC VPX instance on the Microsoft Azure in two ways:

- Through Azure Marketplace. The Citrix ADC VPX virtual appliance is available as an image in the Microsoft Azure Marketplace.
- Using the Citrix ADC Azure Resource Manager (ARM) json template available on GitHub. For more information, see the [GitHub repository for Citrix NetScaler solution templates](#).

The Microsoft Azure stack is an integrated platform of hardware and software that delivers the Microsoft Azure public cloud services in a local data center to let organizations construct hybrid clouds. You can now deploy the Citrix ADC VPX instances on the Microsoft Azure stack.

Prerequisite

You need some prerequisite knowledge before deploying a Citrix VPX instance on Azure.

- Familiarity with Azure terminology and network details. For information, see [Azure terminology](#).
- Knowledge of a Citrix ADC appliance. For detailed information the Citrix ADC appliance, see [Citrix ADC](#)
- Knowledge of Citrix ADC networking. See the [Networking](#) topic.

How a Citrix ADC VPX instance works on Azure

In an on-premises deployment, a Citrix ADC VPX instance requires at least three IP addresses:

- Management IP address, called NSIP address
- Subnet IP (SNIP) address for communicating with the server farm
- Virtual server IP (VIP) address for accepting client requests

For more information, see [Network architecture for Citrix ADC VPX instances on Microsoft Azure](#).

Note:

VPX virtual appliances can be deployed on any instance type that has two or more Intel VT-X cores and more than 2 GB memory. For more information on system requirements, see [Citrix ADC VPX data sheet](#). Currently, Citrix ADC VPX instance supports only the Intel processors.

In an Azure deployment, you can provision a Citrix ADC VPX instance on Azure in three ways:

- Multi-NIC multi-IP architecture
- Single NIC multi IP architecture
- Single NIC single IP

Depending on your need, you can use any of these supported architecture types.

Multi-NIC multi-IP architecture

In this deployment type, you can have more than one network interfaces (NICs) attached to a VPX instance. Any NIC can have one or more IP configurations - static or dynamic public and private IP addresses assigned to it.

For more information, see the following use cases:

- [Configure a high-availability setup with multiple IP addresses and NICs](#)

- [Configure a high-availability setup with multiple IP addresses and NICs by using PowerShell commands](#)

Note:

To avoid MAC moves and interface mutes on Azure environments, Citrix recommends you to create a VLAN per data interface (without tag) of ADC VPX instance and bind the primary IP of NIC in Azure. For more information, see [CTX224626](#) article.

Single NIC multi IP architecture

In this deployment type, one network interfaces (NIC) associated with multiple IP configurations - static or dynamic public and private IP addresses assigned to it.

For more information, see the following use cases:

- [Configure multiple IP addresses for a Citrix ADC VPX standalone instance](#)
- [Configure multiple IP addresses for a Citrix ADC VPX standalone instance by using PowerShell commands](#)

Single NIC single IP

In this deployment type, one network interfaces (NIC) associated with a single IP address, which is used to perform the functions of NSIP, SNIP, and VIP.

For more information, see the following use case:

- [Configure a Citrix ADC VPX standalone instance](#)

Note:

The single IP mode is available only in Azure deployments. This mode is not available for a Citrix ADC VPX instance on your premises, on AWS, or in other type of deployment.

Citrix ADC VPX licensing

A Citrix ADC VPX instance on Azure requires a license. The following licensing options are available for Citrix ADC VPX instances running on Azure.

- **Subscription-based licensing:** Citrix ADC VPX appliances are available as paid instances on Azure Marketplace. Subscription-based licensing is a pay-as-you-go option. Users are charged hourly.

Note:

For subscription-based license instances, your subscription billing applies throughout the license period for a particular license model. Due to cloud restrictions, Azure does not support changing or removing the license model applicable for your subscription. To change or remove a subscription license, delete the existing ADC VM, and recreate a new ADC VM with desired license.

Citrix provides technical support for subscription-based license instances. To file a support case, see [Support for Citrix ADC on Azure –Subscription license with hourly price](#).

- **Bring your own license (BYOL):** If you bring your own license (BYOL), see the VPX Licensing Guide at <http://support.citrix.com/article/CTX122426>. You have to:
 - Use the licensing portal within Citrix website to generate a valid license.
 - Upload the license to the instance.

Note:

In an Azure stack environment, **BYOL** is the only available licensing option.

- **Citrix ADC VPX Check-In/Check-Out licensing:** For more information, see [Citrix ADC VPX Check-In/Check-Out Licensing](#).

Starting with NetScaler release 12.0 56.20, Citrix ADC VPX Express for on-premises and cloud deployments does not require a license file. For more information on ADC VPX Express, see the “Citrix ADC VPX Express license” section in [Citrix ADC licensing overview](#).

The following VPX models and license types are available on Azure Marketplace.

VPX model	License type	Recommended instances		
		VPX 1 NIC/2 NIC	VPX 3 NIC	VPX upto 8 NIC
VPX200	Advanced	Standard_D2s_v4	Standard_DS3_v2	Standard_DS4_v2
VPX1000	Premium	Standard_D4s_v4	Standard_DS3_v2	Standard_DS4_v2
VPX5000	Premium	Standard_D8ds_v5	Standard_D8ds_v5	Standard_DS4_v2
VPX BYOL	Customer Licensed FIPS - Customer Licensed	-	-	-

Note:

The recommended instances for VPX BYOL depends on the VPX license that you have purchased.

Points to note:

- You must enable Azure accelerated networking on NetScaler VPX instances to get the optimal performance on the following VPX models:
 - VPX1000
 - VPX5000

For more information on configuring Accelerated networking, see [Configure a Citrix ADC VPX instance to use Azure accelerated networking] (/en-us/vpx/current-release/deploy-vpx-on-azure/configure-vpx-to-use-azure-accelerated-networking.html)

- The VPX8000 and VPX10000 licenses are available only as BYOL.
- Regardless of the subscription-based hourly license bought from Azure Marketplace, in rare cases, the Citrix ADC VPX instance deployed on Azure might come up with a default Citrix ADC license. This happens due to issues with the Azure Instance Metadata Service (IMDS).
- Do a warm restart, before making any configuration change on the Citrix ADC VPX instance, to enable the correct Citrix ADC VPX license.

Limitations

Running the Citrix ADC VPX load balancing solution on ARM imposes the following limitations:

- The Azure architecture does not accommodate support for the following NetScaler features:
 - IPv6
 - Gratuitous ARP (GARP)
 - L2 Mode
 - Tagged VLAN
 - Dynamic Routing
 - virtual MAC
 - USIP
 - Jumbo Frames
 - Clustering

Note:

With the Citrix Application Delivery Management (ADM) Autoscale feature (cloud deployment), the ADC instances support clustering on all licenses. For information, see [Autoscal-](#)

ing of Citrix ADC VPX in Microsoft Azure using Citrix ADM.

- If you expect that you might have to shut down and temporarily deallocate the Citrix ADC VPX virtual machine at any time, assign a static Internal IP address while creating the virtual machine. If you do not assign a static internal IP address, Azure might assign the virtual machine a different IP address each time it restarts, and the virtual machine might become inaccessible.
- In an Azure deployment, only the following Citrix ADC VPX models are supported: VPX 10, VPX 200, VPX 1000, and VPX 3000. For for information, see the Citrix ADC VPX Data Sheet.

If you use a Citrix ADC VPX instance with a model number higher than VPX 3000, the network throughput might not be the same as specified by the instance's license. However, other features such as SSL throughput and SSL transactions per second might improve.

- The “deployment ID” that is generated by Azure during virtual machine provisioning is not visible to the user in ARM. You cannot use the deployment ID to deploy Citrix ADC VPX appliance on ARM.
- The Citrix ADC VPX instance supports 20 Mb/s throughput and standard edition features when it's initialized.
- The Citrix ADC VPX instances on Azure with accelerated networking enabled, provides better performance. Azure accelerated networking is supported on Citrix ADC VPX instances from release 13.0 build 76.x onwards. To enable accelerated networking on ADC VPX, Citrix recommends you to use an Azure instance type which supports accelerated networking.
- For Citrix Virtual Apps and Citrix Virtual Desktops deployment, a VPN virtual server on a VPX instance can be configured in the following modes:
 - Basic mode, where the `ICAOnly` VPN virtual server parameter is set to ON. The Basic mode works fully on an unlicensed Citrix ADC VPX instance.
 - SmartAccess mode, where the `ICAOnly` VPN virtual server parameter is set to OFF. The SmartAccess mode works for only five Citrix ADC AAA session users on an unlicensed Citrix ADC VPX instance.

Note:

To configure the SmartControl feature, you must apply a Premium license to the Citrix ADC VPX instance.

Azure terminology

September 6, 2024


Some of the Azure terms that are used in the Citrix ADC VPX Azure documentation are listed below.

1. Azure Load Balancer –Azure load balancer is a resource that distributes incoming traffic among computers in a network. Traffic is distributed among virtual machines defined in a load-balancer set. A load balancer can be external or internet-facing, or it can be internal.
2. Azure Resource Manager (ARM) –ARM is the new management framework for services in Azure. Azure Load Balancer is managed using ARM-based APIs and tools.
3. Back-End Address Pool –These are IP addresses associated with the virtual machine NIC (NIC) to which load will be distributed.
4. BLOB - Binary Large Object –Any binary object like a file or an image that can be stored in Azure storage.
5. Front-End IP Configuration –An Azure Load balancer can include one or more front-end IP addresses, also known as a virtual IPs (VIPs). These IP addresses serve as ingress for the traffic.
6. Instance Level Public IP (ILPIP) –An ILPIP is a public IP address that you can assign directly to your virtual machine or role instance, rather than to the cloud service that your virtual machine or role instance resides in. This does not take the place of the VIP (virtual IP) that is assigned to your cloud service. Rather, it's an extra IP address that you can use to connect directly to your virtual machine or role instance.

Note:

In the past, an ILPIP was referred to as a PIP, which stands for public IP.

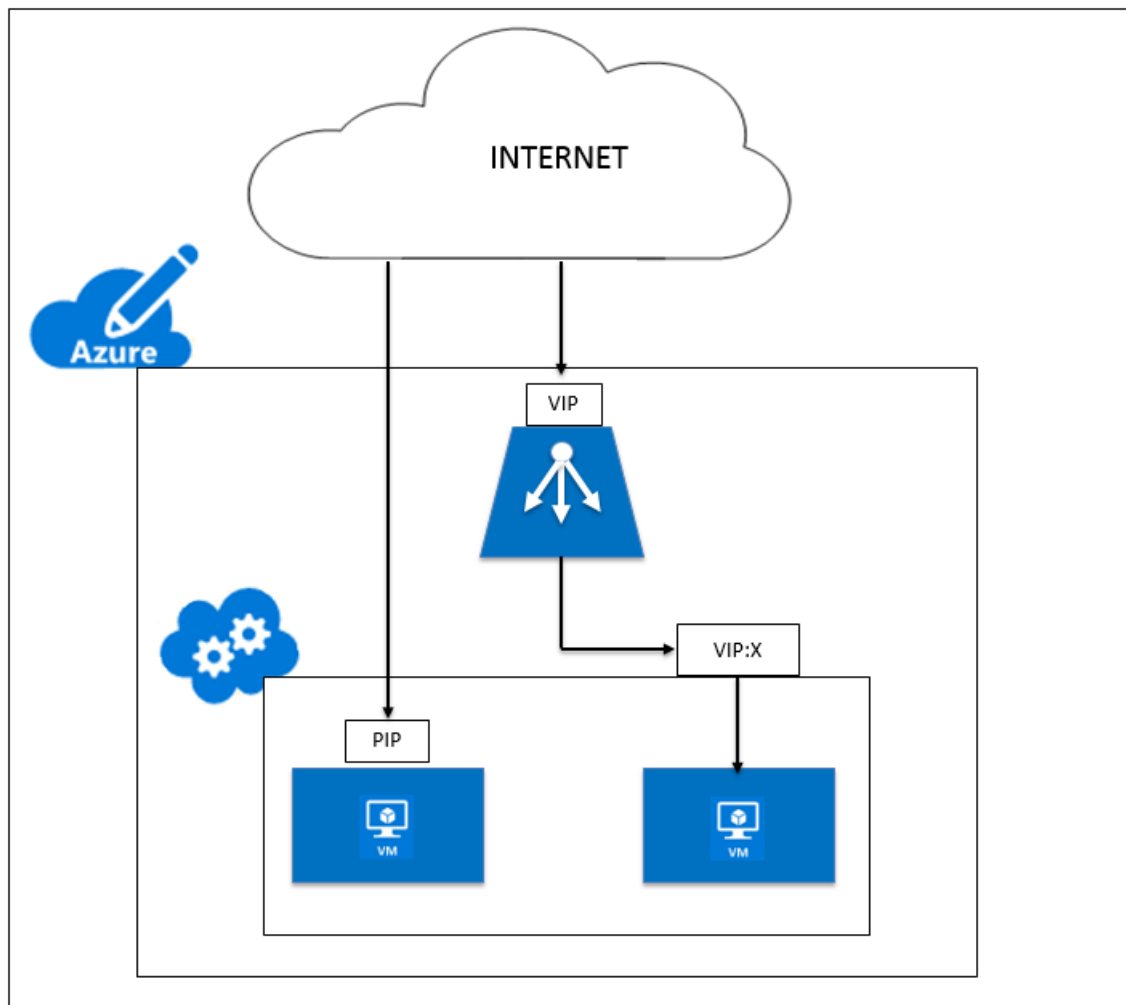
7. Inbound NAT Rules –This contains rules mapping a public port on the load balancer to a port for a specific virtual machine in the back end address pool.
8. IP-Config - It can be defined as an IP address pair (public IP and private IP) associated with an individual NIC. In an IP-Config, the public IP address can be NULL. Each NIC can have multiple IP-Config associated with it, which can be up to 255.
9. Load Balancing Rules –A rule property that maps a given front-end IP and port combination to a set of back-end IP addresses and port combination. With a single definition of a load balancer resource, you can define multiple load balancing rules, each rule reflecting a combination of a front end IP and port and back end IP and port associated with virtual machines.



10. Network security group –Contains a list of Access Control List (ACL) rules that allow or deny network traffic to your virtual machine instances in a virtual network. NSGs can be associated with either subnets or individual virtual machine instances within that subnet. When a network security group is associated with a subnet, the ACL rules apply to all the virtual machine instances in that subnet. In addition, traffic to an individual virtual machine can be restricted further by associating a network security group directly to that virtual machine.
11. Private IP addresses –Used for communication within an Azure virtual network, and your on-premises network when you use a VPN gateway to extend your network to Azure. Private IP addresses allow Azure resources to communicate with other resources in a virtual network or an on-premises network through a VPN gateway or ExpressRoute circuit, without using an Internet-reachable IP address. In the Azure Resource Manager deployment model, a private IP address is associated with the following types of Azure resources –virtual machines, internal load balancers (ILBs), and application gateways.
12. Probes –This contains health probes used to check availability of virtual machines instances in the back end address pool. If a particular virtual machine does not respond to health probes for some time, then it is taken out of traffic serving. Probes enable you to keep track of the

health of virtual instances. If a health probe fails, the virtual instance will be taken out of rotation automatically.

13. Public IP Addresses (PIP) –PIP is used for communication with the Internet, including Azure public-facing services and is associated with virtual machines, Internet-facing load balancers, VPN gateways, and application gateways.
14. Region - An area within a geography that does not cross national borders and that contains one or more data centers. Pricing, regional services, and offer types are exposed at the region level. A region is typically paired with another region, which can be up to several hundred miles away, to form a regional pair. Regional pairs can be used as a mechanism for disaster recovery and high availability scenarios. Also referred to generally as location.
15. Resource Group - A container in Resource Manager holds related resources for an application. The resource group can include all of the resources for an application, or only those resources that are logically grouped together
16. Storage Account –An Azure storage account gives you access to the Azure blob, queue, table, and file services in Azure Storage. Your storage account provides the unique namespace for your Azure storage data objects.
17. Virtual Machine –The software implementation of a physical computer that runs an operating system. Multiple virtual machines can run simultaneously on the same hardware. In Azure, virtual machines are available in a variety of sizes.
18. Virtual Network - An Azure virtual network is a representation of your own network in the cloud. It is a logical isolation of the Azure cloud dedicated to your subscription. You can fully control the IP address blocks, DNS settings, security policies, and route tables within this network. You can also further segment your VNet into subnets and launch Azure IaaS virtual machines and cloud services (PaaS role instances). Additionally, you can connect the virtual network to your on-premises network using one of the connectivity options available in Azure. In essence, you can expand your network to Azure, with complete control on IP address blocks with the benefit of enterprise scale Azure provides.



Network architecture for Citrix ADC VPX instances on Microsoft Azure

September 12, 2024

In Azure Resource Manager (ARM), a Citrix ADC VPX virtual machine (VM) resides in a virtual network. A single network interface can be created in a given subnet of the Virtual Network and can be attached to the VPX instance. You can filter network traffic to and from a VPX instance in an Azure virtual network with a network security group. A network security group contains security rules that allow or deny inbound network traffic to or outbound network traffic from a VPX instance. For more information, see [Security groups](#).

Network security group filters the requests to the Citrix ADC VPX instance, and the VPX instance sends them to the servers. The response from a server follows the same path in reverse. The Network security group can be configured to filter a single VPX VM, or, with subnets and virtual networks, can filter


traffic in deployment of multiple VPX instances.

The NIC contains network configuration details such as the virtual network, subnets, internal IP address, and Public IP address.

While on ARM, it is good to know the following IP addresses that are used to access the VMs deployed with a single NIC and a single IP address:

- Public IP (PIP) address is the internet-facing IP address configured directly on the virtual NIC of the NetScaler VM. This allows you to directly access a VM from the external network.
- Citrix ADC IP (also known as NSIP) address is the internal IP address configured on the VM. It is non-routable.
- Virtual IP address (VIP) is configured by using the NSIP and a port number. Clients access NetScaler services through the PIP address, and when the request reaches the NIC of the NetScaler VPX VM or the Azure load balancer, the VIP gets translated to internal IP (NSIP) and internal port number.
- Internal IP address is the private internal IP address of the VM from the virtual network's address space pool. This IP address cannot be reached from the external network. This IP address is by default dynamic unless you set it to static. Traffic from the internet is routed to this address according to the rules created on the network security group. The network security group integrates with the NIC to selectively send the right type of traffic to the right port on the NIC, which depends on the services configured on the VM.


The following figure shows how traffic flows from a client to a server through a NetScaler VPX instance provisioned in ARM.



Traffic flow through network address translation

You can also request a public IP (PIP) address for your Citrix ADC VPX instance (instance level). If you use this direct PIP at the VM level, you need not define inbound and outbound rules to intercept the network traffic. The incoming request from the Internet is received on the VM directly. Azure performs network address translation (NAT) and forwards the traffic to the internal IP address of the VPX instance.

The following figure shows how Azure performs network address translation to map the NetScaler internal IP address.



In this example, the Public IP assigned to the network security group is 140.x.x.x and the internal IP address is 10.x.x.x. When the inbound and outbound rules are defined, public HTTP port 80 is defined as the port on which the client requests are received, and a corresponding private port, 10080, is defined as the port on which the Citrix ADC VPX instance listens. The client request is received on the Public IP address (140.x.x.x). Azure performs network address translation to map the PIP to the internal IP address 10.x.x.x on port 10080, and forwards the client request.

Note:

Citrix ADC VPX VMs in high availability are controlled by external or internal load balancers that have inbound rules defined on them to control the load balancing traffic. The external traffic is first intercepted by these load balancers and the traffic is diverted according to the load balancing rules configured, which has back-end pools, NAT rules, and health probes defined on the load balancers.

Port usage guidelines

You can configure more inbound and outbound rules in network security group while creating the Citrix ADC VPX instance or after the virtual machine is provisioned. Each inbound and outbound rule

is associated with a public port and a private port.

Before configuring network security group rules, note the following guidelines regarding the port numbers you can use:

1. The Citrix ADC VPX instance reserves the following ports. You cannot define these as private ports when using the Public IP address for requests from the internet.

Ports 21, 22, 80, 443, 8080, 67, 161, 179, 500, 520, 3003, 3008, 3009, 3010, 3011, 4001, 5061, 9000, 7000.

However, if you want internet-facing services such as the VIP to use a standard port (for example, port 443) you have to create port mapping by using the network security group. The standard port is then mapped to a different port that is configured on the NetScaler for this VIP service.

For example, a VIP service might be running on port 8443 on the VPX instance but be mapped to public port 443. So, when the user accesses port 443 through the Public IP, the request is directed to private port 8443.

2. Public IP address does not support protocols in which port mapping is opened dynamically, such as passive FTP or ALG.
3. High availability does not work for traffic that uses a public IP address (PIP) associated with a VPX instance, instead of a PIP configured on the Azure load balancer.

Note:

In Azure Resource Manager, a Citrix ADC VPX instance is associated with two IP addresses - a public IP address (PIP) and an internal IP address. While the external traffic connects to the PIP, the internal IP address or the NSIP is non-routable. To configure VIP in VPX, use the internal IP address and any of the free ports available. Do not use the PIP to configure VIP.

Configure a Citrix ADC VPX standalone instance

December 14, 2024

You can provision a single Citrix ADC VPX instance in Azure Resource Manager (ARM) portal in a standalone mode by creating the virtual machine and configuring other resources.

Before you begin

Ensure that you have the following:

- A Microsoft Azure user account
- Access to Microsoft Azure Resource Manager
- Microsoft Azure SDK
- Microsoft Azure PowerShell

On the [Microsoft Azure Portal](#) page, log on to the Azure Resource Manager portal by providing your user name and password.

Note:

In ARM portal, clicking an option in one pane opens a new pane to the right. Navigate from one pane to another to configure your device.

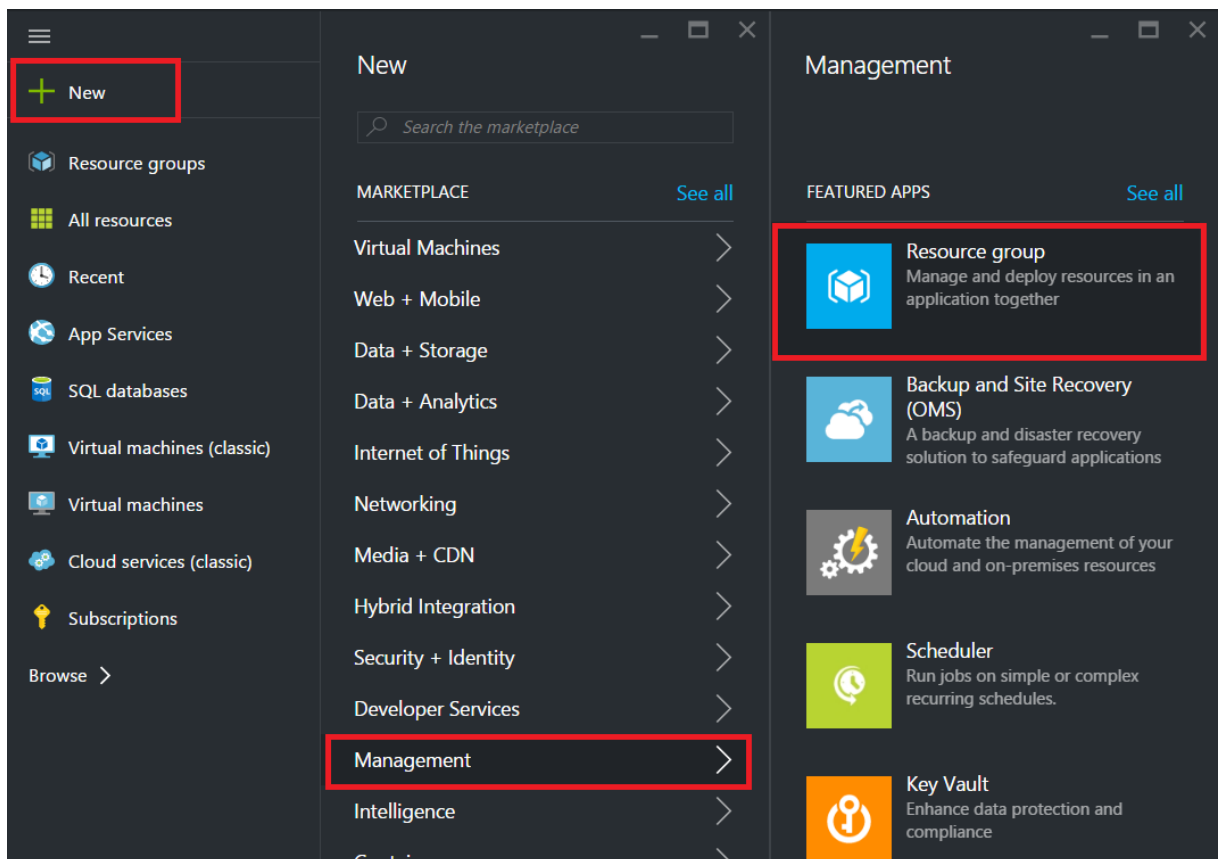
Summary of configuration steps

1. Configure a resource group
2. Configure a network security group
3. Configure virtual network and its subnets
4. Configure a storage account
5. Configure an availability set
6. Configure a Citrix ADC VPX instance.

Configure a resource group

Create a new resource group that is a container for all your resources. Use the resource group to deploy, manage, and monitor your resources as a group.

1. Click **New > Management > Resource group**.
2. In the **Resource group** pane, enter the following details:
 - Resource group name
 - Resource group location
3. Click **Create**.




Configure a network security group

Create a network security group to assign inbound and outbound rules to control the incoming and outgoing traffic within the virtual network. Network security group allows you to define security rules for a single virtual machine and also to define security rules for a virtual network subnet.

1. Click **New > Networking > Network security group**.
2. In the **Create network security group** pane, enter the following details, and then click **Create**.
 - Name - type a name for the security group
 - Resource group - select the resource group from the drop-down list

Note:

Ensure that you have selected the correct location. The list of resources that appear in the drop-down list is different for different locations.




Configure a virtual network and subnets

Virtual networks in ARM provide a layer of security and isolation to your services. VMs and services that are part of the same virtual network can access each other.

For these steps to create a virtual network and subnets.

1. Click **New > Networking > Virtual Network**.
2. In the **Virtual Network** pane, ensure the deployment mode is **Resource Manager** and click **Create**.



3. In the **Create virtual network** pane, enter the following values, and then click **Create**.

- Name of the virtual network
- Address space - type the reserved IP address block for the virtual network
- Subnet - type the name of the first subnet (you create the second subnet later in this step)
- Subnet address range - type the reserved IP address block of the subnet
- Resource group - select the resource group created earlier from the drop-down list

Create virtual network

* Name
NetScalerVNet ✓

* Address space ⓘ
22.22.0.0/16 ✓
22.22.0.0 - 22.22.255.255 (65536 addresses)

* Subnet name
NSFrontEnd ✓

* Subnet address range ⓘ
22.22.1.0/24 ✓
22.22.1.0 - 22.22.1.255 (256 addresses)

* Subscription
Microsoft Azure Enterprise ▼

* Resource group ⓘ
 Create new Use existing
NSDocs ▼


* Location
Southeast Asia ▼

Pin to dashboard

Create [Automation options](#)

Configure the second subnet

1. Select the newly created virtual network from **All resources** pane and in the **Settings** pane, click **Subnets**.



2. Click **+Subnet** and create the second subnet by entering the following details.
 - Name of the second subnet
 - Address range - type the reserved IP address block of the second subnet
 - Network security group - select the network security group from the drop-down list
3. Click **Create**.

Add subnet
NetScalerVNet

* Name
NSBackEnd ✓

* Address range (CIDR block) ⓘ
22.22.2.0/24 ✓
22.22.2.0 - 22.22.2.255 (256 addresses)

Network security group
None >

Route table
None >


OK

Configure a storage account

The ARM IaaS infrastructure storage includes all services where we can store data in the form of blobs, tables, queues, and files. You can also create applications using these forms of storage data in ARM.

Create a storage account to store all your data.

1. Click **+New > Data + Storage > Storage account**.
2. In the **Create storage account** pane, enter the following details:
 - Name of the account
 - Deployment mode - make sure to select **Resource Manager**
 - Account kind - select **General purpose** from the drop-down list
 - Replication - select **Locally redundant storage** from the drop-down list
 - Resource group - select the newly created resource group from the drop-down list
3. Click **Create**.




Configure an availability set

An availability set guarantee that at least one VM is kept up and running in case of planned or un-planned maintenance. Two or more VMs under the same 'availability set' are placed on different fault

domains to achieve redundant services.

1. Click **+New**.
2. Click **See all** in the MARKETPLACE pane and click **Virtual Machines**.
3. Search for availability set, and then select **Availability set** entity from the list displayed.



4. Click **Create**, and in the **Create availability set** pane, enter the following details:
 - Name of the set
 - Resource group - select the newly created resource group from the drop-down list
5. Click **Create**.

Create availability set

* Name
 ✓

Fault domains ⓘ
 3

Update domains ⓘ
 5

* Subscription
 ▼

* Resource group ⓘ
 Create new Use existing
 ▼

* Location
 ▼

Create

Configure a Citrix ADC VPX instance

Create an instance of Citrix ADC VPX in the virtual network. Obtain the Citrix ADC VPX image from the Azure Marketplace, and then use the Azure Resource Manager portal to create a Citrix ADC VPX instance.

Before you begin creating the Citrix ADC VPX instance, make sure that you have created a virtual network with required subnets in which the instance resides. You can create virtual networks during VM

provisioning, but without the flexibility to create different subnets.

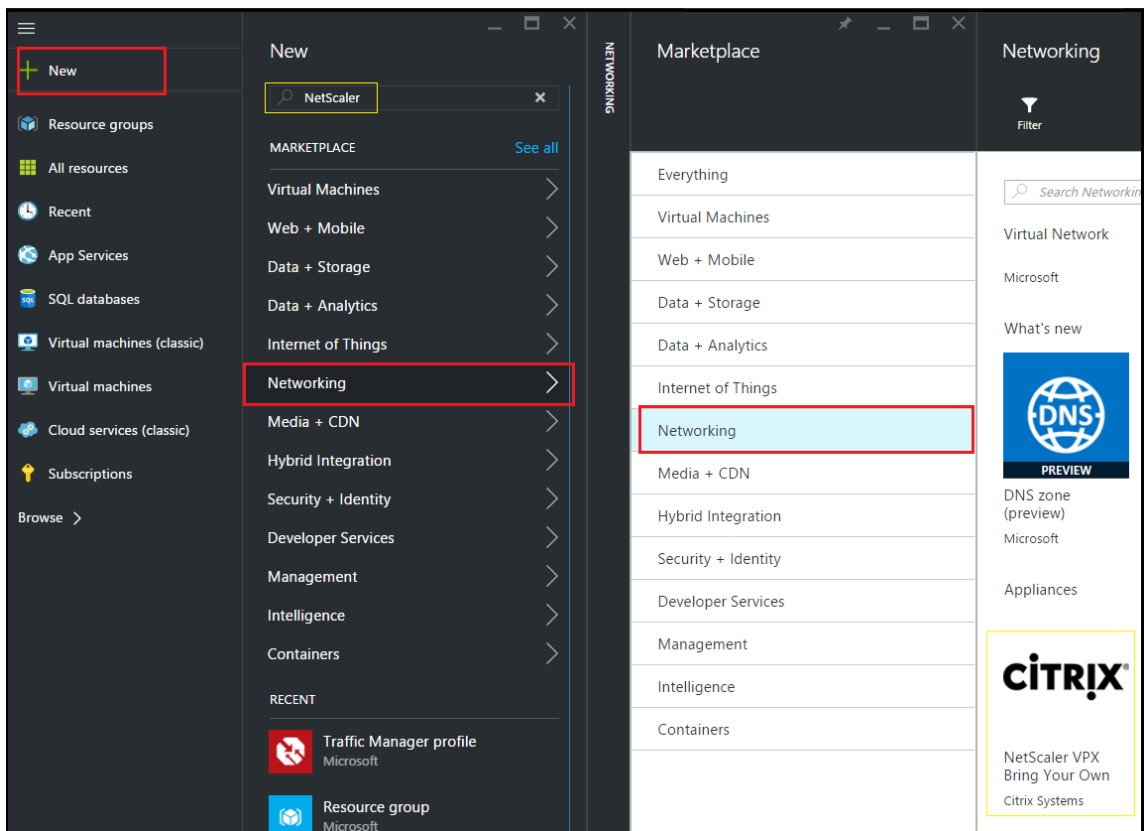
Optionally, configure DNS server and VPN connectivity that allows a virtual machine to access internet resources.

Note:

Citrix recommends that you create resource group, network security group, virtual network, and other entities before you provision the Citrix ADC VPX VM, so that the network information is available during provisioning.

1. Click **+New > Networking**.
2. Click **See All** and in the Networking pane, click **Citrix ADC 13.0**.
3. Select **Citrix ADC 13.0 VPX Bring Your Own License** from the list of software plans.

As a quick way to find any entity on ARM portal, you can also type the name of the entity in the Azure Marketplace search box and press <Enter>. Type NetScaler in the search box to find the Citrix NetScaler images.



Note:

Ensure to select the latest image. Your Citrix NetScaler image might have the release number in the name.

- On the **Citrix ADC VPX Bring Your Own License** page, from the drop-down list, select **Resource Manager** and click **Create**.

The screenshot shows the 'Create virtual machine' wizard in the 'Basics' step. The left sidebar contains a progress indicator with five steps: 1. Basics (Configure basic settings), 2. Size (Choose virtual machine size), 3. Settings (Configure optional features), 4. Summary (NetScaler 11.1 VPX Bring Your ...), and 5. Buy. The main area displays the following configuration options:

- Name:** Citrix-NetScaler-User (with a green checkmark)
- VM disk type:** SSD (dropdown menu)
- User name:** CitrixUser1 (with a green checkmark)
- Authentication type:** SSH public key and Password (radio buttons, with Password selected)
- Password:** (masked with dots, with a green checkmark)
- Confirm password:** (masked with dots, with a green checkmark)
- Subscription:** Microsoft Azure Enterprise (dropdown menu)
- Resource group:** Create new (radio button) and Use existing (radio button, selected). Below it is a dropdown menu showing NetScalerResGroup.
- Location:** Southeast Asia (dropdown menu)

An **OK** button is located at the bottom of the configuration pane.

- In the **Create virtual machine** pane, specify the required values in each section to create a virtual machine. Click **OK** in each section to save your configuration.

Basic:

- Name - specify a name for the Citrix ADC VPX instance
- VM disk type - select SSD (default value) or HDD from the drop-down menu
- User name and Password - specify a user name and password to access the resources in the resource group that you have created

- Authentication Type - select SSH Public Key or Password
- Resource group - select the resource group you have created from the drop-down list

You can create a resource group here, but Citrix recommends that you create a resource group from Resource groups in Azure Resource Manager and then select the group from the drop-down list.

Note:

In an Azure stack environment, in addition to the basic parameters, specify the following parameters:

- Azure stack domain
- Azure stack tenant (Optional)
- Azure client (Optional)
- Azure client secret (Optional)

Size:

Depending on the VM disk type, SDD, or HDD, you selected in Basic settings, the disk sizes are displayed.

- Select a disk size according to your requirement and click **Select**.

Settings:

- Select the default (Standard) disk type
- Storage account - select the storage account
- Virtual network - select the virtual network
- Subnet - set the subnet address
- Public IP address - select the type of IP address assignment
- Network security group - select the security group that you have created. Ensure that inbound and outbound rules are configured in the security group.
- Availability Set - select the availability set from the drop-down menu box

Summary:

The configuration settings are validated and the Summary page displays the result of the validation. If the validation fails, the Summary page displays the reason of the failure. Go back to the particular section and make changes as required. If the validation passes, click **OK**.

Buy:

Review the offer details and legal terms on the Purchase page and click **Purchase**.

For high availability deployment, create two independent instances of Citrix ADC VPX in the same availability set and in the same resource group to deploy them in active-standby configuration.

Configure multiple IP addresses for a Citrix ADC VPX standalone instance

September 9, 2024

This section explains how to configure a standalone Citrix ADC VPX instance with multiple IP addresses, in Azure Resource Manager (ARM). The VPX instance can have one or more NIC attached to it, and each NIC can have one or more static or dynamic public and private IP addresses assigned to it. You can assign multiple IP addresses as NSIP, VIP, SNIP, and so on.

For more information, see the Azure documentation [Assign multiple IP addresses to virtual machines using the Azure portal](#).

If you want to use PowerShell commands, see [Configuring multiple IP addresses for a Citrix ADC VPX instance in standalone mode by using PowerShell commands](#).

Use case

In this use case, a standalone Citrix ADC VPX appliance is configured with a single NIC that is connected to a virtual network (VNET). The NIC is associated with three IP configurations (ipconfig), each server a different purpose - as shown in the table.


IP config	Associated with	Purpose
ipconfig1	Static public IP address; static private IP address	Serves management traffic
ipconfig2	Static public IP address; static private address	Serves client-side traffic
ipconfig3	Static private IP address	Communicates with back-end servers

Note:

`IPConfig-3` is not associated with any public IP address.

Diagram: Topology

Here is the visual representation of the use case.

**Note:**

In a multi-NIC, multi-IP Azure Citrix ADC VPX deployment, the private IP associated with the primary (first) `IPConfig` of the primary (first) NIC is automatically added as the management NSIP of the appliance. The remaining private IP addresses associated with `IPConfigs` need to be added in the VPX instance as a VIP or SNIP by using the `add ns ip` command, according to your requirement.

Before you begin

Before you begin, create a VPX instance by following the steps given at this link:

[Configure a Citrix ADC VPX standalone instance](#)

For this use case, the NSDoc0330VM VPX instance is created.

Procedure to configure multiple IP addresses for a Citrix ADC VPX instance in standalone mode.


For configuring multiple IP addresses for a Citrix ADC VPX appliance in standalone mode:

1. Add IP addresses to the VM

2. Configure Citrix ADC -owned IP addresses


Step 1: Add IP addresses to the VM

1. In the portal, click **More services > type virtual machines** in the filter box, and then click **Virtual machines**.
2. In the **Virtual machines** blade, click the VM you want to add IP addresses to. Click **Network interfaces** in the virtual machine blade that appears, and then select the network interface.



In the blade that appears for the NIC you selected, click **IP configurations**. The existing IP configuration that was assigned when you created the VM, **ipconfig1**, is displayed. For this use case, make sure the IP addresses associated with ipconfig1 are static. Next, create two more IP configurations: ipconfig2 (VIP) and ipconfig3 (SNIP).

To create more **ipconfigs**, create **Add**.



In the **Add IP configuration** window, enter a **Name**, specify allocation method as **Static**, enter an IP address (192.0.0.5 for this use case), and enable **Public IP address**.

Note

Before adding a static private IP address, check for IP address availability and make sure the IP address belongs to the same subnet to which the NIC is attached.

The screenshot shows the 'Add IP configuration' window for a VM named 'nsdoc0330vm923'. The configuration is for a secondary IP address with the name 'ipconfig2'. The type is set to 'Secondary'. A message states 'Primary IP configuration already exists'. Under 'Private IP address settings', the allocation is set to 'Static' and the IP address is '192.0.0.5'. The 'Public IP address' is set to 'Enabled'. At the bottom, there is a button labeled 'Configure required settings' with a right-pointing arrow.

Next, click **Configure required settings** to create a static public IP address for ipconfig2.

By default, public IPs are dynamic. To make sure that the VM always uses the same public IP address, create a static Public IP.

In the Create public IP address blade, add a Name, under Assignment click **Static**. And then click **OK**.

Create public IP address

* Name
 ✓

Assignment
 Dynamic Static

Note:

Even when you set the allocation method to static, you cannot specify the actual IP address assigned to the public IP resource. Instead, it gets allocated from a pool of available IP addresses in the Azure location the resource is created in.

Follow the steps to add one more IP configuration for ipconfig3. Public IP is not mandatory.

Search IP configurations				
NAME	IP VERSION	TYPE	PRIVATE IP ADDRESS	PUBLIC IP ADDRESS
ipconfig1	IPv4	Primary	192.0.0.4 (Static)	13.78.187.150 (NSDoc0330VM-ip)
ipconfig2	IPv4	Secondary	192.0.0.5 (Static)	13.78.183.123 (ipconfig2_PIP2)
ipconfig3	IPv4	Secondary	192.0.0.6 (Static)	-

Step 2: Configure Citrix ADC-owned IP addresses

Configure the Citrix ADC-owned IP addresses by using the GUI or the command `add ns ip`. For more information, see [Configuring Citrix ADC-Owned IP Addresses](#).

Configure a high-availability setup with multiple IP addresses and NICs

September 9, 2024

In a Microsoft Azure deployment, a high-availability configuration of two Citrix ADC VPX instances is achieved by using the Azure Load Balancer (ALB). This is achieved by configuring a health probe on ALB, which monitors each VPX instance by sending a health probe at every 5 seconds to both primary and secondary instances.

In this setup, only the primary node responds to health probes and the secondary does not. Once the primary sends the response to the health probe, the ALB starts sending the data traffic to the instance. If the primary instance misses two consecutive health probes, ALB does not redirect traffic to that instance. On failover, the new primary starts responding to health probes and the ALB redirects traffic to it. The standard VPX high availability failover time is three seconds. The total failover time that might take for traffic switching can be a maximum of 13 seconds.

You can deploy a pair of Citrix ADC VPX instances with multiple NICs in an active-passive high availability (HA) setup on Azure. Each NIC can contain multiple IP addresses.

The following options are available for a multi-NIC high availability deployment:

- High availability using Azure availability set
- High availability using Azure availability zones

For more information about Azure Availability Set and Availability Zones, see the Azure documentation [Manage the availability of Linux virtual machines](#).

High availability using availability set

A high availability setup using a availability set must meet the following requirements:

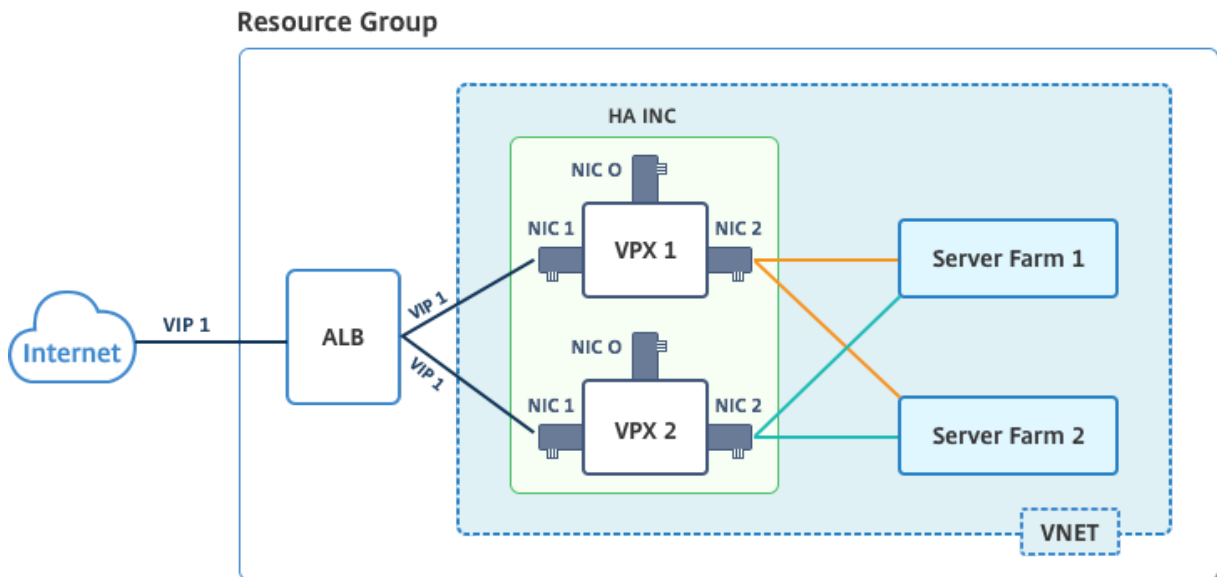
- An HA Independent Network Configuration (INC) configuration
- The Azure Load Balancer (ALB) in Direct Server Return (DSR) mode

All traffic goes through the primary node. The secondary node remains in standby mode until the primary node fails.

Note:

For a Citrix VPX high availability deployment on the Azure cloud to work, you need a floating public IP (PIP) that can be moved between the two VPX nodes. The Azure Load Balancer (ALB) provides that floating PIP, which is moved to the second node automatically in the event of a failover.

Diagram: Example of a high availability deployment architecture, using Azure Availability Set



In an active-passive deployment, the ALB front end public IP (PIP) addresses are added as the VIP addresses in each VPX node. In HA-INC configuration, the VIP addresses are floating and SNIP addresses are instance specific.

You can deploy a VPX pair in active-passive high availability mode in two ways by using:

- **Citrix ADC VPX standard high availability template:** use this option to configure an HA pair with the default option of three subnets and six NICs.
- **Windows PowerShell commands:** use this option to configure an HA pair according to your subnet and NIC requirements.

This topic describes how to deploy a VPX pair in active-passive HA setup by using the Citrix template. If you want to use PowerShell commands, see [Configuring an HA Setup with Multiple IP Addresses and NICs by Using PowerShell Commands](#).

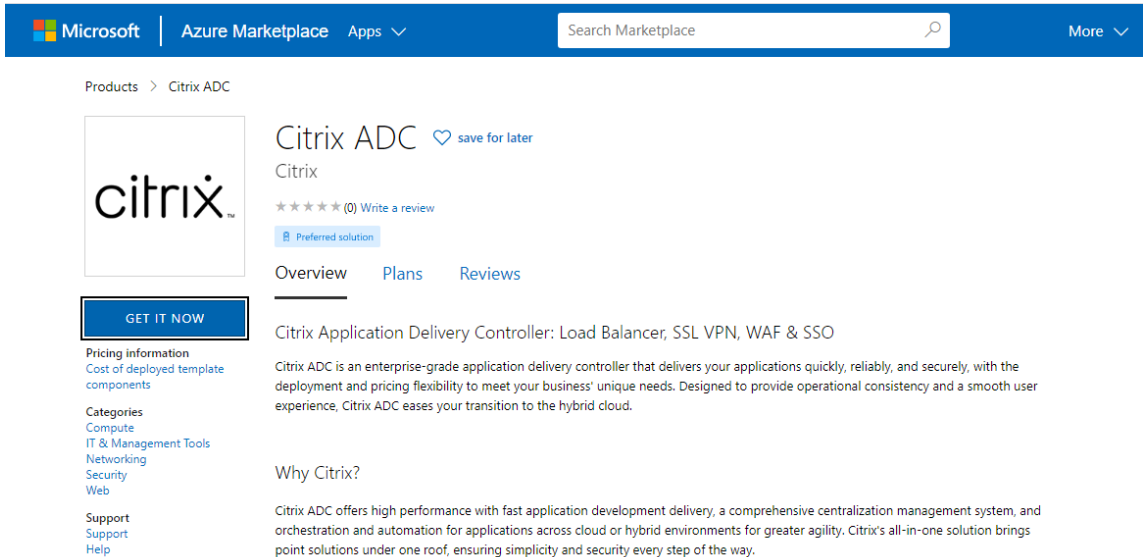
Configure HA-INC nodes by using the Citrix high availability template

You can quickly and efficiently deploy a pair of VPX instances in HA-INC mode by using the standard template. The template creates two nodes, with three subnets and six NICs. The subnets are for management, client, and server-side traffic, and each subnet has two NICs for both the VPX instances.

You can get the Citrix ADC HA Pair template at the [Azure Marketplace](#).

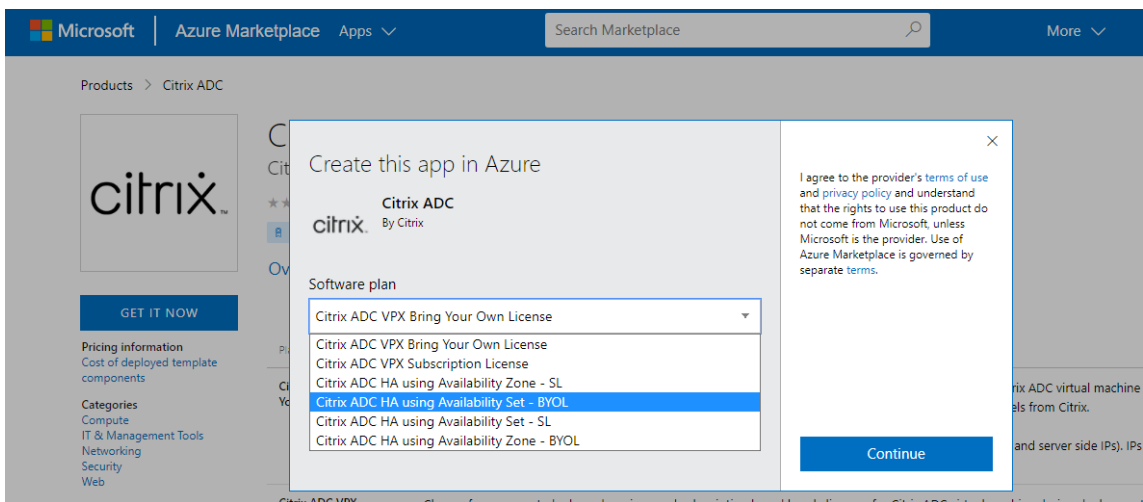
Complete the following steps to launch the template and deploy a high availability VPX pair, by using Azure availability sets.

1. From Azure Marketplace, search **Citrix ADC**.




2. Click **GET IT NOW**.

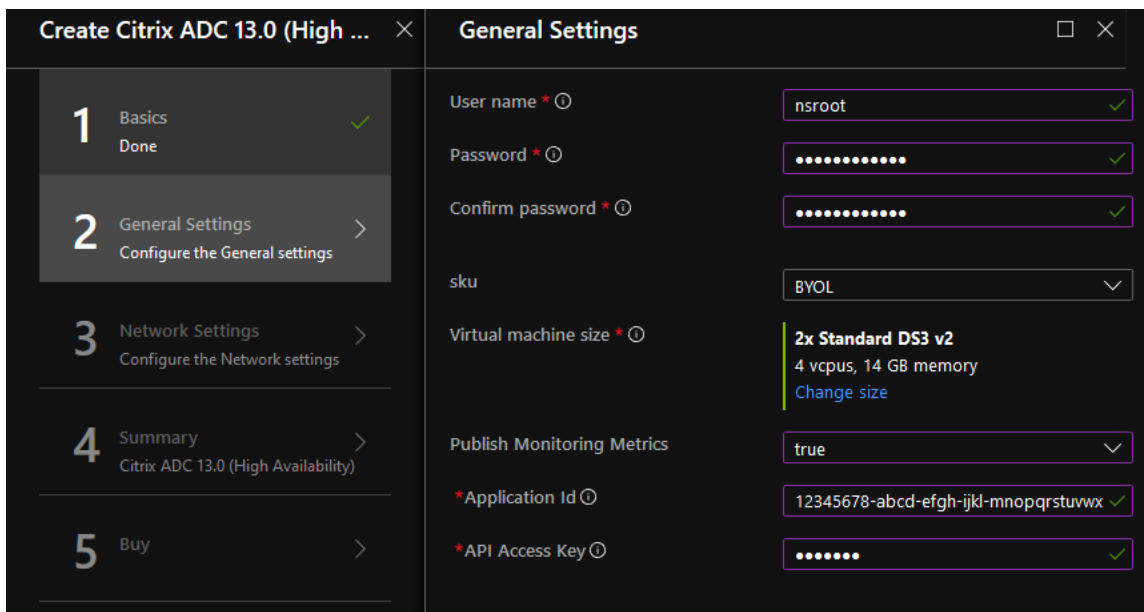
3. Select the required HA deployment along with license, and click **Continue**.



4. The **Basics** page appears. Create a Resource Group and select **OK**.



5. The **General Settings** page appears. Type the details and select **OK**.

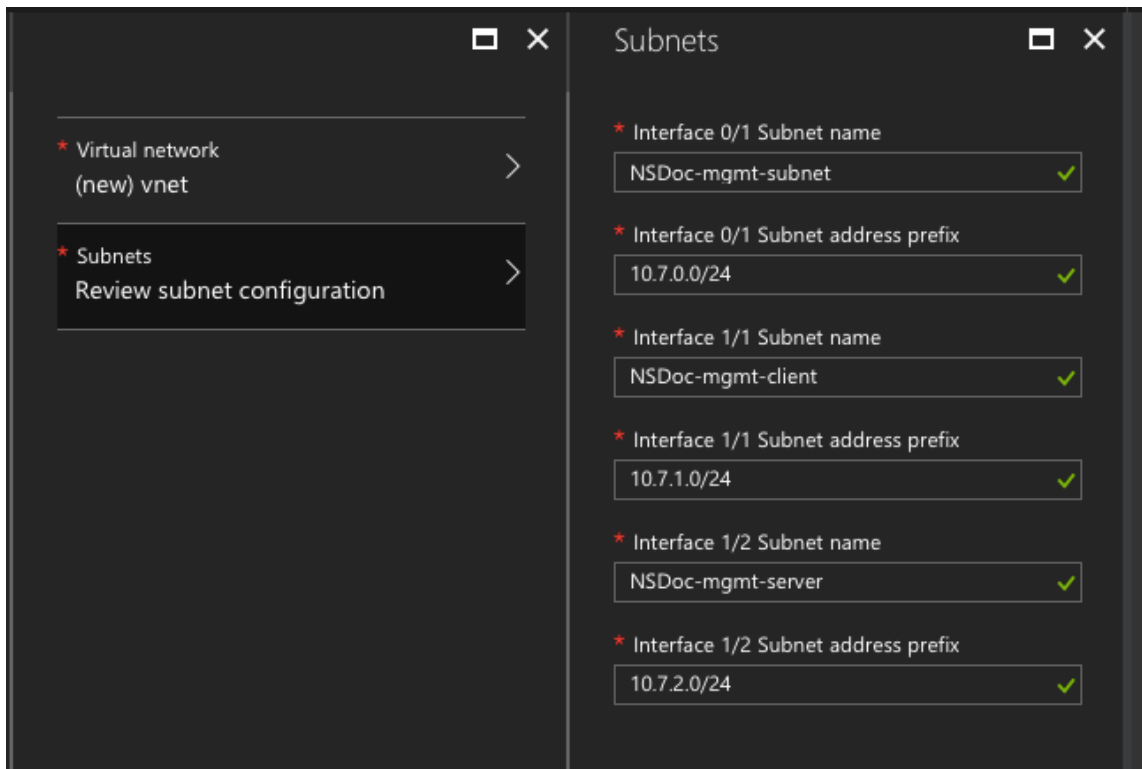


Note:

By default, the **Publishing Monitoring Metrics** option is set to **false**. If you want to enable this option, select **true**.

Create an Azure Active Directory (ADD) application and service principal that can access resources. Assign contributor role to the newly created AAD application. For more information, see [Use portal to create an Azure Active Directory application and service principal that can access resources](#).

- The **Network Settings** page appears. Check the VNet and subnet configurations, edit the required settings, and select **OK**.


























- The **Summary** page appears. Review the configuration and edit accordingly. Select **OK** to confirm.
- The **Buy** page appears. Select **Purchase** to complete the deployment.

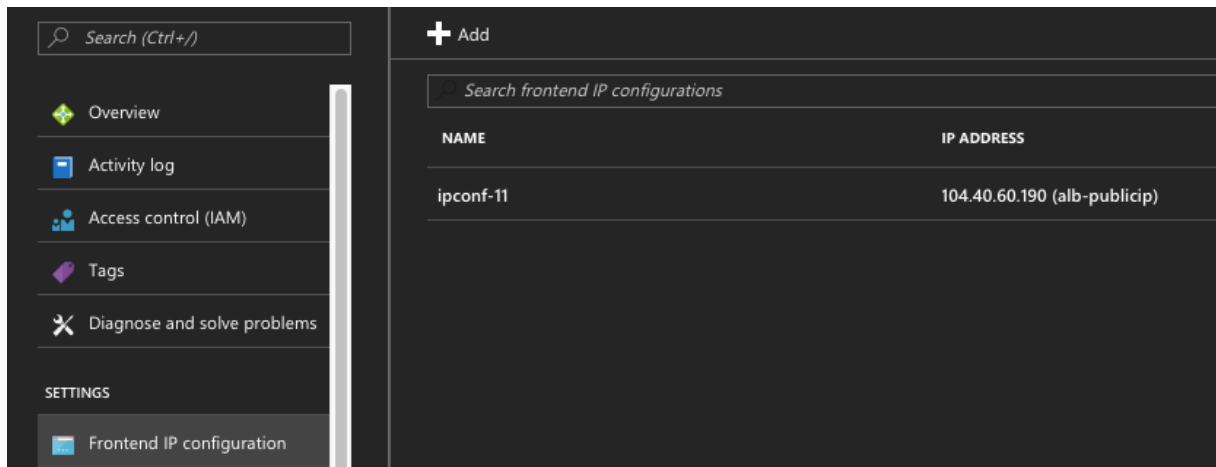
It might take a moment for the Azure Resource Group to be created with the required configurations. After completion, select the **Resource Group** in the Azure portal to see the configuration details, such as LB rules, back-end pools, health probes. The high availability pair appears as ns-vpx0 and ns-vpx1.

If further modifications are required for your HA setup, such as creating more security rules and ports, you can do that from the Azure portal.

23 items Show hidden types ⓘ

<input type="checkbox"/>	NAME ↑↓	TYPE ↑↓
<input type="checkbox"/>	 alb	Load balancer
<input type="checkbox"/>	 alb-publicip	Public IP address
<input type="checkbox"/>	 avl-set	Availability set
<input type="checkbox"/>	 ns-vpx0	Disk
<input type="checkbox"/>	 ns-vpx0	Virtual machine
<input type="checkbox"/>	 ns-vpx0-mgmt-publicip	Public IP address
<input type="checkbox"/>	 ns-vpx1	Disk
<input type="checkbox"/>	 ns-vpx1	Virtual machine
<input type="checkbox"/>	 ns-vpx1-mgmt-publicip	Public IP address
<input type="checkbox"/>	 ns-vpx-nic0-01	Network interface
<input type="checkbox"/>	 ns-vpx-nic0-11	Network interface
<input type="checkbox"/>	 ns-vpx-nic0-12	Network interface
<input type="checkbox"/>	 ns-vpx-nic1-01	Network interface
<input type="checkbox"/>	 ns-vpx-nic1-11	Network interface
<input type="checkbox"/>	 ns-vpx-nic1-12	Network interface
<input type="checkbox"/>	 ns-vpx-nic-nsg0-01	Network security group
<input type="checkbox"/>	 ns-vpx-nic-nsg0-11	Network security group
<input type="checkbox"/>	 ns-vpx-nic-nsg0-12	Network security group
<input type="checkbox"/>	 ns-vpx-nic-nsg1-01	Network security group
<input type="checkbox"/>	 ns-vpx-nic-nsg1-11	Network security group
<input type="checkbox"/>	 ns-vpx-nic-nsg1-12	Network security group
<input type="checkbox"/>	 vnet01	Virtual network
<input type="checkbox"/>	 vpxhamd7fi3wouvrk	Storage account

Next, you need to configure the load-balancing virtual server with the **ALB's Frontend public IP (PIP) address**, on primary node. To find the ALB PIP, select ALB > **Frontend IP configuration**.



See the **Resources** section for more information about how to configure the load-balancing virtual server.

Resources:

The following links provide additional information related to HA deployment and virtual server configuration:

- [Configuring high availability nodes in different subnets](#)
- [Set up basic load balancing](#)


Related resources:

- [Configure a high-availability setup with multiple IP addresses and NICs by using PowerShell commands](#)
- [Configuring GSLB on Active-Standby HA Deployment on Azure](#)

High availability using availability zones

Azure Availability Zones are fault-isolated locations within an Azure region, providing redundant power, cooling, and networking and increasing resiliency. Only specific Azure regions support Availability Zones. For more information, see the Azure documentation [What are Availability Zones in Azure].

Diagram: Example of a high availability deployment architecture, using Azure Availability Zones



You can deploy a VPX pair in high availability mode by using the template called “NetScaler 13.0 HA using Availability Zones,” available in Azure Marketplace.

Complete the following steps to launch the template and deploy a high availability VPX pair, by using Azure Availability Zones.


1. From Azure Marketplace, select and initiate the Citrix solution template.



2. Ensure deployment type is Resource Manager and select **Create**.
3. The **Basics** page appears. Enter the details and click **OK**.

Note:

Ensure that you select an Azure region that supports Availability Zones. For more information about regions that support Availability Zones, see Azure documentation [What are Availability Zones in Azure?](#)



4. The **General Settings** page appears. Type the details and select **OK**.
5. The **Network Setting** page appears. Check the VNet and subnet configurations, edit the required settings, and select **OK**.
6. The **Summary** page appears. Review the configuration and edit accordingly. Select **OK** to confirm.
7. The **Buy** page appears. Select **Purchase** to complete the deployment.

It might take a moment for the Azure Resource Group to be created with the required configurations. After completion, select the **Resource Group** to see the configuration details, such as LB rules, back-end pools, health probes, and so on, in the Azure portal. The high availability pair appears as ns-vpv0 and ns-vpv1. Also, you can see the location under the **Location** column.

Filter by name... All types All locations No grouping

22 items Show hidden types

NAME	TYPE	LOCATION
alb	Load balancer	East US 2
alb-publicip	Public IP address	East US 2
ns-vpx0	Virtual machine	East US 2
ns-vpx0_OsDisk_1_d7b757b8aa804bf1991a083f319e553a	Disk	East US 2
ns-vpx0-mgmt-publicip	Public IP address	East US 2
ns-vpx1	Virtual machine	East US 2
ns-vpx1_OsDisk_1_0c2364d43e2b47fa896bf14b02090ee0	Disk	East US 2
ns-vpx1-mgmt-publicip	Public IP address	East US 2
ns-vpx-nic0-01	Network interface	East US 2
ns-vpx-nic0-11	Network interface	East US 2
ns-vpx-nic0-12	Network interface	East US 2
ns-vpx-nic1-01	Network interface	East US 2
ns-vpx-nic1-11	Network interface	East US 2
ns-vpx-nic1-12	Network interface	East US 2
ns-vpx-nic-nsg0-01	Network security group	East US 2
ns-vpx-nic-nsg0-11	Network security group	East US 2
ns-vpx-nic-nsg0-12	Network security group	East US 2
ns-vpx-nic-nsg1-01	Network security group	East US 2
ns-vpx-nic-nsg1-11	Network security group	East US 2
ns-vpx-nic-nsg1-12	Network security group	East US 2
test1	Virtual network	East US 2
vpxhavdosvod3v5jeu	Storage account	East US 2

If further modifications are required for your HA setup, such as creating more security rules and ports, you can do that from the Azure portal.

Monitor your instances using metrics in Azure monitor

You can use metrics in the Azure monitor data platform to monitor a set of Citrix ADC VPX resources such as CPU, memory utilization, and throughput. Metrics service monitors Citrix ADC VPX resources that run on Azure, in real time. You can use **Metrics Explorer** to access the collected data. For more information, see [Azure Monitor Metrics overview](#).


Points to note

- If you deploy a Citrix ADC VPX instance on Azure by using the Azure Marketplace offer, Metrics service is disabled by default.
- The Metrics service is not supported in Azure CLI.
- Metrics are available for CPU (management and packet CPU usage), memory, and throughput (inbound and outbound).

How to view metrics in Azure monitor

To view metrics in the Azure monitor for your instance, perform these steps:

1. Log on to **Azure Portal > Virtual Machines**.
2. Select the virtual machine that is the Primary Node.
3. In the **Monitoring** section, click **Metrics**.
4. From the **Metric Namespace** drop-down menu, click **Citrix ADC**.
5. Under **All metrics** in **Metrics** drop-down menu, click the metrics you want to view.
6. Click **Add metric** to view another metric on the same chart. Use the Chart options to customize your chart.



Configure a high-availability setup with multiple IP addresses and NICs by using PowerShell commands

November 21, 2024

You can deploy a pair of Citrix ADC VPX instances with multiple NICs in an active-passive high availability (HA) setup on Azure. Each NIC can contain multiple IP addresses.

An active-passive deployment requires:

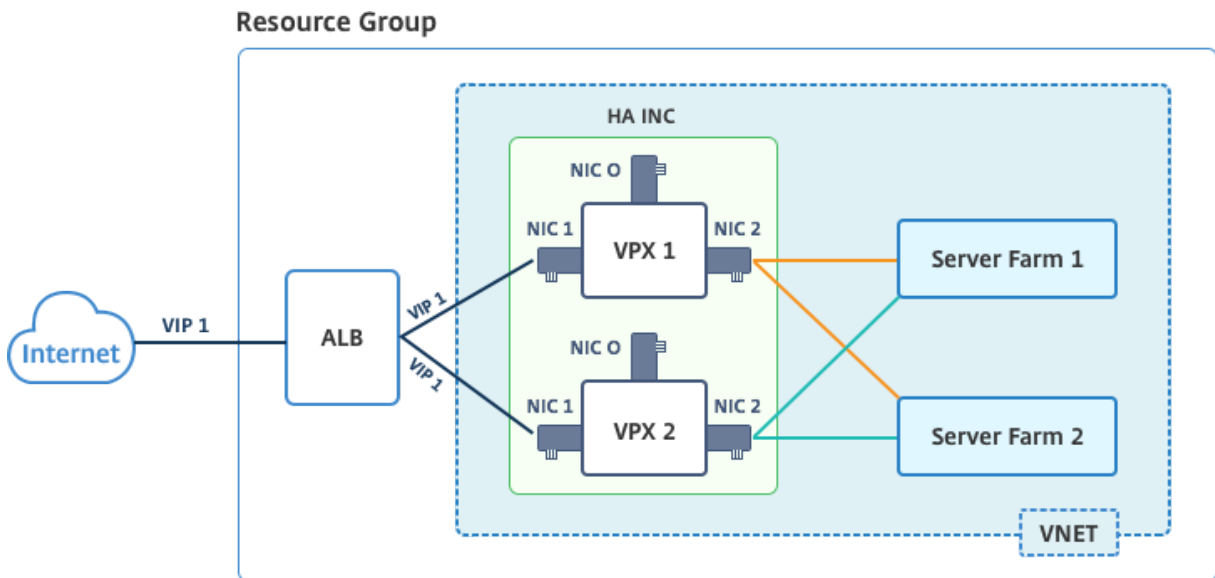
- An HA Independent Network Configuration (INC) configuration
- The Azure Load Balancer (ALB) in Direct Server Return (DSR) mode

All traffic goes through the primary node. The secondary node remains in standby mode until the primary node fails.

Note:

For a Citrix ADC VPX high availability deployment on an Azure cloud to work, you need a floating public IP (PIP) that can be moved between the two high-availability nodes. The Azure Load Balancer (ALB) provides that floating PIP, which is moved to the second node automatically in the event of a failover.

Diagram: Example of an active-passive deployment architecture



In an active-passive deployment, the ALB floating public IP (PIP) addresses are added as the VIP addresses in each VPX node. In HA-INC configuration, the VIP addresses are floating and SNIP addresses are instance specific.

ALB monitors each VPX instance by sending health probe at every 5 seconds and redirects traffic to that instance only that sends health probes response on regular interval. So in an HA setup, the primary node responds to health probes and secondary does not. If the primary instances miss two consecutive health probes, ALB does not redirect traffic to that instance. On failover, the new primary starts responding to health probes and the ALB redirects traffic to it. The standard VPX high availability failover time is three seconds. The total failover time that might take for traffic switching can be maximum of 13 seconds.

You can deploy a VPX pair in active-passive HA setup in two ways by using:

- **Citrix ADC VPX Standard high availability template:** use this option to configure an HA pair with the default option of three subnets and six NICs.
- **Windows PowerShell commands:** use this option to configure an HA pair according to your subnet and NIC requirements.

This topic describes how to deploy a VPX pair in active-passive HA setup by using PowerShell commands. If you want to use the Citrix ADC VPX Standard HA template, see [Configuring an HA Setup with Multiple IP Addresses and NICs](#).

Configure HA-INC nodes by using PowerShell Commands

Scenario: HA-INC PowerShell deployment

In this scenario, you deploy a Citrix ADC VPX pair by using the topology given in the table. Each VPX instance contains three NICs, with each NIC is deployed in a different subnet. Each NIC is assigned an IP configuration.

ALB	VPX1	VPX2
ALB is associated with public IP 3 (pip3)	Management IP is configured with IPConfig1, which includes one public IP (pip1) and one private IP (12.5.2.24); nic1; Mgmtsubnet=12.5.2.0/24	Management IP is configured with IPConfig5, which includes one public IP (pip3) and one private IP (12.5.2.26);nic4;Mgmtsubnet=12.5.2.0/24
LB rules and port configured are HTTP (80),SSL (443), health probe (9000)	Client-side IP is configured with IPConfig3, which includes one private IP(12.5.1.27);nic2; FrontEndsubnet=12.5.1.0/24	Client-side IP is configured with IPConfig7, which includes one private IP (12.5.1.28);nic5;FrontEndsubnet=12.5.1.0/24
-	Server-side IP is configured with IPConfig4, which includes one private IP(12.5.3.24); nic3;BackendSubnet=12.5.3.0/24	Server-side IP is configured with IPConfig8, which includes one private IP(12.5.3.28);nic6;BackendSubnet=12.5.3.0/24
-	Rules and ports for NSG are SSH (22),HTTP (80),HTTPS (443)	-

Parameter settings

The following parameter settings are used in this scenario:

```

1 $locName= "South east Asia"
2
3 $rgName = "MulitIP-MultiNIC-RG"
4
5 $nicName1= "VM1-NIC1"
6

```

```
7 $nicName2 = "VM1-NIC2"
8
9 $nicName3= "VM1-NIC3"
10
11 $nicName4 = "VM2-NIC1"
12
13 $nicName5= "VM2-NIC2"
14
15 $nicName6 = "VM2-NIC3"
16
17 $vNetName = "Azure-MultiIP-ALB-vnet"
18
19 $vNetAddressRange= "12.5.0.0/16"
20
21 $frontEndSubnetName= "frontEndSubnet"
22
23 $frontEndSubnetRange= "12.5.1.0/24"
24
25 $mgmtSubnetName= "mgmtSubnet"
26
27 $mgmtSubnetRange= "12.5.2.0/24"
28
29 $backEndSubnetName = "backEndSubnet"
30
31 $backEndSubnetRange = "12.5.3.0/24"
32
33 $prmStorageAccountName = "multiipmultinicbstorage"
34
35 $avSetName = "multiple-avSet"
36
37 $vmSize= "Standard\_DS4\_V2"
38
39 $publisher = "Citrix"
40
41 $offer = "netscalervpx-120"
42
43 $sku = "netscalerbyol"
44
45 $version="latest"
46
47 $pubIPName1="VPX1MGMT"
48
49 $pubIPName2="VPX2MGMT"
50
51 $pubIPName3="ALBPIP"
52
53 $domName1="vpx1dns"
54
55 $domName2="vpx2dns"
56
57 $domName3="vpxalbdns"
58
59 $vmNamePrefix="VPXMultiIPALB"
```

```
60
61 $osDiskSuffix1="osmultiipalbdiskdb1"
62
63 $osDiskSuffix2="osmultiipalbdiskdb2"
64
65 $lbName= "MultiIPALB"
66
67 $frontEndConfigName1= "FrontEndIP"
68
69 $backendPoolName1= "BackendPoolHttp"
70
71 $lbRuleName1= "LBRuleHttp"
72
73 $healthProbeName= "HealthProbe"
74
75 $nsgName="NSG-MultiIP-ALB"
76
77 $rule1Name="Inbound-HTTP"
78
79 $rule2Name="Inbound-HTTPS"
80
81 $rule3Name="Inbound-SSH"
```

To complete the deployment, complete the following steps by using PowerShell commands:

1. Create a resource group, storage account, and availability set
2. Create a network security group and add rules
3. Create a virtual network and three subnets
4. Create public IP addresses
5. Create IP configurations for VPX1
6. Create IP configurations for VPX2
7. Create NICs for VPX1
8. Create NICs for VPX2
9. Create VPX1
10. Create VPX2
11. Create ALB

Create a resource group, storage account, and availability set.

```
1 New-AzureRmResourceGroup -Name $rgName -Location $locName
2
3
4 $prmStorageAccount=New-AzureRMStorageAccount -Name
   $prmStorageAccountName -ResourceGroupName $rgName -Type Standard_LRS
   -Location $locName
5
6
7 $avSet=New-AzureRMAvailabilitySet -Name $avSetName -ResourceGroupName
   $rgName -Location $locName
```

Create a network security group and add rules.

```

1 $rule1 = New-AzureRmNetworkSecurityRuleConfig -Name $rule1Name -
  Description "Allow HTTP" -Access Allow -Protocol Tcp -Direction
  Inbound -Priority 101
2
3
4 -SourceAddressPrefix Internet -SourcePortRange * -
  DestinationAddressPrefix * -DestinationPortRange 80
5
6
7 $rule2 = New-AzureRmNetworkSecurityRuleConfig -Name $rule2Name -
  Description "Allow HTTPS" -Access Allow -Protocol Tcp -Direction
  Inbound -Priority 110
8
9
10 -SourceAddressPrefix Internet -SourcePortRange * -
  DestinationAddressPrefix * -DestinationPortRange 443
11
12
13 $rule3 = New-AzureRmNetworkSecurityRuleConfig -Name $rule3Name -
  Description "Allow SSH" -Access Allow -Protocol Tcp -Direction
  Inbound -Priority 120
14
15
16 -SourceAddressPrefix Internet -SourcePortRange * -
  DestinationAddressPrefix * -DestinationPortRange 22
17
18
19 $nsg = New-AzureRmNetworkSecurityGroup -ResourceGroupName $rgName -
  Location $locName -Name $nsgName -SecurityRules $rule1,$rule2,$rule3

```

Create a virtual network and three subnets.

```

1 $frontendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
  $frontEndSubnetName -AddressPrefix $frontEndSubnetRange (this
  parameter value should be as per your requirement)
2
3
4 $mgmtSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name $mgmtSubnetName
  -AddressPrefix $mgmtSubnetRange
5
6
7 $backendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
  $backEndSubnetName -AddressPrefix $backEndSubnetRange
8
9
10 $vnet =New-AzureRmVirtualNetwork -Name $vNetName -ResourceGroupName
  $rgName -Location $locName -AddressPrefix $vNetAddressRange -Subnet
  $frontendSubnet,$backendSubnet, $mgmtSubnet
11
12
13 $subnetName ="frontEndSubnet"
14

```

```
15
16 \${subnet1}=\${vnet.Subnets}|?{
17   \${\}.Name -eq \${subnetName} }
18
19
20
21 \${subnetName}="backEndSubnet"
22
23
24 \${subnet2}=\${vnet.Subnets}|?{
25   \${\}.Name -eq \${subnetName} }
26
27
28
29 \${subnetName}="mgmtSubnet"
30
31
32 \${subnet3}=\${vnet.Subnets}|?{
33   \${\}.Name -eq \${subnetName} }
```

Create public IP addresses.

```
1 $pip1=New-AzureRmPublicIpAddress -Name $pubIPName1 -ResourceGroupName
   $rgName -DomainNameLabel $domName1 -Location $locName -
   AllocationMethod Dynamic
2
3 $pip2=New-AzureRmPublicIpAddress -Name $pubIPName2 -ResourceGroupName
   $rgName -DomainNameLabel $domName2 -Location $locName -
   AllocationMethod Dynamic
4
5 $pip3=New-AzureRmPublicIpAddress -Name $pubIPName3 -ResourceGroupName
   $rgName -DomainNameLabel $domName3 -Location $locName -
   AllocationMethod Dynamic
```

Create IP configurations for VPX1.

```
1 $IPConfigName1 = "IPConfig1"
2
3
4 $IPAddress = "12.5.2.24"
5
6
7 $IPConfig1=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName1 -
   Subnet $subnet3 -PrivateIpAddress $IPAddress -PublicIpAddress $pip1
   -Primary
8
9
10 $IPConfigName3="IPConfig-3"
11
12
13 $IPAddress="12.5.1.27"
14
15
```

```
16 $IPConfig3=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName3 -
    Subnet $subnet1 -PrivateIpAddress $IPAddress -Primary
17
18
19 $IPConfigName4 = "IPConfig-4"
20
21
22 $IPAddress = "12.5.3.24"
23
24
25 $IPConfig4 = New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName4 -
    Subnet $subnet2 -PrivateIpAddress $IPAddress -Primary
```

Create IP configurations for VPX2.

```
1 $IpConfigName5 = "IPConfig5"
2
3
4 $IPAddress="12.5.2.26"
5
6
7 $IPConfig5=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName5 -
    Subnet $subnet3 -PrivateIpAddress $IPAddress -PublicIpAddress $pip2
    -Primary
8
9
10 $IPConfigName7="IPConfig-7"
11
12
13 $IPAddress="12.5.1.28"
14
15
16 $IPConfig7=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName7 -
    Subnet $subnet1 -PrivateIpAddress $IPAddress -Primary
17
18
19 $IPConfigName8="IPConfig-8"
20
21
22 $IPAddress="12.5.3.28"
23
24
25 $IPConfig8=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName8 -
    Subnet $subnet2 -PrivateIpAddress $IPAddress -Primary
```

Create NICs for VPX1.

```
1 $nic1=New-AzureRmNetworkInterface -Name $nicName1 -ResourceGroupName
    $rgName -Location $locName -IpConfiguration $IpConfig1 -
    NetworkSecurityGroupId $nsg.Id
2
3
4 $nic2=New-AzureRmNetworkInterface -Name $nicName2 -ResourceGroupName
```

```

    $rgName -Location $locName -IpConfiguration $IpConfig3 -
    NetworkSecurityGroupId $nsg.Id
5
6
7 $nic3=New-AzureRmNetworkInterface -Name $nicName3 -ResourceGroupName
    $rgName -Location $locName -IpConfiguration $IpConfig4 -
    NetworkSecurityGroupId $nsg.Id

```

Create NICs for VPX2.

```

1 $nic4=New-AzureRmNetworkInterface -Name $nicName4 -ResourceGroupName
    $rgName -Location $locName -IpConfiguration $IpConfig5 -
    NetworkSecurityGroupId $nsg.Id
2
3
4 $nic5=New-AzureRmNetworkInterface -Name $nicName5 -ResourceGroupName
    $rgName -Location $locName -IpConfiguration $IpConfig7 -
    NetworkSecurityGroupId $nsg.Id
5
6
7 $nic6=New-AzureRmNetworkInterface -Name $nicName6 -ResourceGroupName
    $rgName -Location $locName -IpConfiguration $IpConfig8 -
    NetworkSecurityGroupId $nsg.Id

```

Create VPX1.

This step includes the following substeps:

- Create VM config object
- Set credentials, OS, and image
- Add NICs
- Specify OS disk and create VM

```

1 $suffixNumber = 1
2
3 $vmName=$vmNamePrefix + $suffixNumber
4
5 $vmConfig=New-AzureRMVMConfig -VMName $vmName -VMSize $vmSize -
    AvailabilitySetId $avSet.Id
6
7 $cred=Get-Credential -Message "Type the name and password for VPX
    login."
8
9 $vmConfig=Set-AzureRMVMOperatingSystem -VM $vmConfig -Linux -
    ComputerName $vmName -Credential $cred
10
11 $vmConfig=Set-AzureRMVMSourceImage -VM $vmConfig -PublisherName
    $publisher -Offer $offer -Skus $sku -Version $version
12
13 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic1.
    Id -Primary

```

```

14
15 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic2.
    Id
16
17 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic3.
    Id
18
19 $osDiskName=$vmName + "-" + $osDiskSuffix1
20
21 $osVhdUri=$prmStorageAccount.PrimaryEndpoints.Blob.ToString() + "
    vhd/" + $osDiskName + ".vhd"
22
23 $vmConfig=Set-AzureRMVMOSDisk -VM $vmConfig -Name $osDiskName -
    VhdUri $osVhdUri -CreateOption fromImage
24
25 Set-AzureRmVMPlan -VM $vmConfig -Publisher $publisher -Product
    $offer -Name $sku
26
27 New-AzureRMVM -VM $vmConfig -ResourceGroupName $rgName -Location
    $locName

```

Create VPX2.

```

1  ``
2  $suffixNumber=2
3
4
5  $vmName=$vmNamePrefix + $suffixNumber
6
7
8  $vmConfig=New-AzureRMVMConfig -VMName $vmName -VMSize $vmSize -
    AvailabilitySetId $avSet.Id
9
10
11 $cred=Get-Credential -Message "Type the name and password for VPX login
    ."
12
13
14 $vmConfig=Set-AzureRMVMOperatingSystem -VM $vmConfig -Linux -
    ComputerName $vmName -Credential $cred
15
16
17 $vmConfig=Set-AzureRMVMSourceImage -VM $vmConfig -PublisherName
    $publisher -Offer $offer -Skus $sku -Version $version
18
19
20 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic4.Id -
    Primary
21
22
23 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic5.Id
24
25

```



```
26 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic6.Id
27
28
29 $osDiskName=$vmName + "-" + $osDiskSuffix2
30
31
32 $osVhdUri=$prmStorageAccount.PrimaryEndpoints.Blob.ToString() + "vhds/"
    + $osDiskName + ".vhd"
33
34
35 $vmConfig=Set-AzureRMVMOSDisk -VM $vmConfig -Name $osDiskName -VhdUri
    $osVhdUri -CreateOption fromImage
36
37
38 Set-AzureRmVMPlan -VM $vmConfig -Publisher $publisher -Product $offer -
    Name $sku
39
40
41 New-AzureRMVM -VM $vmConfig -ResourceGroupName $rgName -Location
    $locName
42 ```
```

To view private and public IP addresses assigned to the NICs, type the following commands:

```
1 ```
2 $nic1.IPConfig
3
4
5 $nic2.IPConfig
6
7
8 $nic3.IPConfig
9
10
11 $nic4.IPConfig
12
13
14 $nic5.IPConfig
15
16
17 $nic6.IPConfig
18 ```
```

Create Azure load balance (ALB).

This step includes the following substeps:

- Create front end IP config
- Create health probe
- Create back end address pool
- Create load-balancing rules (HTTP and SSL)

- Create ALB with front end IP config, back end address pool, and LB rule
- Associate IP config with back end pools

```
$frontEndIP1=New-AzureRmLoadBalancerFrontendIpConfig -Name
$frontEndConfigName1 -PublicIpAddress $pip3

$healthProbe=New-AzureRmLoadBalancerProbeConfig -Name $healthProbeName
-Protocol Tcp -Port 9000 -IntervalInSeconds 5 -ProbeCount 2

$beAddressPool1=New-AzureRmLoadBalancerBackendAddressPoolConfig -
Name $backendPoolName1

$lbRule1=New-AzureRmLoadBalancerRuleConfig -Name $lbRuleName1
-FrontendIpConfiguration $frontEndIP1 -BackendAddressPool
$beAddressPool1 -Probe $healthProbe -Protocol Tcp -FrontendPort
80 -BackendPort 80 -EnableFloatingIP

$lb=New-AzureRmLoadBalancer -ResourceGroupName $rgName -Name
$lbName -Location $locName -FrontendIpConfiguration $frontEndIP1
-LoadBalancingRule $lbRule1 -BackendAddressPool $beAddressPool1 -
Probe $healthProbe

$nic2.IpConfigurations[0].LoadBalancerBackendAddressPools.Add($lb
.BackendAddressPools[0])

$nic5.IpConfigurations[0].LoadBalancerBackendAddressPools.Add($lb
.BackendAddressPools[0])

$lb=$lb | Set-AzureRmLoadBalancer

$nic2=$nic2 | Set-AzureRmNetworkInterface

$nic5=$nic5 | Set-AzureRmNetworkInterface
```

After you've successfully deployed the Citrix ADC VPX pair, log on to each VPX instance to configure HA-INC, and SNIP and VIP addresses.

1. Type the following command to add HA nodes.

```
add ha node 1 PeerNodeNSIP -inc Enabled
```

2. Add private IP addresses of client-side NICs as SNIPs for VPX1 (NIC2) and VPX2 (NIC5)

```
add nsip privateIPofNIC2 255.255.255.0 -type SNIP
add nsip privateIPofNIC5 255.255.255.0 -type SNIP
```

3. Add load-balancing virtual server on the primary node with front-end IP address (public IP) of ALB.

```
add lb virtual server v1 HTTP FrontEndIPofALB 80
```

Related resources:

[Configuring GSLB on Active-Standby HA Deployment on Azure](#)

Configure a Citrix ADC VPX instance to use Azure accelerated networking

September 12, 2024

Accelerated networking enables the single root I/O virtualization (SR-IOV) virtual function (VF) NIC to a virtual machine, which improves the networking performance. You can use this feature with heavy workloads that need to send or receive data at higher throughput with reliable streaming and lower CPU utilization.

When a NIC is enabled with accelerated networking, Azure bundles the NIC's existing para virtualized (PV) interface with an SR-IOV VF interface. The support of SR-IOV VF interface enables and enhances the throughput of the Citrix ADC VPX instance.

Accelerated networking provides the following benefits:

- Lower latency
- Higher packets per second (pps) performance
- Enhanced throughput
- Reduced jitter
- Decreased CPU utilization

Note:

Azure accelerated networking is supported on Citrix ADC VPX instances from release 13.0 build 76.29 onwards.

Prerequisites

- Ensure that your VM size matches the requirements for Azure accelerated networking.
- Stop VMs (individual or in an availability set) before enabling accelerated networking on any NIC.

Limitations

Accelerated networking can be enabled only on some instance types. For more information, see [Supported instance types](#).

NICs supported for accelerated networking

Azure provides Mellanox ConnectX3 and ConnectX4 NICs in the SR-IOV mode for accelerated networking.

When accelerated networking is enabled on a Citrix ADC VPX interface, Azure bundles either ConnectX3 or ConnectX4 interface with the existing PV interface of a Citrix ADC VPX appliance.

For more information about enabling accelerated networking before attaching an interface to a VM, see [Create a network interface with accelerated networking](#).


For more information about enabling accelerated networking on an existing interface on a VM, see [Enable existing interfaces on a VM](#).

How to enable accelerated networking on Citrix ADC VPX instance using the Azure console


You can enable accelerated networking on a specific interface using the Azure console or the Azure PowerShell.

Do the following steps to enable accelerated networking by using Azure availability sets or availability zones.

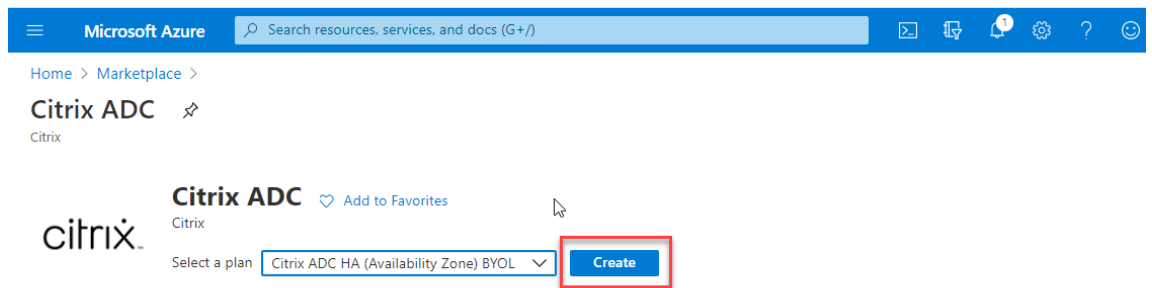
1. Log in to [Azure portal](#), and navigate to **Azure Marketplace**.



2. From the **Azure Marketplace**, search **Citrix ADC**.



3. Select a non-FIPS Citrix ADC plan along with license, and click **Create**.



The **Create Citrix ADC** page appears.

4. In the **Basics** tab, create a Resource Group. Under the **Parameters** tab, enter details for the Region, Admin user name, Admin Password, license type (VM SKU), and other fields.

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

Home > Citrix ADC >

Create Citrix ADC

Basics VM Configurations Network and Additional Settings Review + create

Project details

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

Subscription * ⓘ NSDev Platform CA

Resource group * ⓘ (New) test-aan-new
[Create new](#)

Instance details

Region * ⓘ South India

Citrix ADC Release Version * ⓘ
 12.1
 13.0

License Subscription Model * ⓘ
 10 Mbps
 200 Mbps
 1000 Mbps
 3000 Mbps

License Subscription Edition * ⓘ
 Standard
 Enterprise
 Platinum

Virtual Machine name * ⓘ citrix-adc-vpx

Administrator account

Username * ⓘ ✓

Authentication type * ⓘ
 Password
 SSH Public Key

Password * ⓘ ✓

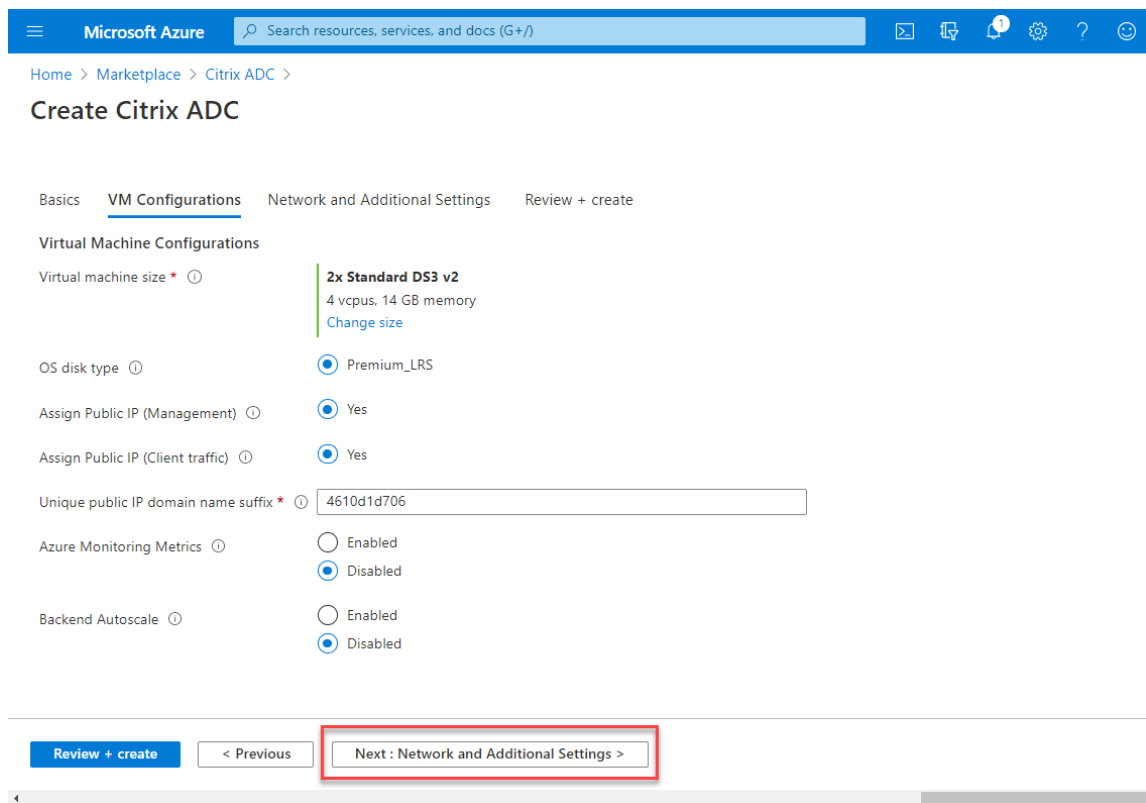
Confirm password * ⓘ ✓

[Review + create](#) < Previous **Next : VM Configurations >**

5. Click **Next : VM Configurations >**.

On the **VM Configurations** page, perform the following:

- a) Configure public IP domain name suffix.
- b) Enable or disable **Azure Monitoring Metrics**.
- c) Enable or disable **Backend Autoscale**.



6. Click **Next: Network and Additional settings >**.

On the **Network and Additional Settings** page, create a Boot diagnostics account and configure the network settings.

Under the **Accelerated Networking** section, you have the option to enable or disable the accelerated networking separately for the Management interface, Client interface, and Server interface.

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

Home > Marketplace > Citrix ADC >

Create Citrix ADC

Basics VM Configurations **Network and Additional Settings** Review + create

Boot diagnostics

Diagnostic storage account * ⓘ (new) citrixadcvp4610d1d706 [Create New](#)

Network Settings

Configure virtual networks

Virtual network * ⓘ (new) citrix-adc-vpx-virtual-network [Create new](#)

Management Subnet * ⓘ (new) 01-management-subnet (172.17.40.0/24) [Create new](#)

Client Subnet * ⓘ (new) 11-client-subnet (172.17.41.0/24) [Create new](#)

Server Subnet * ⓘ (new) 12-server-subnet (172.17.42.0/24) [Create new](#)

Accelerated Networking

Accelerated Networking (Management Interface) ⓘ On Off

Accelerated Networking (Client Interface) ⓘ On Off

Accelerated Networking (Server Interface) ⓘ On Off

VM 1 of HA Pair -> Public IP (Management)

Management Public IP (NSIP) of VM 1 * ⓘ (new) citrix-adc-vpx-nsip-0 [Create new](#)

Management Domain Name of VM 1 ⓘ citrix-adc-vpx-nsip-0-4610d1d706 ✓
.southindia.cloudapp.azure.com

VM 2 of HA Pair -> Public IP (Management)

Management Public IP (NSIP) of VM 2 * ⓘ (new) citrix-adc-vpx-nsip-1 [Create new](#)

Management Domain Name of VM 2 ⓘ citrix-adc-vpx-nsip-1-4610d1d706 ✓
.southindia.cloudapp.azure.com

Public IP (Client-side)

Client-side Public IP (VIP) * ⓘ (new) citrix-adc-vpx-vip [Create new](#)

Client-side Domain Name ⓘ citrix-adc-vpx-vip-4610d1d706 ✓
.southindia.cloudapp.azure.com

Public Inbound Ports (Management only)

Ports open for Management public IP ⓘ None ssh (22) ssh (22), http (80), https (443)

[Review + create](#) < Previous **Next : Review + create >**

7. Click **Next: Review + create >**.

After the validation is successful, review the basic settings, VM configurations, network and additional settings, and click **Create**. It might take some time for the Azure Resource Group to be created with the required configurations.

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

Home > Marketplace > Citrix ADC >

Create Citrix ADC

Validation Passed

Basics VM Configurations Network and Additional Settings Review + create

PRODUCT DETAILS

Citrix ADC
by Citrix
[Terms of use](#) | [Privacy policy](#)

TERMS

By clicking "Create", I (a) agree to the legal terms and privacy statement(s) associated with the Marketplace offering(s) listed above; (b) authorize Microsoft to bill my current payment method for the fees associated with the offering(s), with the same billing frequency as my Azure subscription; and (c) agree that Microsoft may share my contact, usage and transactional information with the provider(s) of the offering(s) for support, billing and other transactional activities. Microsoft does not provide rights for third-party offerings. See the [Azure Marketplace Terms](#) for additional details.

Basics


Subscription	NSDev Platform CA
Resource group	test-aan
Region	South Central US
Citrix ADC Release Version	13.0
License Subscription	Bring Your Own License
Virtual Machine name prefix	citrix-adc-vpx
Username	
Password (Use a domain name suffix)	*****
Azure Monitoring Metrics	Disabled
Backend Autoscale	Disabled

Network and Additional Settings


Diagnostic storage account	citrixadcpx4610d1d706
Virtual network	citrix-adc-vpx-virtual-network
Management Subnet	01-management-subnet
Address prefix (Management Subnet)	172.17.40.0/24
Client Subnet	11-client-subnet
Address prefix (Client Subnet)	172.17.41.0/24
Server Subnet	12-server-subnet
Address prefix (Server Subnet)	172.17.42.0/24
Accelerated Networking (Management Interface)	On
Accelerated Networking (Client Interface)	On
Accelerated Networking (Server Interface)	On
Public IP address	citrix-adc-vpx-nsip-0
Domain name label	citrix-adc-vpx-nsip-0-4610d1d706
Public IP address	citrix-adc-vpx-nsip-1
Domain name label	citrix-adc-vpx-nsip-1-4610d1d706
Public IP address	citrix-adc-vpx-vip
Domain name label	citrix-adc-vpx-vip-4610d1d706
Ports open for Management public IP	ssh (22)

Create < Previous Next Download a template for automation

8. After the deployment is complete, select the **Resource Group** to see the configuration details.



- To verify the Accelerated Networking configurations, select **Virtual machine > Networking**. The Accelerated Networking status is displayed as **Enabled** or **Disabled** for each NIC.



Enable accelerated networking using Azure PowerShell

If you need to enable accelerated networking after the VM creation, you can do so using Azure PowerShell.

Note:

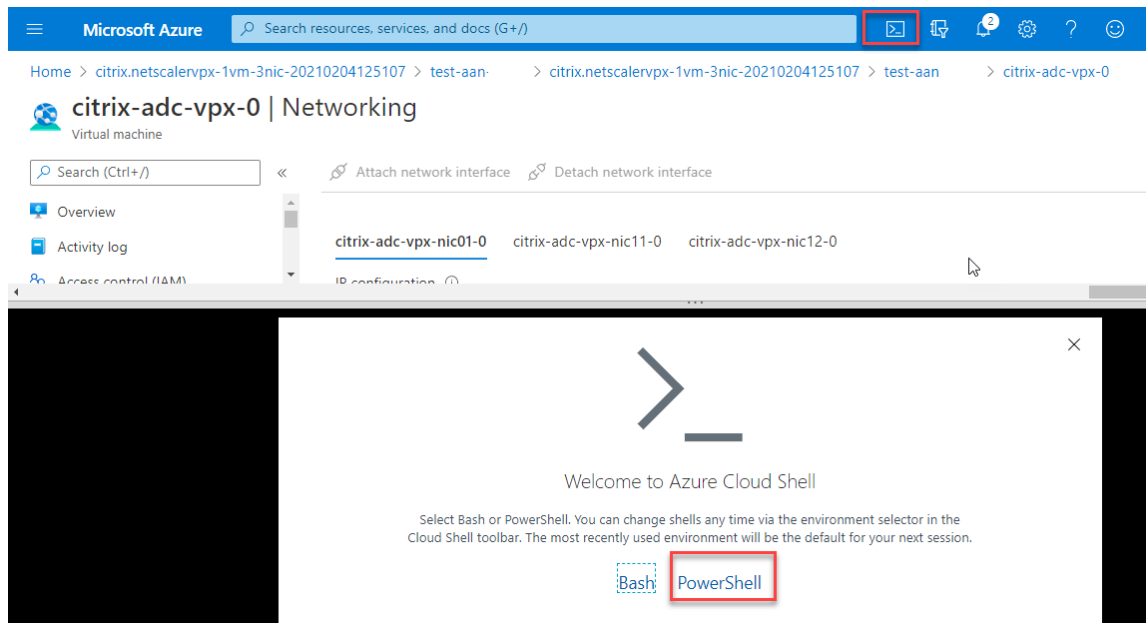
Ensure to stop the VM before you enable Accelerated Networking using Azure PowerShell.

Perform the following steps to enable accelerated networking by using Azure PowerShell.

1. Navigate to **Azure portal**, click the **PowerShell** icon on the right-hand top corner.

Note:

If you are in the Bash mode, change to the PowerShell mode.



2. At the command prompt, run the following command:

```
1 az network nic update --name <nic-name> --accelerated-networking [
  true | false] --resource-group <resourcegroup-name>
```

The accelerated networking parameter accepts either of the following values:

- **True:** Enables accelerated networking on the specified NIC.
- **False:** Disables accelerated networking on the specified NIC.

To enable accelerated networking on a specific NIC:


```
1 az network nic update --name citrix-adc-vpx-nic01-0 --accelerated-
  networking true --resource-group rsgp1-aan
```

To disable accelerated networking on a specific NIC:

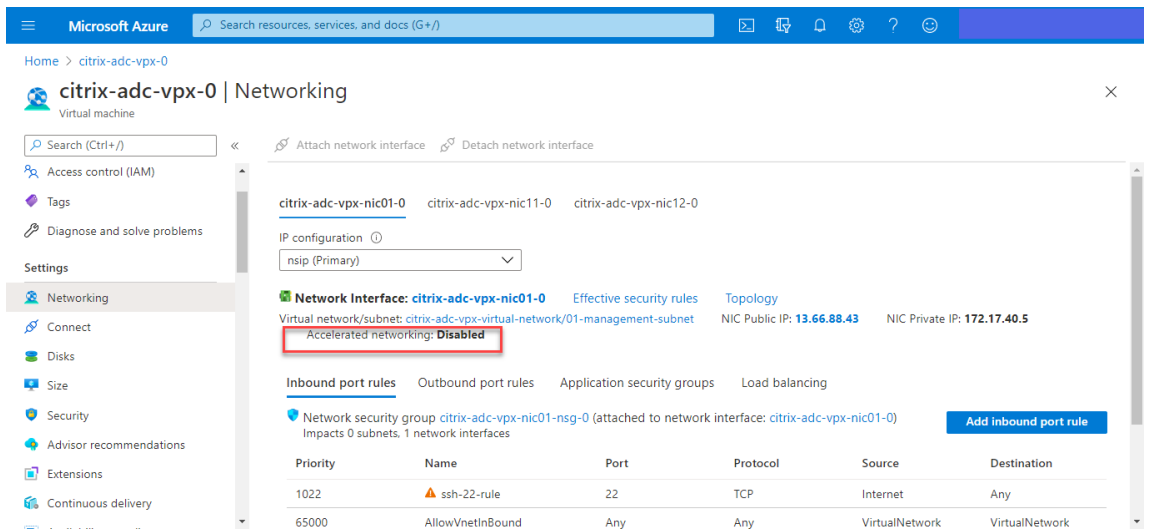
```
1 az network nic update --name citrix-adc-vpx-nic01-0 --accelerated-
  networking false --resource-group rsgp1-aan
```

3. To verify the Accelerated Networking status after the deployment is completed, Navigate to **VM > Networking**.

In the following example, you can see that Accelerated Networking is **Enabled**.



In the following example, you can see that Accelerated Networking is **Disabled**.



To verify accelerated networking on an interface by using FreeBSD Shell of Citrix ADC

You can log in to FreeBSD shell of Citrix ADC, and run the following commands to verify the accelerated networking status.

Example for ConnectX3 NIC:

The following example shows the “ifconfig” command output of the Mellanox ConnectX3 NIC. The “50/n” indicates the VF interfaces of the Mellanox ConnectX3 NICs. 0/1 and 1/1 indicates the PV interfaces of the Citrix ADC VPX instance. You can observe that both PV interface (1/1) and CX3 VF interface (50/1) have the same MAC addresses (00:22:48:1c:99:3e). This indicates that the two interfaces are bundled together.

```

root@nvr-us-cx3# ifconfig
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> metric 0 mtu 1500
    options=3<RXCSUM,TXCSUM>
    inet 127.0.0.1 netmask 0xff000000
    inet6 ::1 prefixlen 128
    inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1
    nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
0/1: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu 1500
    options=80019<RXCSUM,VLAN_MTU,VLAN_HWTAGGING,LINKSTATE>
    ether 00:0d:3a:98:71:be
    inet 172.16.27.11 netmask 0xfffff00 broadcast 172.16.27.255
    inet6 fe80::20d:3aff:fe98:71be%0/1 prefixlen 64 autoconf scopeid 0x2
    nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
    media: Ethernet autoselect (10Gbase-T <full-duplex>)
    status: active
1/1: flags=8802<BROADCAST,SIMPLEX,MULTICAST> metric 0 mtu 1500
    options=80019<RXCSUM,VLAN_MTU,VLAN_HWTAGGING,LINKSTATE>
    ether 00:22:48:1c:99:3e
    media: Ethernet autoselect (10Gbase-T <full-duplex>)
    status: active
50/1: flags=8842<BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu 1500
    options=900b8<VLAN_MTU,VLAN_HWTAGGING,JUMBO_MTU,VLAN_HWCSUM,VLAN_HWFILTER,LINKSTATE>
    ether 00:22:48:1c:99:3e
    media: Ethernet autoselect (<unknown subtype>)
    status: active

```

Example for ConnectX4 NIC:

The following example shows the “ifconfig” command output of the Mellanox ConnectX4 NIC. The “100/n” indicates the VF interfaces of the Mellanox ConnectX4 NICs. 0/1, 1/1, and 1/2 indicates the PV interfaces of Citrix ADC VPX instance.

You can observe that both PV interface (1/1) and CX4 VF interface (100/1) have the same MAC addresses (00:0d:3a:9b:f2:1d). This indicates that the two interfaces are bundled together. Similarly, the PV interface (1/2) and CX4 VF interface (100/2) have the same MAC addresses (00:0d:3a:1e:d2:23).

```

root@SmartNIC-CX4-NS-DUT-NEW1# ifconfig
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> metric 0 mtu 1500
options=3<RXCSUM,TXCSUM>
inet 127.0.0.1 netmask 0xff000000
inet6 ::1 prefixlen 128
inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1
nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
1/1: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu 1500
options=80019<RXCSUM,VLAN_MTU,VLAN_HWTAGGING,LINKSTATE>
ether 00:0d:3a:9b:f2:1d
inet 10.0.1.29 netmask 0xfffff00 broadcast 10.0.1.255
inet6 fe80::20d:3aff:fe9b:f21d%0/1 prefixlen 64 scopeid 0x2
nd6 options=3<PERFORMNUD,ACCEPT_RTADV>
media: Ethernet autoselect (10Gbase-T <full-duplex>)
status: active

1/2: flags=8802<BROADCAST,SIMPLEX,MULTICAST> metric 0 mtu 1500
options=80019<RXCSUM,VLAN_MTU,VLAN_HWTAGGING,LINKSTATE>
ether 00:0d:3a:1e:d2:23
media: Ethernet autoselect (10Gbase-T <full-duplex>)
status: active

100/1: flags=8a03<UP,BROADCAST,ALLMULTI,SIMPLEX,MULTICAST> metric 0 mtu 1500
ether 00:0d:3a:9b:f2:1d
media: Ethernet autoselect <full-duplex rxpause txpause> (autoselect
<full-duplex rxpause>)
status: active

100/2: flags=8a03<UP,BROADCAST,ALLMULTI,SIMPLEX,MULTICAST> metric 0 mtu 1500
ether 00:0d:3a:1e:d2:23
media: Ethernet autoselect <full-duplex rxpause txpause> (autoselect
<full-duplex rxpause>)
status: active

```

To verify accelerated networking on an interface by using ADC CLI

Example for ConnectX3 NIC:

The following show interface command output indicates that the PV interface 1/1 is bundled with virtual function 50/1, which is an SR-IOV VF NIC. The MAC addresses of both 1/1 and 50/1 NICs are the same. After accelerated networking is enabled, the data of the 1/1 interface is sent through datapath of the 50/1 interface, which is a ConnectX3 interface. You can see that the “show interface” output of PV interface (1/1) points to the VF (50/1). Similarly, the “show interface” output of VF interface (50/1) points to the PV interface (1/1).

```

> show interface 1/1
Interface 1/1 (NetScaler Virtual Interface, SmartNIC, VF 50/1 Datapath 50/1) #1
Flags=0xe060 <ENABLED, UP, UP, HEARTBEAT, 802.1q>
MTU=1500, native vlan=1, MAC=00:22:48:1c:99:3e, uptime 0h00m07s
LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
NIC: InDisc(0) OutDisc(0) FcTls(0) Stalls(0) Hangs(0) Muted(0)
Bandwidth thresholds are not set.

Done

> show interface 50/1
Interface 50/1 (CX3 VF Interface, SmartNIC, PV 1/1) #2
Flags=0xe400 <ENABLED, UP, UP, 802.1q>
MTU=1500, native vlan=1, MAC=00:22:48:1c:99:3e, uptime 0h00m08s
Actual: media NONE, speed 50000, duplex FULL, FcTl NONE, throughput 50000
LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
NIC: InDisc(0) OutDisc(0) FcTls(0) Stalls(0) Hangs(0) Muted(0)
Bandwidth thresholds are not set.

```

Example for ConnectX4 NIC:

The following show interface command output indicates that the PV interface 1/1 is bundled with virtual function 100/1, which is an SR-IOV VF NIC. The MAC addresses of both 1/1 and 100/1 NICs are the same. After accelerated networking is enabled, the data of 1/1 interface is sent through the data path of 100/1 interface, which is a ConnectX4 interface. You can see that the “show interface” output of PV interface (1/1) points to the VF (100/1). Similarly, the “show interface” output of VF interface (100/1) points to the PV interface (1/1).


```

> show interface 1/1
1) Interface 1/1 (NetScaler Virtual Interface, SmartNIC, VF 100/1, Datapath 100/1) #0
   flags=0xe060 <ENABLED, UP, UP, HEARTBEAT, 802.1q>
   MTU=1500, native vlan=10, MAC=00:0d:3a:9b:f2:1d, uptime 10h49m10s
   LLDP Mode: NONE, LR Priority: 1024

   RX: Pkts(310366) Bytes(98476082) Errs(0) Drops(0) Stalls(0)
   TX: Pkts(44) Bytes(6368) Errs(0) Drops(0) Stalls(0)
   NIC: InDisc(0) OutDisc(0) Fcfls(0) Stalls(0) Hangs(0) Muted(0)
   Bandwidth thresholds are not set.

Done
> show interface 100/1
1) Interface 100/1 (CX4 VF Interface, SmartNIC, PV 1/1) #3
   flags=0xe460 <ENABLED, UP, UP, 802.1q>
   MTU=1500, native vlan=10, MAC=00:0d:3a:9b:f2:1d, uptime 10h49m11s
   Actual: media FIBER, speed NONE, duplex FULL, fctl NONE, throughput
0
   LLDP Mode: NONE, LR Priority: 1024

   RX: Pkts(1135870) Bytes(1487381079) Errs(0) Drops(0) Stalls(0)
   TX: Pkts(1143020) Bytes(143165922) Errs(0) Drops(0) Stalls(0)
   NIC: InDisc(0) OutDisc(0) Fcfls(0) Stalls(0) Hangs(0) Muted(0)
   Bandwidth thresholds are not set.

Done
>

```

Points to note in Citrix ADC

- PV interface is considered as the primary or main interface for all the necessary operations. Configurations must be performed on PV interfaces only.
- All the 'set' operations on a VF interface are blocked except the following:
 - enable interface
 - disable interface
 - reset interface
 - clear stats

Note:

Citrix recommends that you do not perform any operations on the VF interface.

- You can verify the binding of PV interface with VF interface using the `show interface` command.

Configure a VLAN to a PV interface

When a PV interface is bound to a VLAN, the associated accelerated VF interface is also bound to the same VLAN as the PV interface. In this example, the PV interface (1/1) is bound to VLAN (20). The VF

interface (100/1) that is bundled with the PV interface (1/1) is also bound to VLAN 20.

Example:

1. Create a VLAN.

```
1 add vlan 20
```

2. Bind a VLAN to the PV interface.

```
1 bind vlan 20 -ifnum 1/1
2
3 show vlan
4
5 1) VLAN ID: 1
6    Link-local IPv6 addr: fe80::20d:3aff:fe9b:f21d/64
7    Interfaces : L0/1
8
9 2) VLAN ID: 10    VLAN Alias Name:
10   Interfaces : 0/1 100/1
11   IPs : 10.0.1.29  Mask: 255.255.255.0
12
13 3) VLAN ID: 20    VLAN Alias Name:
14   Interfaces : 1/1 100/2
```

Note:

VLAN binding operation is not permitted on an accelerated VF interface.

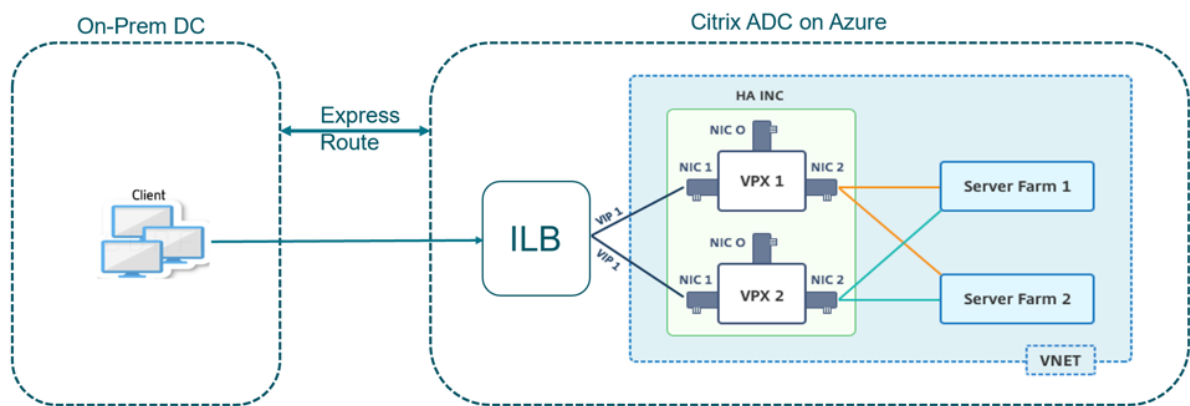
```
1 bind vlan 1 -ifnum 100/1
2 ERROR: Operation not permitted
```

Configure HA-INC nodes by using the Citrix high availability template with Azure ILB

September 9, 2024


You can quickly and efficiently deploy a pair of VPX instances in HA-INC mode by using the standard template for intranet applications. The Azure internal load balancer (ILB) uses an internal or private IP address for the front end as shown in Figure 1. The template creates two nodes, with three subnets and six NICs. The subnets are for management, client, and server-side traffic with each subnet belonging to a different NIC on each device.

Figure 1: Citrix ADC HA pair for clients in an internal network



You can also use this deployment when the Citrix ADC HA pair is behind a firewall as shown in Figure 2. The public IP address belongs to the firewall and is NAT'd to the front-end IP address of the ILB.

Figure 2: Citrix ADC HA pair with firewall having public IP address



You can get the Citrix ADC HA pair template for intranet applications at the [Azure portal](#).

Complete the following steps to launch the template and deploy a high availability VPX pair by using Azure Availability Sets.

1. From the Azure portal, navigate to the **Custom deployment** page.
2. The **Basics** page appears. Create a Resource Group. Under the **Parameters** tab, enter details for the Region, Admin user name, Admin Password, license type (VM sku), and other fields.

Custom deployment
Deploy from a custom template

12 resources

[Edit template](#) [Edit parameters](#)

Deployment scope
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](#)

Parameters

Region * ⓘ

Admin Username ⓘ ✓

Admin Password * ⓘ ✓

Vm Size ⓘ

Vm Sku ⓘ

Vnet Name ⓘ

Vnet Resource Group ⓘ

Vnet New Or Existing

Subnet Name-01 ⓘ

Subnet Name-11 ⓘ

Subnet Name-12 ⓘ

Subnet Address Prefix-01 ⓘ

Subnet Address Prefix-11 ⓘ

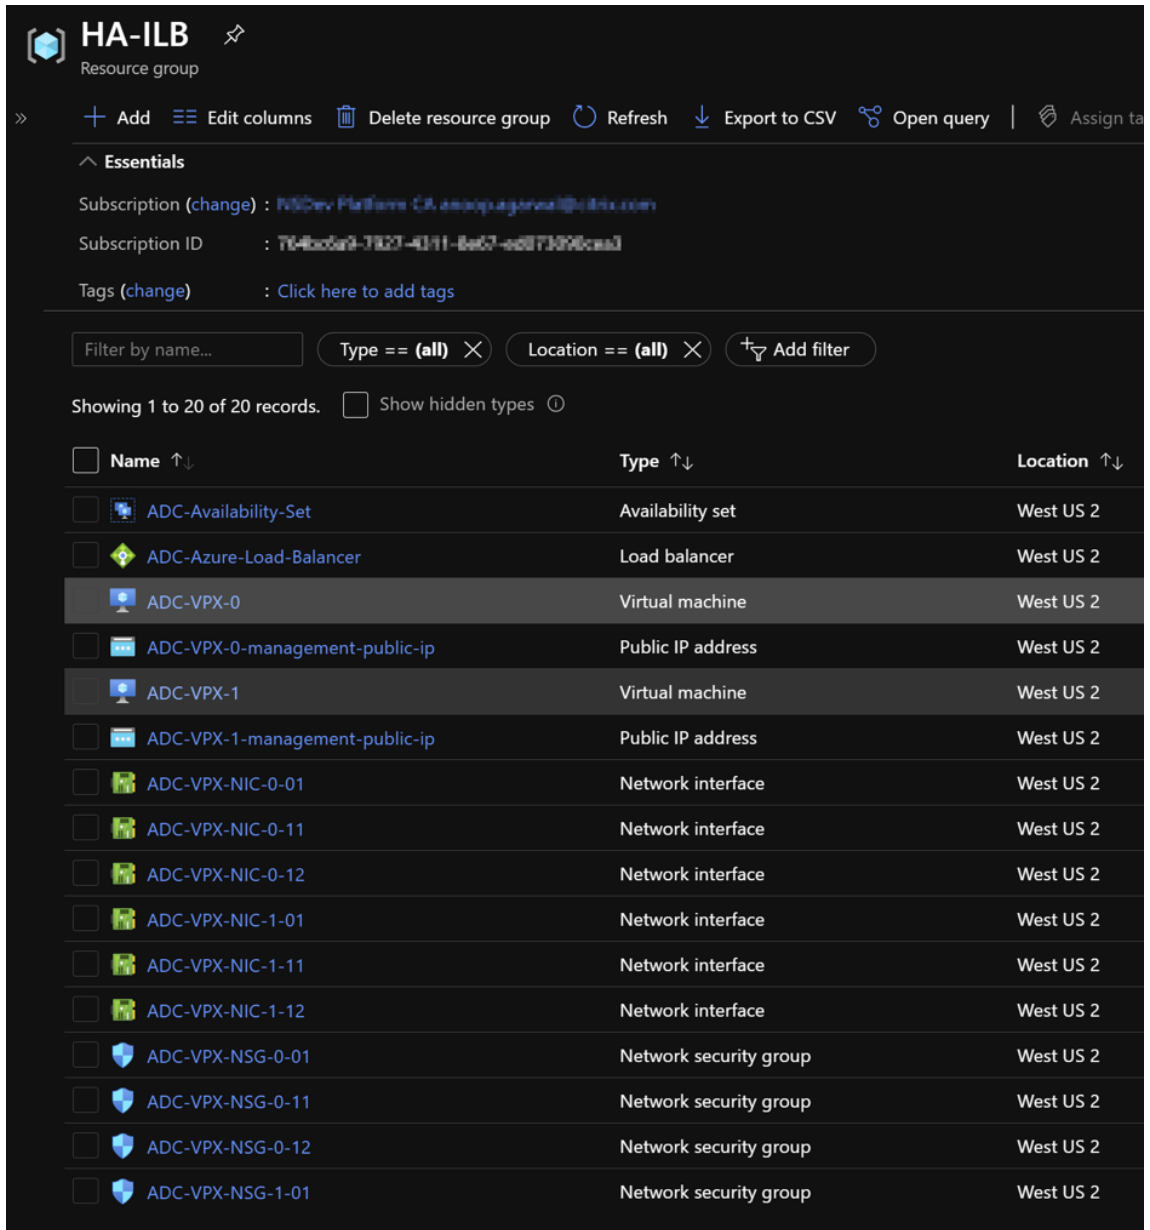
[Review + create](#) [< Previous](#) **[Next : Review + create >](#)**

3. Click **Next : Review + create >**.

It might take a moment for the Azure Resource Group to be created with the required configurations. After completion, select the Resource Group in the Azure portal to see the configuration details, such as LB rules, back-end pools, health probes. The high availability pair appears as ADC-VPX-0 and ADC-VPX-1.

If further modifications are required for your HA setup, such as creating more security rules and ports, you can do that from the Azure portal.

Once the required configuration is complete, the following resources are created.



4. You must log on to **ADC-VPX-0** and **ADC-VPX-1** nodes to validate the following configuration:

- NSIP addresses for both nodes must be in the management subnet.
- On the primary (ADC-VPX-0) and secondary (ADC-VPX-1) nodes, you must see two SNIP addresses. One SNIP (client subnet) is used for responding to ILB probes and the other SNIP (server subnet) is used for back-end server communication.

Note:

In the HA-INC mode, the SNIP address of the ADC-VPX-0 and ADC-VPX-1 VMs are different

while in the same subnet, unlike with the classic on-premises ADC HA deployment where both are the same.

To support deployments when the VPX pair SNIP is in different subnets, or anytime the VIP is not in the same subnet as a SNIP, you must either enable Mac-Based Forwarding (MBF), or add a static host route for each VIP to each VPX node.

On the primary node (ADC-VPX-0)

```
> sh ip
-----
1) 10.11.0.5 0 NetScaler IP Active Enabled Enabled NA Enabled
2) 10.11.1.5 0 SNIP Active Enabled Enabled NA Enabled
3) 10.11.3.4 0 SNIP Active Enabled Enabled NA Enabled
Done
>
>
```

```
> sh ha node
1) Node ID: 0
   IP: 10.11.0.5 (ADC-VPX-0)
   Node State: UP
   Master State: Primary
   Fail-Safe Mode: OFF
   INC State: ENABLED
   Sync State: ENABLED
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
   Sync Status Strict Mode: DISABLED
   Hello Interval: 200 msec
   Dead Interval: 3 secs
   Node in this Master State for: 0:0:20:26 (days:hrs:min:sec)
2) Node ID: 1
   IP: 10.11.0.4
   Node State: UP
   Master State: Secondary
   Fail-Safe Mode: OFF
   INC State: ENABLED
   Sync State: SUCCESS
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
Done
>
>
```

On the secondary node (ADC-VPX-1)

```

> sh ip
-----
Ipaddress      Traffic Domain  Type           Mode   Arp   Icmp   Vserver  State
-----
1) 10.11.0.4     0              NetScaler IP   Active Enabled Enabled NA      Enabled
2) 10.11.1.6     0              SNIP           Active Enabled Enabled NA      Enabled
3) 10.11.3.5     0              SNIP           Active Enabled Enabled NA      Enabled
Done
>


```

```


> sh ha node
1) Node ID: 0
   IP: 10.11.0.4 (ADC-VPX-1)
   Node State: UP
   Master State: Secondary
   Fail-Safe Mode: OFF
   INC State: ENABLED
   Sync State: SUCCESS
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
   Sync Status Strict Mode: DISABLED
   Hello Interval: 200 msec
   Dead Interval: 3 secs
   Node in this Master State for: 0:0:24:18 (days:hrs:min:sec)
2) Node ID: 1
   IP: 10.11.0.5
   Node State: UP
   Master State: Primary
   Fail-Safe Mode: OFF
   INC State: ENABLED
   Sync State: ENABLED
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
Done
>

```

5. After the primary and secondary nodes are UP and the Synchronization status is **SUCCESS**, you must configure the load balancing virtual server or the gateway virtual server on the primary node (ADC-VPX-0) with the private floating IP (FIP) address of the ADC Azure load balancer. For more information, see the [Sample configuration](#) section.
6. To find the private IP address of ADC Azure load balancer, navigate to **Azure portal > ADC Azure Load Balancer > Frontend IP configuration**.



7. In the **Azure Load Balancer** configuration page, the ARM template deployment helps create the LB rule, back-end pools, and health probes.



- The LB Rule (LbRule1) uses port 80, by default.

lbRule1
ADC-Azure-Load-Balancer

Save Discard Delete

i A load balancing rule distributes incoming traffic that is sent to a selected IP address and port combination across a group of backend pool instances. Only backend instances that the health probe considers healthy receive new traffic.

Name *
lbRule1

IP Version *
 IPv4 IPv6

Frontend IP address * ⓘ
10.11.1.4 (ADC-Load-Balancer-Frontend-IP-Configuration-rule) ✓

Protocol
 TCP UDP

Port *
80

Backend port * ⓘ
80

- Edit the rule to use port 443, and save the changes.

Note:

For enhanced security, Citrix recommends you to use SSL port 443 for LB virtual server or Gateway virtual server.

lbRule1

ADC-Azure-Load-Balancer

Save Discard Delete

i A load balancing rule distributes incoming traffic that is sent to a selected IP address and port combination across a group of backend pool instances. Only backend instances that the health probe considers healthy receive new traffic.

Name *

IP Version *
 IPv4 IPv6

Frontend IP address * ⓘ

Protocol
 TCP UDP

Port *
 ✓

Backend port * ⓘ

Backend pool ⓘ

Health probe ⓘ


Session persistence ⓘ

Idle timeout (minutes) ⓘ
 4


Floating IP ⓘ
Enabled

To add more VIP addresses on the ADC, perform the following steps:

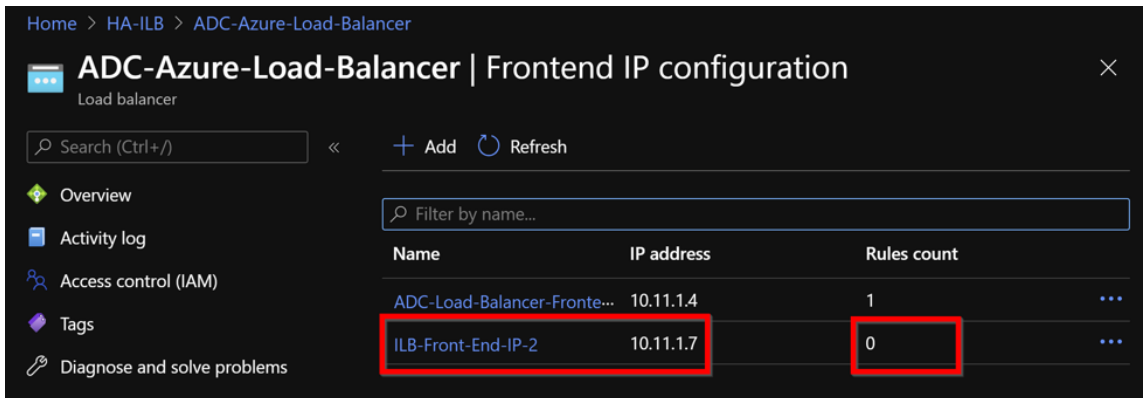
1. Navigate to **Azure Load Balancer > Frontend IP configuration**, and click **Add** to create a new internal load balancer IP address.




2. In the **Add frontend IP address** page, enter a name, choose the client subnet, assign either dynamic or static IP address, and click **Add**.



3. The front-end IP address is created but an LB Rule is not associated. Create a new load balancing rule, and associate it with the front-end IP address.



4. In the **Azure Load Balancer** page, select **Load balancing rules**, and then click **Add**.



5. Create a new LB Rule by choosing the new front-end IP address and the port. **Floating IP** field must be set to **Enabled**.

Home > HA-ILB > ADC-Azure-Load-Balancer >

Add load balancing rule

ADC-Azure-Load-Balancer

i A load balancing rule distributes incoming traffic that is sent to a selected IP address and port combination across a group of backend pool instances. Only backend instances that the health probe considers healthy receive new traffic.

1 Name *
lbrule2 ✓

IP Version *
 IPv4 IPv6

2 Frontend IP address * ⓘ
10.11.1.7 (ILB-Front-End-IP-2) ✓

Protocol
 TCP UDP

3 Port * **3**
443 ✓

4 Backend port * ⓘ **4**
443 ✓

5 Backend pool ⓘ **5**
ADC-Load-Balancer-Backend-rule (2 virtual machines) ✓

Health probe ⓘ
ADC-Load-Balancer-Health-Probe-rule (TCP:9000) ✓

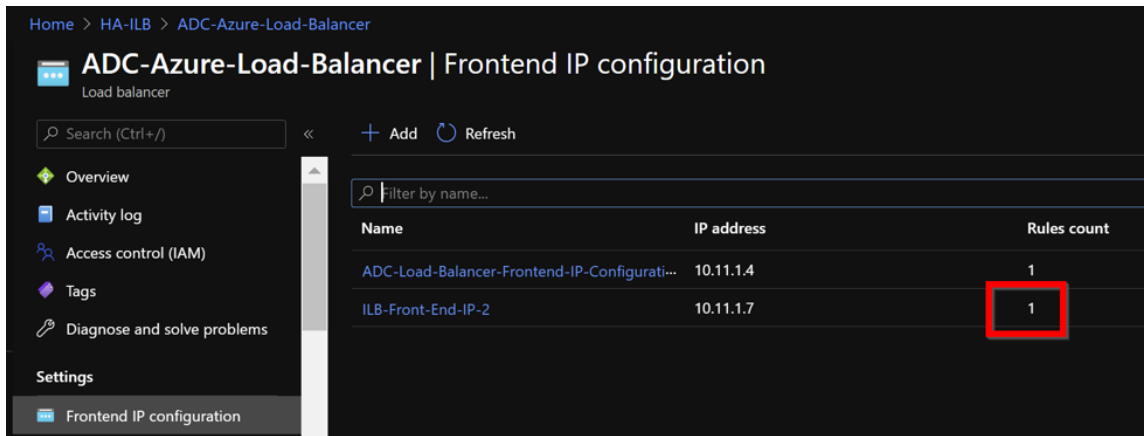
Session persistence ⓘ
None ✓

Idle timeout (minutes) ⓘ
0 4

6 Floating IP ⓘ **6**
Disabled Enabled

7 OK **7**

6. Now the **Frontend IP configuration** shows the LB rule that is applied.



Sample configuration

To configure a gateway VPN virtual server and load balancing virtual server, run the following commands on the primary node (ADC-VPX-0). The configuration auto synchronizes to the secondary node (ADC-VPX-1).

Gateway sample configuration

```

1 enable feature aaa LB SSL SSLVPN
2 enable ns mode MBF
3 add vpn vserver vpn_ssl SSL 10.11.1.4 443
4 add ssl certKey ckp -cert wild-cgwsanity.cer -key wild-cgwsanity.key
5 bind ssl vserver vpn_ssl -certkeyName ckp
    
```

Load balancing sample configuration

```

1 enable feature LB SSL
2 enable ns mode MBF
3 add lb vserver lb_vs1 SSL 10.11.1.7 443
4 bind ssl vserver lb_vs1 -certkeyName ckp
    
```

You can now access the load balancing or VPN virtual server using the fully qualified domain name (FQDN) associated with the internal IP address of ILB.

See the **Resources** section for more information about how to configure the load-balancing virtual server.

Resources:

The following links provide additional information related to HA deployment and virtual server configuration:

- [Configuring high availability nodes in different subnets](#)
- [Set up basic load balancing](#)

Related resources:

- [Configure a high-availability setup with multiple IP addresses and NICs by using PowerShell commands](#)
- [Configuring GSLB on Active-Standby HA Deployment on Azure](#)

Configure HA-INC nodes by using the Citrix high availability template for internet-facing applications

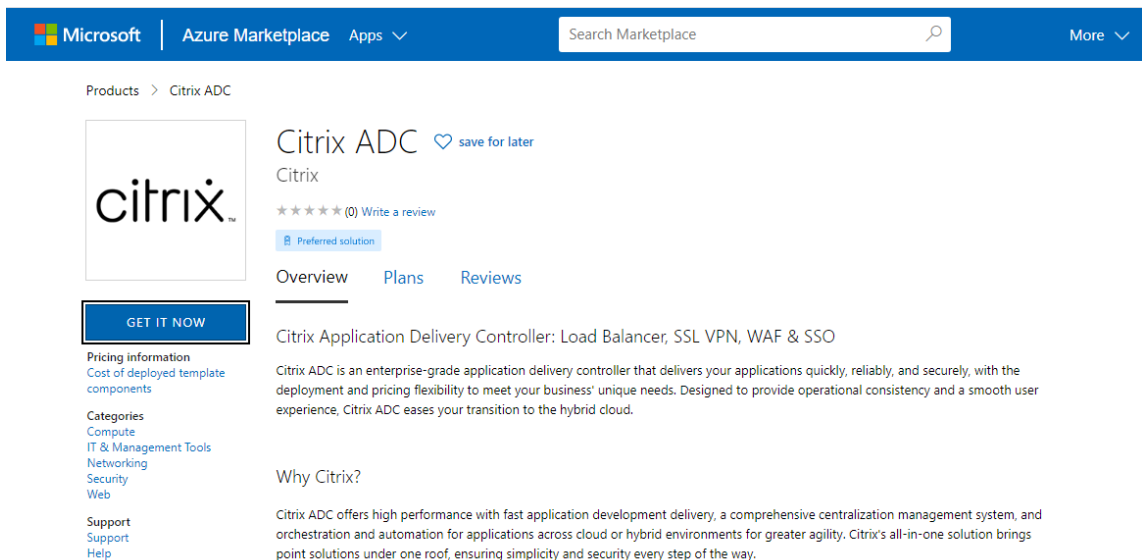
September 9, 2024

You can quickly and efficiently deploy a pair of VPX instances in HA-INC mode by using the standard template for internet-facing applications. The Azure load balancer (ALB) uses a public IP address for the front end. The template creates two nodes, with three subnets and six NICs. The subnets are for management, client, and server-side traffic. Each subnet has two NICs for both the VPX instances.


You can get the Citrix ADC HA pair template for internet-facing applications at the [Azure Marketplace](#).

Complete the following steps to launch the template and deploy a high availability VPX pair by using Azure availability sets or availability zone.

1. From the Azure Marketplace, search **Citrix ADC**.
2. Click **GET IT NOW**.



3. Select the required HA deployment along with license, and click **Continue**.



4. The **Basics** page appears. Create a Resource Group. Under the **Parameters** tab, enter details for the Region, Admin user name, Admin Password, license type (VM SKU), and other fields.

Create Citrix ADC

[Basics](#) [VM Configurations](#) [Network and Additional Settings](#) [Review + create](#)

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

Region * ⓘ

Citrix ADC Release Version * ⓘ 12.1 13.0

License Subscription ⓘ Bring Your Own License

Virtual Machine name * ⓘ

Administrator account

Username * ⓘ ✓

Authentication type * ⓘ Password SSH Public Key

Password * ⓘ ✓

Confirm password * ✓ ✓ Password

[Review + create](#) [< Previous](#) [Next : VM Configurations >](#)

5. Click **Next : VM Configurations >**.

Create Citrix ADC

Basics VM Configurations Network and Additional Settings Review + create

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

Region * ⓘ

Citrix ADC Release Version * ⓘ 12.1 13.0

License Subscription ⓘ Bring Your Own License

Virtual Machine name * ⓘ

Administrator account

Username * ⓘ ✓

Authentication type * ⓘ Password SSH Public Key

Password * ⓘ ✓

Confirm password * ✓ ✓ Password

[Review + create](#)

[< Previous](#)

[Next : VM Configurations >](#)

6. On the **VM Configurations** page, perform the following:

- Configure public IP domain name suffix
- Enable or disable **Azure Monitoring Metrics**
- Enable or disable **Backend Autoscale**

7. Click **Next: Network and Additional settings >**

Create Citrix ADC

Virtual machine size * ⓘ	1x Standard DS3 v2 4 vcpus, 14 GB memory Change size
OS disk type ⓘ	<input checked="" type="radio"/> Premium_LRS
Assign Public IP (Management) ⓘ	<input checked="" type="radio"/> Yes
Assign Public IP (Client traffic) ⓘ	<input checked="" type="radio"/> Yes
Unique public IP domain name suffix * ⓘ	<input type="text" value="d7a2c4d49e"/>
Azure Monitoring Metrics ⓘ	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled
Backend Autoscale ⓘ	<input type="radio"/> Enabled <input checked="" type="radio"/> Disabled

[Review + create](#) [< Previous](#) [Next : Network and Additional Settings >](#)

8. On **Network and Additional Settings** page, create Boot diagnostics account and configure the network settings.

Create Citrix ADC

[Basics](#)
 [VM Configurations](#)
 [Network and Additional Settings](#)
 [Review + create](#)

Boot diagnostics

Diagnostic storage account * ⓘ [Create New](#)

Network Settings

Configure virtual networks

Virtual network * ⓘ [Create new](#)

Management Subnet * ⓘ

Client Subnet * ⓘ

Server Subnet * ⓘ

Public IP (Management)

Management Public IP (NSIP) * ⓘ [Create new](#)

Management Domain Name ⓘ [.southindia.cloudapp.azure.com](#)

Public IP (Clientside)

Clientside Public IP (VIP) * ⓘ [Create new](#)

Clientside Domain Name ⓘ [.southindia.cloudapp.azure.com](#)

Public Inbound Ports (Management only)

Ports open for Management public IP ⓘ None ssh (22) ssh (22), http (80), https (443)

[Review + create](#)
 [< Previous](#)
 [Next : Review + create >](#)

9. Click **Next: Review + create >**.

10. Review the basic settings, VM configuration, network and additional settings, and click **Create**.

It might take a moment for the Azure Resource Group to be created with the required configura-

tions. After completion, select the Resource Group in the Azure portal to see the configuration details, such as LB rules, back-end pools, and health probes. The high availability pair appears as **citrix-adc-vpx-0** and **citrix-adc-vpx-1**.

If further modifications are required for your HA setup, such as creating more security rules and ports, you can do that from the Azure portal.

Once the required configuration is complete, the following resources are created.

Home > citrix.netscalervpx-1vm-3nic-20201006140352 >

Test_HA_Internet_App Resource group

» + Add Edit columns Delete resource group Refresh Export to CSV Open query Assign tags Move

Essentials

Filter by name... Type == all Location == all Add filter

Showing 1 to 23 of 23 records. Show hidden types

Name	Type
citrix-adc-vpx-0	Virtual machine
citrix-adc-vpx-0_OsDisk_1_6749f4a73c534051b0602ba6e3ec2cf8	Disk
citrix-adc-vpx-1	Virtual machine
citrix-adc-vpx-1_OsDisk_1_8fde7770497b4dbdba385715e81505c9	Disk
citrix-adc-vpx-nic01-0	Network interface
citrix-adc-vpx-nic01-1	Network interface
citrix-adc-vpx-nic01-nsg-0	Network security group
citrix-adc-vpx-nic01-nsg-1	Network security group
citrix-adc-vpx-nic11-0	Network interface
citrix-adc-vpx-nic11-1	Network interface
citrix-adc-vpx-nic11-nsg-0	Network security group
citrix-adc-vpx-nic11-nsg-1	Network security group
citrix-adc-vpx-nic12-0	Network interface
citrix-adc-vpx-nic12-1	Network interface
citrix-adc-vpx-nic12-nsg-0	Network security group
citrix-adc-vpx-nic12-nsg-1	Network security group
citrix-adc-vpx-nsip-0	Public IP address
citrix-adc-vpx-nsip-1	Public IP address
citrix-adc-vpx-vip	Public IP address
citrix-adc-vpx-vip-load-balancer	Load balancer
citrix-adc-vpx-virtual-network	Virtual network
citrix-adc-vpx-vm-availability-set	Availability set
citrixadcpx9db3901a6a	Storage account

11. You must log on to **citrix-adc-vpx-0** and **citrix-adc-vpx-1** nodes to validate the following configuration:

- NSIP addresses for both nodes must be in the management subnet.
- On the primary (citrix-adc-vpx-0) and secondary (citrix-adc-vpx-1) nodes, you must see

two SNIP addresses. One SNIP (client subnet) is used for responding to the ALB probes and the other SNIP (server subnet) is used for back-end server communication.

Note:

In the HA-INC mode, the SNIP addresses of the citrix-adc-vpx-0 and citrix-adc-vpx-1 VMs are different, unlike with the classic on-premises ADC high availability deployment where both are the same.

On the primary node (citrix-adc-vpx-0)

```
> sh ip
-----
1) 10.18.0.4 0 NetScaler IP Active Enabled Enabled NA Enabled
2) 10.18.1.5 0 SNIP Active Enabled Enabled NA Enabled
3) 10.18.2.4 0 SNIP Active Enabled Enabled NA Enabled
Done
```

```
> sh ha node
1) Node ID: 0
   IP: 10.18.0.4 (ns-vpx0)
   Node State: UP
   Master State: Primary
   Fail-Safe Mode: OFF
   INC State: ENABLED
   Sync State: ENABLED
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
   Sync Status Strict Mode: DISABLED
   Hello Interval: 200 msec
   Dead Interval: 3 secs
   Node in this Master State for: 0:3:34:21 (days:hrs:min:sec)
2) Node ID: 1
   IP: 10.18.0.5
   Node State: UP
   Master State: Secondary
   Fail-Safe Mode: OFF
   INC State: ENABLED
   Sync State: SUCCESS
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
Done
```

On the secondary node (citrix-adc-vpx-1)

```

> show ip
-----
1) 10.18.0.5      0      NetScaler IP      Active  Enabled  Enabled  NA      Enabled
2) 10.18.1.4      0      SNIP              Active  Enabled  Enabled  NA      Enabled
3) 10.18.2.5      0      SNIP              Active  Enabled  Enabled  NA      Enabled
Done
>

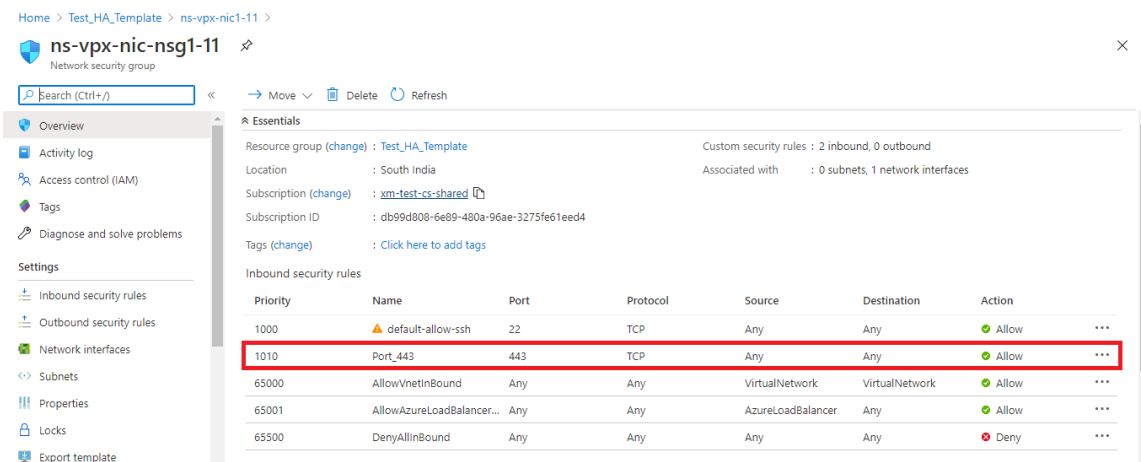
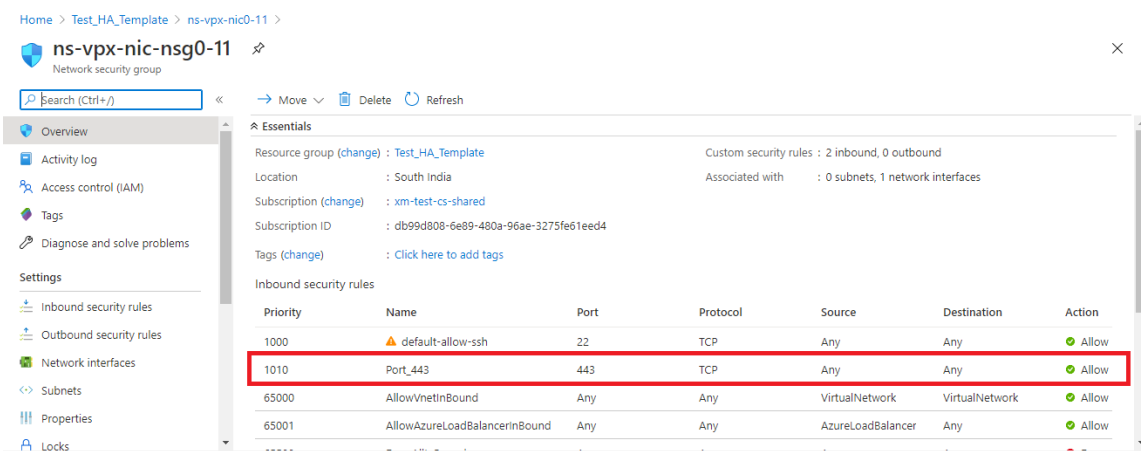
> sh ha node
1) Node ID:      0
   IP:          10.18.0.5 (ns-vpx1)
   Node State:  UP
   Master State: Secondary
   Fail-Safe Mode: OFF
   INC State:   ENABLED
   Sync State:  SUCCESS
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
   Sync Status Strict Mode: DISABLED
   Hello Interval: 200 msec
   Dead Interval: 3 secs
   Node in this Master State for: 0:3:23:51 (days:hrs:min:sec)
2) Node ID:      1
   IP:          10.18.0.4
   Node State:  UP
   Master State: Primary
   Fail-Safe Mode: OFF
   INC State:   ENABLED
   Sync State:  ENABLED
   Propagation: ENABLED
   Enabled Interfaces : 0/1 1/1 1/2
   Disabled Interfaces : None
   HA MON ON Interfaces : None
   HA HEARTBEAT OFF Interfaces : None
   Interfaces on which heartbeats are not seen : 1/1 1/2
   Interfaces causing Partial Failure: None
   SSL Card Status: NOT PRESENT
Done
>

```

12. After the primary and secondary nodes are UP and the Synchronization status is **SUCCESS**, you must configure the load balancing virtual server or the gateway virtual server on the primary node (citrix-adc-vpx-0) with the public IP address of the ALB virtual server. For more information, see the [Sample configuration](#) section.
13. To find the public IP address of ALB virtual server, navigate to **Azure portal > Azure Load Balancer > Frontend IP configuration**.



14. Add the inbound security rule for virtual server port 443 on the network security group of both the client interfaces.



15. Configure the ALB port that you want to access, and create inbound security rule for the specified port. The Backend port is your load balancing virtual server port or the VPN virtual server port.

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

Home > Test_HA_Template > alb >

lbRule1

alb

Save Discard Delete

IPv4 IPv6

Frontend IP address * 52.172.55.197 (jipconf-11)

Protocol TCP UDP

Port * 443

Backend port * 443

Backend pool bepool-11 (2 virtual machines)


Health probe probe-11 (TCP:9000)

Session persistence None

Idle timeout (minutes) 4

Floating IP (direct server return) Enabled

- Now, you can access the load balancing virtual server or the VPN virtual server using the fully qualified domain name (FQDN) associated with the ALB public IP address.



Sample configuration

To configure a gateway VPN virtual server and load balancing virtual server, run the following commands on the primary node (ADC-VPX-0). The configuration auto synchronizes to the secondary node (ADC-VPX-1).

Gateway sample configuration

```
1 enable feature aaa LB SSL SSLVPN
2 add ip 52.172.55.197 255.255.255.0 -type VIP
3 add vpn vserver vpn_ssl SSL 52.172.55.197 443
4 add ssl certKey ckp -cert cgwsanity.cer -key cgwsanity.key
5 bind ssl vserver vpn_ssl -certkeyName ckp
```

Load balancing sample configuration

```
1 enable feature LB SSL
2 enable ns mode MBF
3 add lb vserver lb_vs1 SSL 52.172.55.197 443
4 bind ssl vserver lb_vs1 -certkeyName ckp
```

You can now access the load balancing or VPN virtual server using the FQDN associated with the public IP address of ALB.

See the **Resources** section for more information about how to configure the load balancing virtual server.

Resources:

The following links provide additional information related to HA deployment and virtual server configuration:

- [Create virtual servers](#)
- [Set up basic load balancing](#)

Configure a high-availability setup with Azure external and internal load balancers simultaneously

September 9, 2024

The high availability pair on Azure supports both external and internal load balancers simultaneously.

You have the following two options to configure a high availability pair using both Azure external and internal load balancers:

- Using two LB virtual servers on the Citrix ADC appliance.
- Using one LB virtual server and an IP set. The single LB virtual server serves traffic to multiple IPs, which are defined by the IPset.

Perform the following steps to configure a high availability pair on Azure using both the external and internal load balancers simultaneously:

For Steps 1 and 2, use the Azure portal. For Steps 3 and 4, use the Citrix ADC VPX GUI or the CLI.

Step 1. Configure an Azure load balancer, either an external load balancer or an internal load balancer.

For more information on configuring high-availability setup with Azure external load balancers, see [Configure a high-availability setup with multiple IP addresses and NIC](#).

For more information on configuring high-availability setup with Azure internal load balancers, see [Configure HA-INC nodes by using the Citrix high availability template with Azure ILB](#).

Step 2. Create an extra load balancer (ILB) in your resource group. In Step 1, if you have created an external load balancer, you now create an internal load balancer and conversely.

- To create an internal load balancer, choose the load balancer type as **Internal**. For the **Subnet** field, you must choose your Citrix ADC client subnet. You can choose to provide a static IP address in that subnet, provided there are no conflicts. Otherwise, choose the dynamic IP address.

[Home](#) > [ansible_rg_ganeshb_1611818039](#) > [New](#) > [Load Balancer](#) >

Create load balancer

Project details	
Subscription *	<input type="text" value=""/>
Resource group *	<input type="text" value=""/> Create new
Instance details	
Name *	<input type="text" value="internal-load-balancer"/>
Region *	<input type="text" value="(US) West US 2"/>
Type * ⓘ	<input checked="" type="radio"/> Internal <input type="radio"/> Public
SKU * ⓘ	<input checked="" type="radio"/> Basic <input type="radio"/> Standard
Configure virtual network.	
Virtual network * ⓘ	<input type="text" value="automation_network"/>
Subnet *	<input type="text" value="ClientSubnet (192.168.2.0/24)"/> Manage subnet configuration
IP address assignment *	<input type="radio"/> Static <input checked="" type="radio"/> Dynamic

[Download a template for automation](#)

- To create an external load balancer, choose the load balancer type as **Public** and create the public IP address here.

The screenshot shows the 'Create load balancer' page in the Microsoft Azure portal. The page is titled 'Create load balancer' and includes a search bar at the top. The configuration options are as follows:

- Type ***: Internal Public
- SKU ***: Standard Basic
- Tier ***: Regional Global
- Public IP address**: Create new Use existing
- Public IP address name ***: [Empty text input field]
- Public IP address SKU**: Standard
- IP address assignment**: Dynamic Static
- Availability zone ***: [Dropdown menu]
- Add a public IPv6 address**: No Yes
- Routing preference**: Microsoft network Internet

At the bottom of the page, there are navigation buttons: 'Review + create', '< Previous', 'Next : Tags >', and 'Download a template for automation'.

1. After you have created the Azure Load Balancer, navigate to **Frontend IP configuration** and note down the IP address shown here. You must use this IP address while creating the ADC load balancing virtual server as in Step 3.



2. In the **Azure Load Balancer configuration** page, the ARM template deployment helps create the LB rule, back-end pools, and health probes.
3. Add the high availability pair client NICs to the backend pool for the ILB.
4. Create a health probe (TCP, 9000 port)
5. Create two load balancing rules:
 - One LB rule for HTTP traffic (webapp use case) on port 80. The rule must also use the backend port 80. Select the created backend pool and the health probe. Floating IP must be enabled.
 - Another LB rule for HTTPS or CVAD traffic on port 443. The process is the same as the HTTP traffic.

Step 3. On the primary node of Citrix ADC appliance, create a load balancing virtual server for ILB.

1. Add a load balancing virtual server.

```
1 add lb vsrver <name> <serviceType> [<ILB Frontend IP address>] [<port>]
```

Example:

```
1 add lb vsrver vsrver_name HTTP 52.172.96.71 80
```

Note:

Use the load balancer frontend IP address, which is associated with the additional Load balancer that you create in Step 2.

2. Bind a service to a load balancing virtual server.

```
1 bind lb vservice <name> <serviceName>
```

Example:

```
1 bind lb vservice Vserver-LB-1 Service-HTTP-1
```

For more information, see [Set up basic load balancing](#)

Step 4: As an alternative to Step 3, you can create a load balancing virtual server for ILB using IPsets.

1. Add an IP address of type virtual server IP (VIP).

```
1 add nsip <ILB Frontend IP address> -type <type>
```

Example:

```
1 add nsip 52.172.96.71 -type vip
```

2. Add an IPset on both primary and secondary nodes.

```
1 add ipset <name>
```

Example:

```
1 add ipset ipset1
```

3. Bind IP addresses to the IP set.

```
1 bind ipset <name> <ILB Frontend IP address>
```

Example:

```
1 bind ipset ipset1 52.172.96.71
```

4. Set the existing LB virtual server to use the IPset.

```
1 set lb vservice <vserver name> -ipset <ipset name>
```

Example:

```
1 set lb vservice vserver_name -ipset ipset1
```

For more information, see [Configure a multi-IP virtual server](#).

Install a Citrix ADC VPX instance on Azure VMware Solution

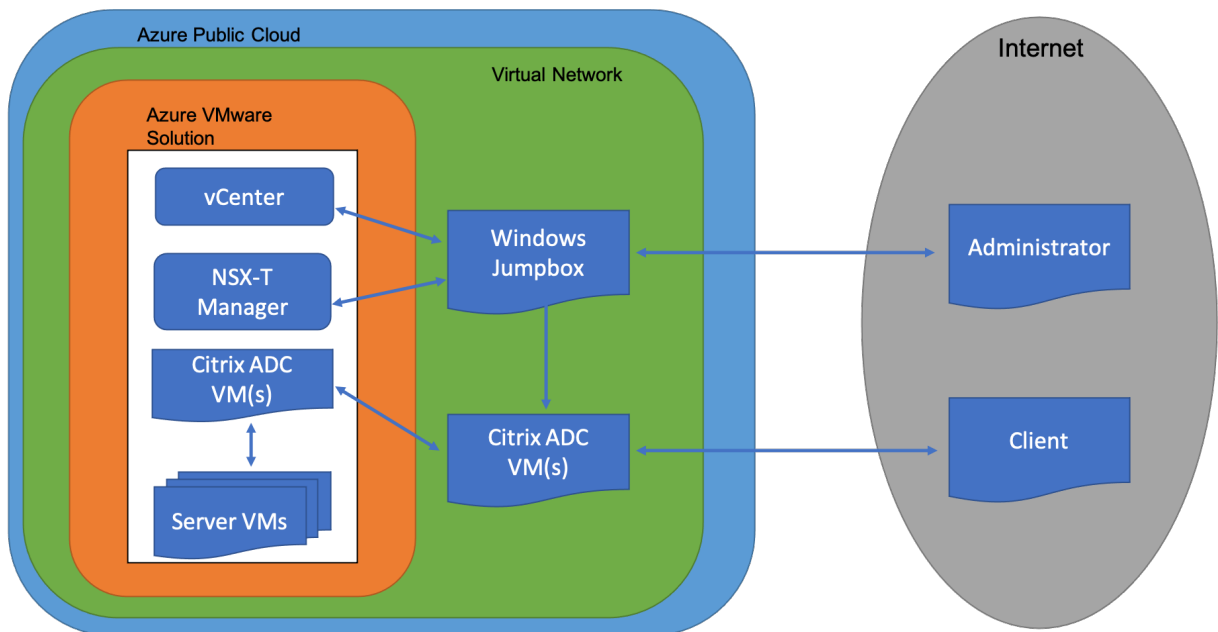
September 6, 2024

Azure VMware Solution (AVS) provides you with private clouds that contain vSphere clusters, built from dedicated bare-metal Azure infrastructure. The minimum initial deployment is three hosts, but additional hosts can be added one at a time, up to a maximum of 16 hosts per cluster. All provisioned private clouds have vCenter Server, vSAN, vSphere, and NSX-T.

The VMware Cloud (VMC) on Azure enables you to create cloud software-defined data centers (SDDC) on Azure with the number of ESX hosts that you want. The VMC on Azure supports Citrix ADC VPX deployments. VMC provides a user interface same as on-prem vCenter. It functions similar to the ESX-based Citrix ADC VPX deployments.

The following diagram shows the Azure VMware solution on the Azure public cloud that an administrator or a client can access over the internet. An administrator can create, manage, and configure workload or server VMs using Azure VMware solution. The admin can access the AVS’s web-based vCenter and NSX-T Manager from a Windows Jumpbox. You can create the Citrix ADC VPX instances (standalone or high availability pair) and server VMs within Azure VMware Solution using vCenter, and manage the corresponding networking using NSX-T manager. The Citrix ADC VPX instance on AVS works similar to the on-prem VMware cluster of hosts. AVS is managed from a Windows Jumpbox that is created in the same virtual network.

A client can only access the AVS service by connecting to the VIP of ADC. Another Citrix ADC VPX instance outside Azure VMware Solution but in the same Azure virtual network helps add the VIP of the Citrix ADC VPX instance within Azure VMware Solution as a service. As per requirement, you can configure the Citrix ADC VPX instance to provide service over the internet.



Prerequisites

Before you begin installing a virtual appliance, do the following:

- For more information on Azure VMware solution and its prerequisites, see [Azure VMware Solution documentation](#).
- For more information on deploying Azure VMware solution, see [Deploy an Azure VMware Solution private cloud](#).
- For more information on creating a Windows Jump box VM to access and manage Azure VMware Solution, see [Access an Azure VMware Solution private cloud](#)
- In Windows Jump box VM, download the Citrix ADC VPX appliance setup files.
- Create appropriate NSX-T network segments on VMware SDDC to which the virtual machines connect. For more information, see [Add a network segment in Azure VMware Solution](#)
- Obtain VPX license files.
- Virtual machines (VMs) created or migrated to the Azure VMware Solution private cloud must be attached to a network segment.

VMware cloud hardware requirements

The following table lists the virtual computing resources that the VMware SDDC must provide for each VPX nCore virtual appliance.

Table 1. Minimum virtual computing resources required for running a Citrix ADC VPX instance

Component	Requirement
Memory	2 GB
Virtual CPU (vCPU)	2
Virtual network interfaces	In VMware SDDC, you can install a maximum of 10 virtual network interfaces if the VPX hardware is upgraded to version 7 or higher.
Disk space	20 GB

Note:

This is in addition to any disk requirements for the hypervisor.

For production use of the VPX virtual appliance, the full memory allocation must be reserved.

OVF Tool 1.0 system requirements

OVF Tool is a client application that can run on Windows and Linux systems. The following table describes the system requirements for installing OVF tool.

Table 2. System requirements for OVF tool installation

Component	Requirement
Operating system	For detailed requirements from VMware, search for the “OVF Tool User Guide” PDF file at http://kb.vmware.com/ .
CPU	750 MHz minimum, 1 GHz or faster recommended
RAM	1 GB Minimum, 2 GB recommended
NIC	100 Mbps or faster NIC

For information about installing OVF, search for the “OVF Tool User Guide” PDF file at <http://kb.vmware.com/>.

Downloading the Citrix ADC VPX setup files

The Citrix ADC VPX instance setup package for VMware ESX follows the Open Virtual Machine (OVF) format standard. You can download the files from the Citrix website. You need a Citrix account to log on. If you do not have a Citrix account, access the home page at <http://www.citrix.com>. Click the **New Users link**, and follow the instructions to create a new Citrix account.

Once logged on, navigate the following path from the Citrix home page:

Citrix.com > **Downloads** > **Citrix ADC** > **Virtual Appliances**.


Copy the following files to a workstation on the same network as the ESX server. Copy all three files into the same folder.

- NSVPX-ESX-<release number>-<build number>-disk1.vmdk (for example, NSVPX-ESX-13.0-79.64-disk1.vmdk)
- NSVPX-ESX-<release number>-<build number>.ovf (for example, NSVPX-ESX-13.0-79.64.ovf)
- NSVPX-ESX-<release number>-<build number>.mf (for example, NSVPX-ESX-13.0-79.64.mf)

Deploy Azure VMware solution

1. Log in to your [Microsoft Azure portal](#), and navigate to **Azure Marketplace**.

2. From the **Azure Marketplace**, search **Azure VMware Solution** and click **Create**.



3. In the **Create a private cloud** page, enter the following details:

- Select a minimum of 3 ESXi hosts to create the default cluster of your private cloud.
- For the **Address block** field, use **/22** address space.
- For the **Virtual Network**, make sure that the CIDR range doesn't overlap with any of your on-premises or other Azure subnets (virtual networks) or with the gateway subnet.
- Gateway subnet is used to express route the connection with private cloud.

[Home](#) >

Create a private cloud

Azure settings

Subscription * ⓘ

Resource group * ⓘ
[Create new](#)

Location * ⓘ

General

Resource name * ⓘ

SKU * ⓘ


ESXi hosts * ⓘ

\$11,929.68
estimated monthly total


Address block * ⓘ

Virtual Network
[Create new](#)
Only Virtual Networks with a valid subnet with the name "GatewaySubnet" are available for selection. For details about adding subnet in a virtual network, refer to details [here](#)


4. Click **Review + Create**.
5. Review the settings. If you must change any settings, click **Previous**.



6. Click **Create**. Private cloud provisioning process starts. It can take up to two hours for the private cloud to be provisioned.



7. Click **Go to resource**, to verify the private cloud that is created.



Note:

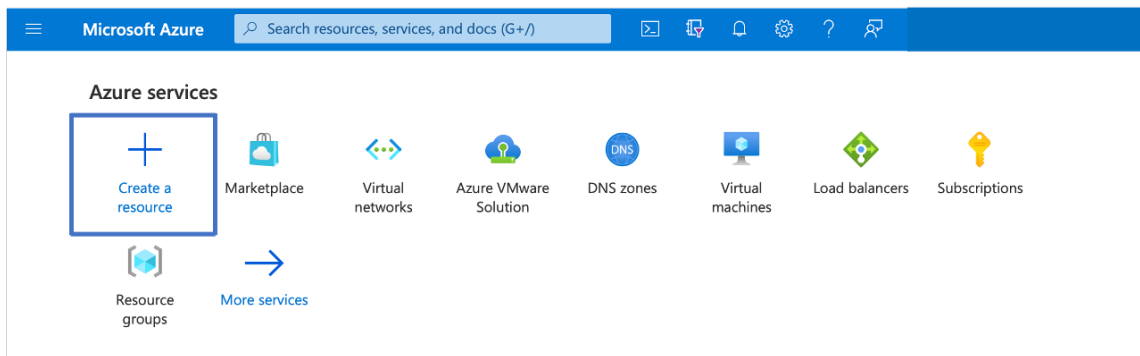
To access this resource, you need a VM in Windows that acts as a Jump box.

Connect to an Azure virtual machine running Windows


This procedure shows you how to use the Azure portal to deploy a virtual machine (VM) in Azure that runs Windows Server 2019. To see your VM in action, you then RDP to the VM and install the IIS web server.

To access the private cloud that you have created, you need to create a Windows Jump box within the same virtual network.

1. Go to the **Azure portal**, and click **Create a Resource**.



2. Search for **Microsoft Windows 10**, and click **Create**.



3. Create a virtual machine (VM) that runs Windows Server 2019. The **Create a virtual machine** page appears. Enter all the details in **Basics** tab, and select the **Licensing** check box. Leave the remaining defaults and then select the **Review + create** button at the bottom of the page.

Home > Create a resource > Microsoft Windows 10 >

Create a virtual machine

Basics Disks Networking Management Advanced Tags Review + create

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization. [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 * [Create new](#)

Instance details

Virtual machine name *

Region *

Availability options

Image * [See all images](#)

Azure Spot instance

Size * [See all sizes](#)

Administrator account

Username *

Password *

Confirm password *

Inbound port rules
Select which virtual machine network ports are accessible from the public internet. You can specify more limited or granular network access on the Networking tab.

Public inbound ports * None Allow selected ports

Select inbound ports *

⚠ This will allow all IP addresses to access your virtual machine. This is only recommended for testing. Use the Advanced controls in the Networking tab to create rules to limit inbound traffic to known IP addresses.

Licensing

I confirm I have an eligible Windows 10 license with multi-tenant hosting rights. [Review multi-tenant hosting rights for Windows 10 compliance](#)

[Review + create](#) < Previous Next: Disks >

4. After validation runs, select the **Create** button at the bottom of the page.
5. After the deployment is complete, select **Go to resource**.
6. Go to the Windows VM that you have created. Use the public IP address of the Windows VM and connect using RDP.


Use the **Connect** button in the Azure portal to start a Remote Desktop (RDP) session from a Windows desktop. First you connect to the virtual machine, and then you sign on.

To connect to a Windows VM from a Mac, you must install an RDP client for Mac such as Microsoft

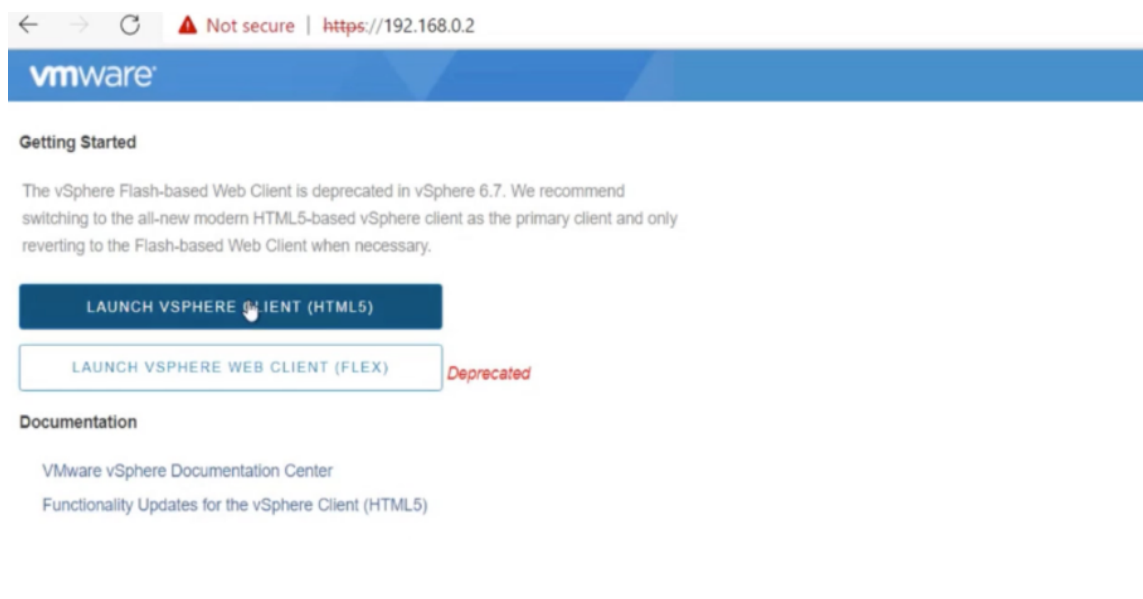
Remote Desktop. For more information, see [How to connect and sign on to an Azure virtual machine running Windows](#).

Access your Private Cloud vCenter portal


1. In your Azure VMware Solution private cloud, under **Manage**, select **Identity**. Make note of the vCenter credentials.




2. Launch the vSphere client by typing the vCenter web client URL.



3. Log in to VMware vSphere using the vCenter credentials of your Azure VMware Solution private cloud.



4. In the vSphere client, you can verify the ESXi hosts that you created in Azure portal.




For more information, see [Access your Private Cloud vCenter portal.](#)

Create an NSX-T segment in the Azure portal


You can create and configure an NSX-T segment from the Azure VMware Solution console in the Azure portal. These segments are connected to the default Tier-1 gateway, and the workloads on these seg-

ments get East-West and North-South connectivity. Once you create the segment, it displays in NSX-T Manager and vCenter.


1. In your Azure VMware Solution private cloud, under **Workload Networking**, select **Segments > Add**. Provide the details for the new logical segment and select **OK**. You can create three separate segments for Client, Management, and Server interfaces.




2. In your Azure VMware Solution private cloud, under **Manage**, select **Identity**. Make note of the NSX-T Manager credentials.



3. Launch the VMware NSX-T Manager by typing the NSX-T web client URL.



4. In the NSX-T manager, under **Networking > Segments**, you can see all the segments that you have created. You can also verify the subnets.




For more information, see [Create an NSX-T segment in the Azure portal](#).

Install a Citrix ADC VPX instance on VMware cloud

After you have installed and configured VMware Software-Defined Data Center (SDDC), you can use the SDDC to install virtual appliances on the VMware cloud. The number of virtual appliances that you can install depends on the amount of memory available on the SDDC.

To install Citrix ADC VPX instances on VMware cloud, perform these steps in Windows Jumpbox VM:

1. Download the Citrix ADC VPX instance setup files for ESXi host from the Citrix Downloads site.

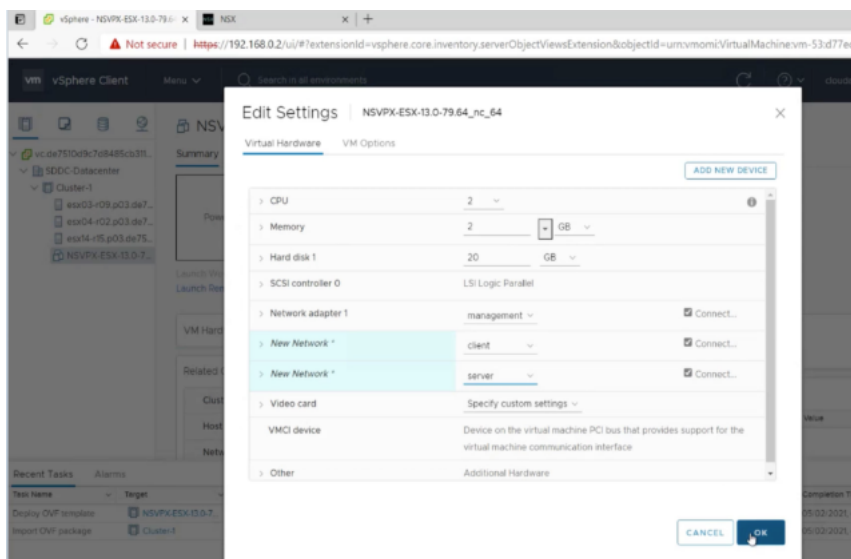


2. Open VMware SDDC in the Windows Jumpbox.
3. In the **User Name** and **Password** fields, type the administrator credentials, and then click **Login**.
4. On the **File** menu, click **Deploy OVF Template**.
5. In the **Deploy OVF Template** dialog box, in **Deploy from file** field, browse to the location at which you saved the Citrix ADC VPX instance setup files, select the .ovf file, and click **Next**.


Note:

By default, the Citrix ADC VPX instance uses E1000 network interfaces. To deploy ADC with the VMXNET3 interface, modify the OVF to use VMXNET3 interface instead of E1000. Availability of VMXNET3 interface is limited by Azure infrastructure and might not be available in Azure VMware Solution.


6. Map the networks shown in the virtual appliance OVF template to the networks that you configured on the VMware SDDC. Click **OK**.




7. Click **Finish** to start installing a virtual appliance on VMware SDDC.



8. You are now ready to start the Citrix ADC VPX instance. In the navigation pane, select the Citrix ADC VPX instance that you have installed and, from the right-click menu, select **Power On**. Click the **Console** tab to emulate a console port.



9. You are now connected to the Citrix ADC VM from the vSphere client.



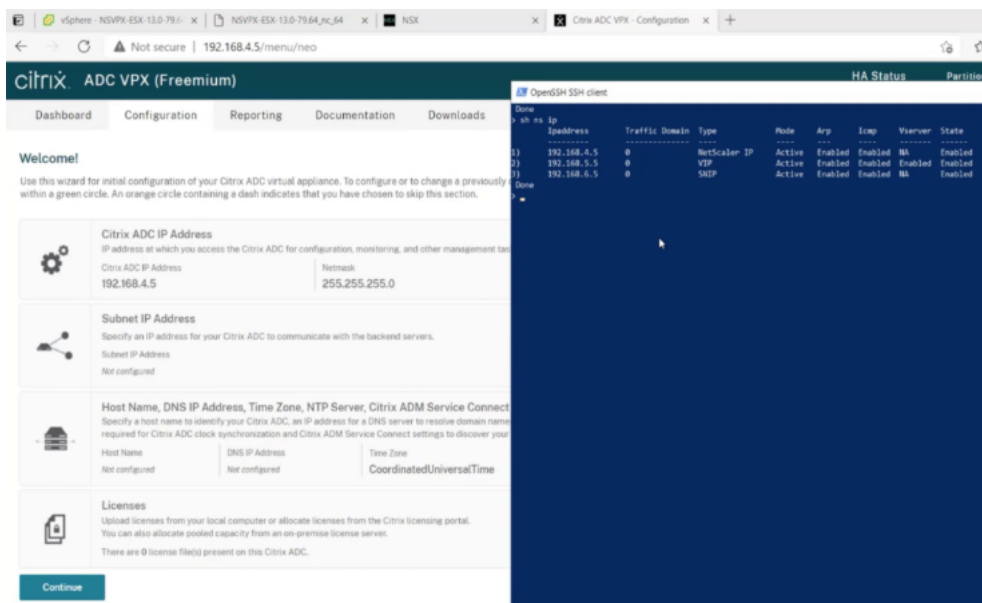
10. To access the Citrix ADC appliance by using the SSH keys, type the following command in the CLI:

```
1 ssh nsroot@<management IP address>
```

Example:

```
1 ssh nsroot@192.168.4.5
```

11. You can verify the ADC configuration by using the `show ns ip` command.



Add Azure Autoscale settings


September 6, 2024

Efficient hosting of applications in a cloud involves easy and cost-effective management of resources depending on the application demand. To meet increasing demand, you have to scale network resources upward. Whether demand subsides, you must scale down to avoid the unnecessary cost of idle resources. To minimize the cost of running the application, you have to constantly monitor traffic, memory and CPU use, and so on. However, monitoring traffic manually is cumbersome. For the application environment to scale up or down dynamically, you must automate the processes of monitoring traffic and of scaling resources up and down whenever necessary.

You can use Autoscale with Azure virtual machine scale sets (VMSS) for VPX multi-IP standalone and high availability deployment on Azure.

Integrated with the Azure virtual machine scale sets (VMSS) and Autoscale feature, the Citrix ADC VPX instance provides the following advantages:

- **Load balance and management:** Auto configures servers to scale up and scale down, depending on demand. The VPX instance auto detects the VMSS Autoscale setting in the back-end subnet in the same resource group as the VPX instance and allows the user to select the VMSS Autoscale setting to balance the load. All of this is done by auto configuring Citrix ADC virtual and subnet IP addresses on the VPX instance.
- **High availability:** Detects Autoscale groups in the same resource group and load-balance servers.
- **Better network availability:** The VPX instance supports back-end servers on different virtual networks (VNets).



For more information, see the following Azure topic

- [Virtual Machine Scale Sets Documentation](#)

- [Overview of Autoscale in Microsoft Azure Virtual Machines, Cloud Services, and Web Apps](#)

Before you begin

1. Read Azure-related usage guidelines. For more information, see [Deploy a Citrix ADC VPX instance on Microsoft Azure](#).
2. Create one or more Citrix ADC VPX instances with three network interfaces on Azure according to your requirement (standalone or high availability deployment).
3. Open the TCP 9001 port on the network security group of the 0/1 interface of the VPX instance. The VPX instance uses this port to receive the scale-out and scale-in notification.
4. Create an Azure virtual machine scale set (VMSS) in the same resource group. If you don't have an existing VMSS configuration, complete the following tasks:
 - a) Create a VMSS
 - b) Enable Autoscale on VMSS
 - c) Create scale-in and scale-out policy in VMSS Autoscale setting

For more information, see [Overview of Autoscale with Azure virtual machine scale sets](#).

5. Create an Azure Active Directory (AAD) application and service principal that can access resources. Assign contributor role to the newly created AAD application. For more information, see [Use portal to create an Azure Active Directory application and service principal that can access resources](#).

Add VMSS to a Citrix ADC VPX instance

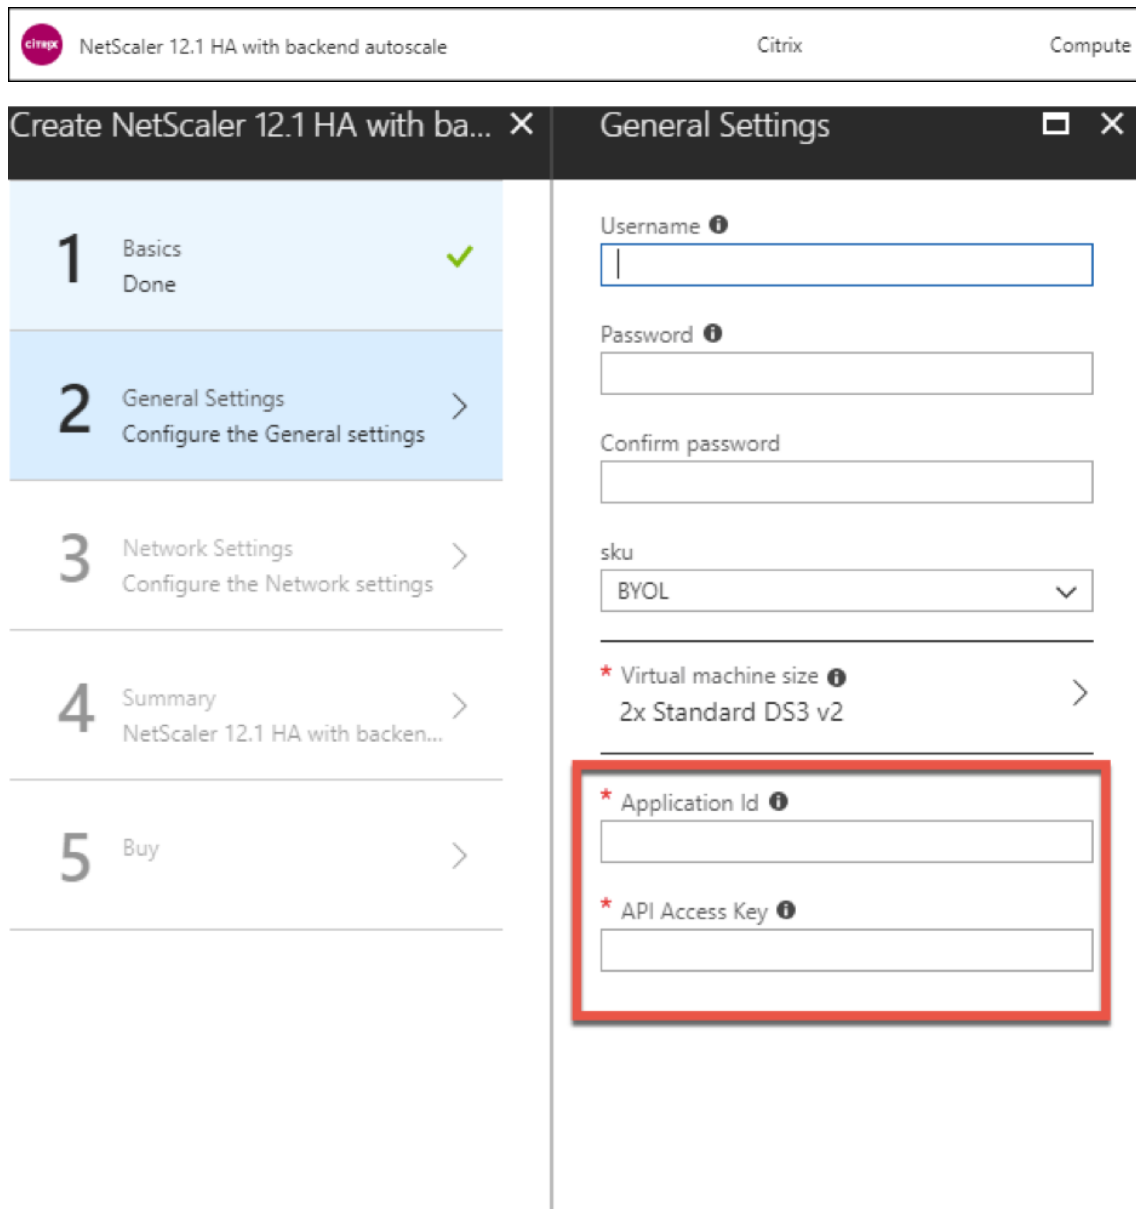
You can add the Autoscale setting to a VPX instance with a single click by using the GUI. Complete these steps to add the Autoscale setting to the VPX instance:

1. Log on to the VPX instance.
2. When you log on to the Citrix ADC VPX instance for the first time, you see the Set Credentials page. Add the required Azure credentials for the Autoscale feature to work.

The screenshot shows the Citrix NetScaler VPX AZURE Configuration page. At the top, there is a dark blue header with the text "Citrix NetScaler VPX AZURE". Below the header, there are two tabs: "Dashboard" and "Configuration". The "Configuration" tab is active. Below the tabs, there is a blue back arrow icon followed by the text "Set Credentials". Below this, there are three input fields: "Tenant ID", "Application ID", and "Application Secret". At the bottom of the form, there are two buttons: "OK" and "Cancel".

The Set Credential page appears only when the application ID and API access key are not set or the correct application ID and API access keys (same as application secret) is not set in the Azure portal.

When you deploy the “NetScaler 12.1 HA with back end Autoscale” offer from the Azure Marketplace, the Azure portal prompts for Azure service principal credentials (application ID and API access key).



For information about how to create an application ID see [Adding an application](#) and to create an access key or application secret see [Configure a client application to access web APIs](#).

3. In the default cloud profile page, enter the details, as shown in the following example, and click Create.

Dashboard Configuration

Name
 ?

Virtual Server IP Address*
 ▼

Load Balancing Server Protocol*
 ▼

Load Balancing Server Port*

Auto Scale Setting*
 ▼


Auto Scale Setting Protocol
 ▼

Auto Scale Setting Port*

Points to keep in mind while creating a cloud profile

- The virtual server IP address is auto-populated from the free IP address available to the VPX instance. For more information, see [Assign multiple IP addresses to virtual machines using the Azure portal](#).
- Autoscale setting is prepopulated from the VMSS Autoscale setting configured in the current resource group on your Azure account. For more information, see [Overview of Autoscale with Azure virtual machine scale sets](#).
- While selecting the Auto Scaling Group protocol and port, ensure your servers listen on those protocol and ports and you bind the correct monitor in the service group. By default, the TCP monitor is used.
- For SSL Protocol type Autos Scaling, after you create the Cloud Profile the load balance virtual server or service group will be down because of a missing certificate. You can bind the certificate to the virtual server or service group manually.

After the first time logon, if you want to create a cloud profile, on the GUI go to System > Azure > Cloud Profile and click Add.



The Create Cloud Profile configuration page appears.

Create Cloud Profile

Name
CloudProfile

Virtual Server IP Address*
10.1.27.5

Load Balancing Server Protocol
HTTP

Load Balancing Server Port
80

Auto Scale Setting*
SharePoint

Auto Scale Setting Protocol
HTTP

Auto Scale Setting Port
80

Create Close

Cloud Profile creates a Citrix ADC load-balancing (LB) virtual server (virtual server) and a service group with members (servers) as the servers of the Auto Scaling Group. Your back-end servers must be reachable through the SNIP configured on the VPX instance.

To view autoscale-related information in the Azure portal, go to All service > Virtual machine scale set > Select Virtual machine scale set > Scaling.

Azure tags for Citrix ADC VPX deployment


September 12, 2024

In the Azure cloud portal, you can tag resources with a name: value pair (such as Dept: Finance) to categorize and view resources across resource groups and, within the portal, across subscriptions. Tagging is helpful when you need to organize resources for billing or management or automation.

How Azure tag works for VPX deployment

For Citrix ADC VPX standalone and high-availability instances deployed on Azure Cloud, now you can create load balancing service groups associated with an Azure tag. The VPX instance constantly monitors Azure virtual machines (back-end servers) and network interfaces (NICs), or both, with the respective tag and updates the service group accordingly.

The VPX instance creates the service group that load balances the back-end servers using tags. The instance queries the Azure API for all resources that are tagged with a particular tag name and tag value. Depending on the assigned poll period (by default 60 seconds), the VPX instance periodically polls the Azure API and retrieves the resources available with the tag name and tag values assigned in the VPX GUI. Whenever a VM or NIC with the appropriate tag is added or deleted, the ADC detects the respective change and adds or deletes the VM or NIC IP address from the service group automatically.



Before you begin

Before creating Citrix ADC load balancing service groups, add a tag to the servers in Azure. You can assign the tag to either the virtual machine or to NIC.

Edit tags
Tags for demoGroup

NAME	VALUE
Dept ▼	Finance ▼ 🗑️
Environment ▼	Production ▼ 🗑️
<i>name</i> ▼	<i>value</i> ▼ + 🗑️

2 to be added

Save
Cancel

For more information about adding Azure tags, see Microsoft document [Use tags to organize your Azure resources](#).

Note:

ADC CLI commands to add Azure tag settings support tag names and tag values that start only with numerals or alphabets and not other keyboard characters.

How to add Azure tag settings by using VPX GUI

You can add the Azure tag cloud profile to a VPX instance by using the VPX GUI so that the instance can load balance the back-end servers using the specified tag. Follow these steps:

1. From the VPX GUI, go to **Configuration > Azure > Cloud Profile**.
2. Click Add to create a cloud profile. The cloud profile window opens.

Create Cloud Profile

Name

Virtual Server IP Address*

Type

Azure Tag Name

Azure Tag Value

Azure Poll Periods

Load Balancing Server Protocol

Load Balancing Server Port

Azure Tag Setting*

Azure Tag Setting Protocol

Azure Tag Setting Port

1. Enter values for the following fields:

- Name: Add a name for your profile
- Virtual Server IP Address: The virtual server IP address is auto-populated from the free IP address available to the VPX instance. For more information, see [Assign multiple IP addresses to virtual machines using the Azure portal](#).
- Type: From the menu, select AZURETAGS.
- Azure Tag Name: Enter the name that you have assigned to the VMs or NICs in the Azure portal.
- Azure Tag Value: Enter the value that you have assigned to the VMs or NICs in Azure portal.
- Azure Poll Periods: By default the poll period is 60 seconds, which is the minimum value. You can change it according to your requirement.
- Load Balancing Server Protocol: Select the protocol that your load balancer listens on.
- Load Balancing Server Port: Select the port that your load balancer listens on.
- Azure tag setting: The name of the service group that will be created for this cloud profile.
- Azure Tag Setting Protocol: Select the protocol that your back-end servers listen on.
- Azure Tag Setting Port: Select the port that your back-end servers listen on.

2. Click **Create**.

A load-balancer virtual server and a service group are created for the tagged VMs or NICs. To see the load balancer virtual server, from the VPX GUI, navigate to **Traffic Management > Load Balancing > Virtual Servers**.

How to add Azure tag settings by using VPX CLI

Type the following command on Citrix ADC CLI to create a cloud profile for Azure tags.

```
1 add cloud profile `<profile name>` -type azuretags -vServerName `<vserver name>` -serviceType HTTP -IPAddress `<vserver IP address>` -port 80 -serviceGroupName `<service group name>` -boundServiceGroupSvcType HTTP -vsrvbindsvcpport 80 -azureTagName `<Azure tag specified on Azure portal>` -azureTagValue `<Azure value specified on the Azure portal>` -azurePollPeriod 60
```

Important:

You must save all configurations; otherwise, the configurations are lost after you restart the instance. Type `save config`.

Example 1: Here's a sample command for a cloud profile for HTTP traffic of all Azure VMs/NICs tagged with the "myTagName/myTagValue" pair:

```

1 add cloud profile MyTagCloudProfile -type azuretags -vServerName
  MyTagVServer -serviceType HTTP -IPAddress 40.115.116.57 -port 80 -
  serviceGroupName MyTagsServiceGroup -boundServiceGroupSvcType HTTP -
  vsvrbindsvcport 80 -azureTagName myTagName -azureTagValue myTagValue
  -azurePollPeriod 60
2 Done

```

To display the cloud profile, type `show cloudprofile`.

Example 2: The following CLI command prints information about the newly added cloud profile in example 1.

```

1 show cloudprofile
2 1)   Name: MyTagCloudProfile Type: azuretags           VServerName:
      MyTagVServer ServiceType: HTTP           IPAddress: 52.178.209.133
      Port: 80           ServiceGroupName: MyTagsServiceGroup
      BoundServiceGroupSvcType: HTTP
3     Vsvrbindsvcport: 80   AzureTagName: myTagName AzureTagValue:
      myTagValue AzurePollPeriod: 60   GraceFul: NO
      Delay: 60

```

To remove a cloud profile, type `rm cloud profile <cloud profile name>`

Example 3: The following command removes the cloud profile created in example 1.

```

1 > rm cloudprofile MyTagCloudProfile
2 Done

```

Troubleshooting

Issue: In very rare cases, the “rm cloud profile” CLI command might fail to remove service group and servers associated with the deleted cloud profile. This happens when the command is issued seconds before the poll period of the cloud profile being deleted elapses.

Solution: Manually delete the remaining service groups by entering the following CLI command for each of the remaining service groups:

```

1 #> rm servicegroup <serviceGroupName>

```

Also remove each of the remain servers by entering the following CLI command for each of the remaining servers:

```

1 #> rm server <name>

```

Issue: If you add an Azure tag setting to a VPX instance by using CLI, the `rain_tags` process continues to run on an HA pair node after a warm reboot.

Solution: Manually terminate the process on the secondary node after a warm reboot. From the CLI of the secondary HA node exit to the shell prompt:

```
1 #> shell
```

Use the following command to kill the rain_tags process:

```
1 # PID=`ps -aux | grep rain_tags | awk '{
2   print $2 }
3   `; kill -9 $PID
```

Issue: Back-end servers might not be reachable and reported as DOWN by the VPX instance, in spite of being healthy.

Solution: Make sure that the VPX instance can reach the tagged IP address corresponding to the back-end server. For a tagged NIC, this is the NIC IP address; whereas for a tagged VM, this is the VM's primary IP address. If the VM/NIC resides on a different Azure VNet, make sure that VNet peering is enabled.

Configure GSLB on Citrix ADC VPX instances

November 21, 2024

Citrix ADC appliances configured for global server load balancing (GSLB) provide disaster recovery and continuous availability of applications by protecting against points of failure in a WAN. GSLB can balance the load across data centers by directing client requests to the closest or best performing data center, or to surviving data centers if there is an outage.

This section describes how to enable GSLB on VPX instances on two sites in a Microsoft Azure environment, by using Windows PowerShell commands.

Note:

For more information about GSLB, see [Global Server Load Balancing](#).

You can configure GSLB on a Citrix ADC VPX instance on Azure, in two steps:

1. Create a VPX instance with multiple NICs and multiple IP addresses, on each site.
2. Enable GSLB on the VPX instances.

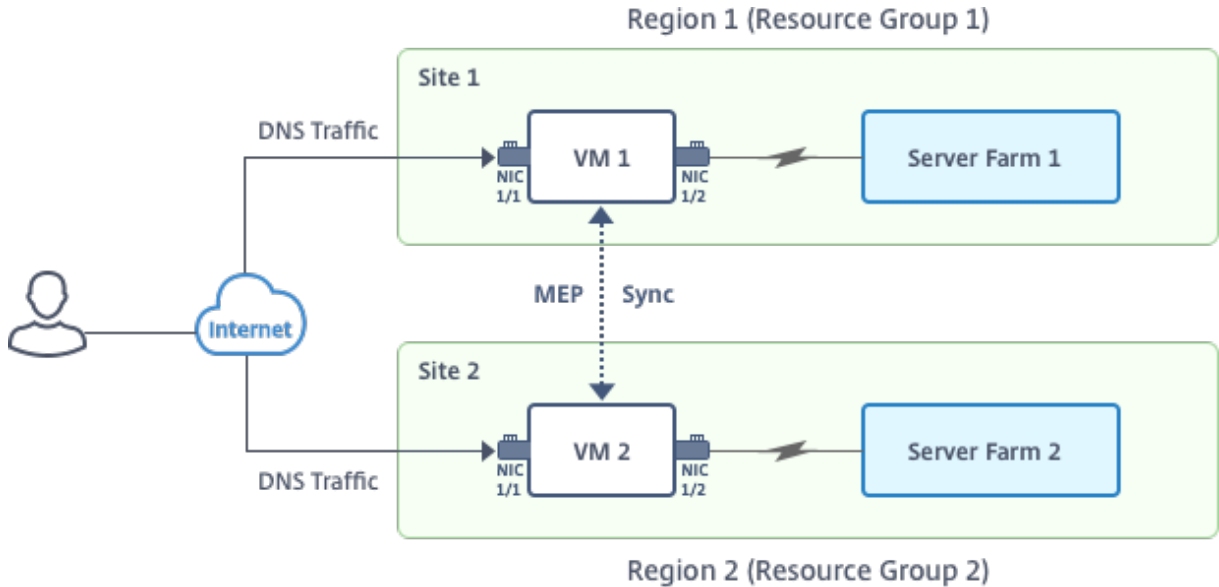
Note:

For more information about configuring multiple NICs and IP addresses see: [Configure multiple IP addresses for a Citrix ADC VPX instance in standalone mode by using PowerShell commands](#)

Scenario

This scenario includes two sites - Site 1 and Site 2. Each site has a VM (VM1 and VM2) configured with multiple NICs, multiple IP addresses, and GSLB.

Figure. GSLB setup implemented across two sites - Site 1 and Site 2.



In this scenario, each VM has three NICs - NIC 0/1, 1/1, and 1/2. Each NIC can have multiple private and public IP addresses. The NICs are configured for the following purposes.

- NIC 0/1: to serve management traffic
- NIC 1/1: to serve client-side traffic
- NIC 1/2: to communicate with back-end servers

For information about the IP addresses configured on each NIC in this scenario, see the IP configuration details section.

Parameters

Following are sample parameters settings for this scenario in this document.

```

1 $location="West Central US"
2
3 $vnetName="NSVPX-vnet"
4
5 $RGName="multiIP-RG"
6
7 $prmStorageAccountName="multiipstorageacctnt"
8
9 $avSetName="MultiIP-avset"
10

```

```
11 $vmSize="Standard\_DS3\_V2"
```

Note:

The minimum requirement for a VPX instance is 2 vCPUs and 2 GB RAM.

```
1 $publisher="citrix"
2
3 $offer="netscalervpx111"
4
5 $sku="netscalerbyol"
6
7 $version="latest"
8
9 $vmNamePrefix="MultiIPVPX"
10
11 $nicNamePrefix="MultiipVPX"
12
13 $osDiskSuffix="osdiskdb"
14
15 $numberOfVMs=1
16
17 $ipAddressPrefix="10.0.0."
18
19 $ipAddressPrefix1="10.0.1."
20
21 $ipAddressPrefix2="10.0.2."
22
23 $pubIPName1="MultiIP-pip1"
24
25 $pubIPName2="MultiIP-pip2"
26
27 $IPConfigName1="IPConfig1"
28
29 $IPConfigName2="IPConfig-2"
30
31 $IPConfigName3="IPConfig-3"
32
33 $IPConfigName4="IPConfig-4"
34
35 $frontendSubnetName="default"
36
37 $backendSubnetName1="subnet\_1"
38
39 $backendSubnetName2="subnet\_2"
40
41 $suffixNumber=10
```

Create a VM

Follow steps 1–10 to create VM1 with multiple NICs and multiple IP addresses, by using PowerShell commands:

1. [Create resource group](#)
2. [Create storage account](#)
3. [Create availability set](#)
4. [Create virtual network](#)
5. [Create public IP address](#)
6. [Create NICs](#)
7. [Create VM config object](#)
8. [Get credentials and set OS properties for the VM](#)
9. [Add NICs](#)
10. [Specify OS disk and create VM](#)

After you complete all the steps and commands to create VM1, repeat these steps to create VM2 with parameters specific to it.

Create resource group

```
1 New-AzureRMResourceGroup -Name $RGName -Location $location
```

Create storage account

```
1 $prmStorageAccount=New-AzureRMStorageAccount -Name
  $prmStorageAccountName -ResourceGroupName $RGName -Type Standard_LRS
  -Location $location
```

Create availability set

```
1 $avSet=New-AzureRMAvailabilitySet -Name $avSetName -ResourceGroupName
  $RGName -Location $location
```

Create virtual network

1. Add subnets.

```

1 $subnet1=New-AzureRmVirtualNetworkSubnetConfig -Name
   $frontendSubnetName -AddressPrefix "10.0.0.0/24"
2 $subnet2=New-AzureRmVirtualNetworkSubnetConfig -Name
   $backendSubnetName1 -AddressPrefix "10.0.1.0/24"
3 $subnet3=New-AzureRmVirtualNetworkSubnetConfig -Name
   $backendSubnetName2 -AddressPrefix "10.0.2.0/24"

```

2. Add virtual network object.

```

1 $vnet=New-AzureRmVirtualNetwork -Name $vnetName -ResourceGroupName
   $RGName -Location $location -AddressPrefix 10.0.0.0/16 -Subnet
   $subnet1, $subnet2, $subnet3

```

3. Retrieve subnets.

```

1 $frontendSubnet=$vnet.Subnets|?{
2   $_.Name -eq $frontendSubnetName }
3
4 $backendSubnet1=$vnet.Subnets|?{
5   $_.Name -eq $backendSubnetName1 }
6
7 $backendSubnet2=$vnet.Subnets|?{
8   $_.Name -eq $backendSubnetName2 }

```

Create public IP address

```

1 $pip1=New-AzureRmPublicIpAddress -Name $pubIPName1 -ResourceGroupName
   $RGName -Location $location -AllocationMethod Dynamic
2 $pip2=New-AzureRmPublicIpAddress -Name $pubIPName2 -ResourceGroupName
   $RGName -Location $location -AllocationMethod Dynamic

```

Create NICs

Create NIC 0/1

```

1 $nic1Name=$nicNamePrefix + $suffixNumber + "-Mgmt"
2 $ipAddress1=$ipAddressPrefix + $suffixNumber
3 $IPConfig1=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName1 -
   SubnetId $frontendSubnet.Id -PublicIpAddress $pip1 -PrivateIpAddress
   $ipAddress1 -Primary
4 $nic1=New-AzureRMNetworkInterface -Name $nic1Name -ResourceGroupName
   $RGName -Location $location -IpConfiguration $IpConfig1

```

Create NIC 1/1

```

1 $nic2Name $nicNamePrefix + $suffixNumber + "--frontend"
2 $ipAddress2=$ipAddressPrefix1 + ($suffixNumber)
3 $ipAddress3=$ipAddressPrefix1 + ($suffixNumber + 1)
4 $IPConfig2=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName2 -
    PublicIpAddress $pip2 -SubnetId $backendSubnet1.Id -
    PrivateIpAddress $ipAddress2 -Primary
5 $IPConfig3=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName3 -
    SubnetId $backendSubnet1.Id -PrivateIpAddress $ipAddress3
6 nic2=New-AzureRMNetworkInterface -Name $nic2Name -ResourceGroupName
    $RGName -Location $location -IpConfiguration $IpConfig2, $IpConfig3

```

Create NIC 1/2

```

1 $nic3Name=$nicNamePrefix + $suffixNumber + "--backend"
2 $ipAddress4=$ipAddressPrefix2 + ($suffixNumber)
3 $IPConfig4=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName4 -
    SubnetId $backendSubnet2.Id -PrivateIpAddress $ipAddress4 -Primary
4 $nic3=New-AzureRMNetworkInterface -Name $nic3Name -ResourceGroupName
    $RGName -Location $location -IpConfiguration $IpConfig4

```

Create VM config object

```

1 $vmName=$vmNamePrefix
2 $vmConfig=New-AzureRMVMConfig -VMName $vmName -VMSize $vmSize -
    AvailabilitySetId $avSet.Id

```

Get credentials and set OS properties

```

1 $cred=Get-Credential -Message "Type the name and password for VPX login
    ."
2 $vmConfig=Set-AzureRMVMOperatingSystem -VM $vmConfig -Linux -
    ComputerName $vmName -Credential $cred
3 $vmConfig=Set-AzureRMVMSourceImage -VM $vmConfig -PublisherName
    $publisher -Offer $offer -Skus $sku -Version $version

```

Add NICs

```

1 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic1.Id -
    Primary
2 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic2.Id
3 $vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic3.Id

```

Specify OS disk and create VM


```

1 $osDiskName=$vmName + "-" + $osDiskSuffix
2 $osVhdUri=$prmStorageAccount.PrimaryEndpoints.Blob.ToString() + "vhds/"
  + $osDiskName + ".vhd"
3 $vmConfig=Set-AzureRMVMOsdisk -VM $vmConfig -Name $osDiskName -VhdUri
  $osVhdUri -CreateOption fromImage
4 Set-AzureRmVMPlan -VM $vmConfig -Publisher $publisher -Product $offer -
  Name $sku
5 New-AzureRMVM -VM $vmConfig -ResourceGroupName $RGName -Location
  $location

```

Note:

Repeat steps 1–10 listed in “Create Multi-NIC VMs by Using PowerShell Commands” to create VM2 with parameters specific to VM2.

IP configuration details

The following IP addresses are used.

Table 1. IP addresses used in VM1

NIC	Private IP	Public IP (PIP)	Description
0/1	10.0.0.10	PIP1	Configured as NSIP (management IP)
1/1	10.0.1.10	PIP2	Configured as SNIP/GSLB Site IP
-	10.0.1.11	-	Configured as LB server IP. Public IP is not mandatory
1/2	10.0.2.10	-	Configured as SNIP for sending monitor probes to services; public IP is not mandatory

Table 2. IP addresses used in VM2

NIC	Internal IP	Public IP (PIP)	Description
0/1	20.0.0.10	PIP4	Configured as NSIP (management IP)

NIC	Internal IP	Public IP (PIP)	Description
1/1	20.0.1.10	PIP5	Configured as SNIP/GSLB Site IP
-	20.0.1.11	-	Configured as LB server IP. Public IP is not mandatory
1/2	20.0.2.10	-	Configured as SNIP for sending monitor probes to services; public IP is not mandatory

Here are sample configurations for this scenario, showing the IP addresses and initial LB configurations as created through the Citrix ADC VPX CLI for VM1 and VM2.

Here's an example configuration on VM1.

```

1 add ns ip 10.0.1.10 255.255.255.0 -mgmtAccess ENABLED
2 Add nsip 10.0.2.10 255.255.255.0
3 add service svc1 10.0.1.10 ADNS 53
4 add lb vserver v1 HTTP 10.0.1.11 80
5 add service s1 10.0.2.120 http 80
6 Add service s2 10.0.2.121 http 80
7 Bind lb vs v1 s[1-2]
```

Here's an example configuration on VM2.

```

1 add ns ip 20.0.1.10 255.255.255.0 -mgmtAccess ENABLED
2 Add nsip 20.0.2.10 255.255.255.0
3 add service svc1 20.0.1.10 ADNS 53
4 add lb vserver v1 HTTP 20.0.1.11 80
5 Add service s1 20.0.2.90 http 80
6 Add service s2 20.0.2.91 http 80
7 Bind lb vs v1 s[1-2]
```

Configure GSLB sites and other settings

Perform the tasks described in the following topic to configure the two GSLB sites and other necessary settings:

[Global Server Load Balancing](#)

Here's an example GSLB configuration on VM1 and VM2.

```

1 enable ns feature LB GSLB
```

```
2 add gslb site site1 10.0.1.10 -publicIP PIP2
3 add gslb site site2 20.0.1.10 -publicIP PIP5
4 add gslb service site1_gslb_http_svc1 10.0.1.11 HTTP 80 -publicIP PIP3
  -publicPort 80 -siteName site1
5 add gslb service site2_gslb_http_svc1 20.0.1.11 HTTP 80 -publicIP PIP6
  -publicPort 80 -siteName site2
6 add gslb vserver gslb_http_vip1 HTTP
7 bind gslb vserver gslb_http_vip1 -serviceName site2_gslb_http_svc1
8 bind gslb vserver gslb_http_vip1 -serviceName site1_gslb_http_svc1
9 bind gslb vserver gslb_http_vip1 -domainName www.gslbindia.com -TTL 5
```

You've configured GSLB on Citrix ADC VPX instances running on Azure.

Configure GSLB on an active-standby high-availability setup

September 9, 2024

You can configure global server load balancing (GSLB) on active-standby HA deployment on Azure in three steps:

1. Create a VPX HA pair on each GSLB site. See [Configure a high-availability setup with multiple IP addresses and NICs](#) for information about how to create an HA pair.
2. Configure the Azure Load Balancer (ALB) with the front-end IP address and rules to allow GSLB and DNS traffic.

This step involves the following substeps. See the scenario in this section for the PowerShell commands used to complete these substeps.

- a. Create a front-end `IPconfig` for GSLB site.
- b. Create a back-end address pool with IP address of NIC 1/1 of nodes in HA.
- c. Create load-balancing rules for following:

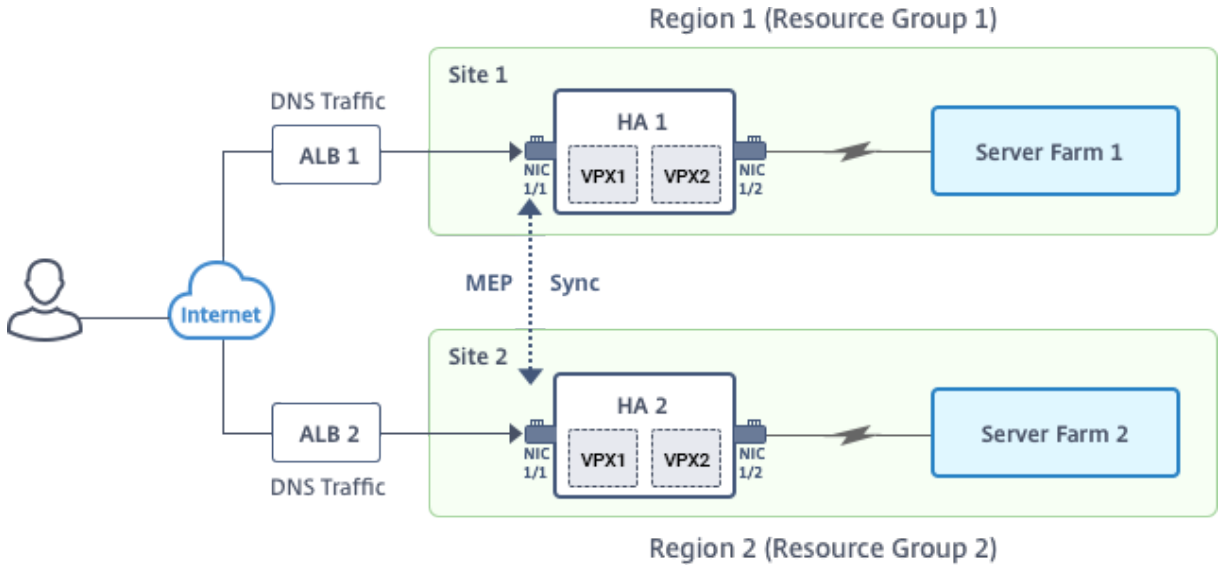
```
1 TCP/3009 - gslb communication
2 TCP/3008 - gslb communication
3 UDP/53 - DNS communication
```

- d. Associate back-end address pool with the LB rules created in step c.
 - e. Update the network security group of NIC 1/1 of nodes in both the HA pair to allow the traffic for TCP 3008, TCP 3009 and UDP 53 ports.
3. Enable GSLB on each HA pair.

Scenario

This scenario includes two sites - Site 1 and Site 2. Each site has an HA pair (HA1 and HA2) configured with multiple NICs, multiple IP addresses, and GSLB.

Figure: GLSB on Active-Standy HA Deployment on Azure



In this scenario, each VM has three NICs - NIC 0/1, 1/1, and 1/2. The NICs are configured for the following purposes.

NIC 0/1: to serve management traffic

NIC 1/1: to serve client-side traffic

NIC 1/2: to communicate with back-end servers

Parameter Settings

Following are sample parameters settings for the ALB. You can use different settings if you want.

```

1 $locName="South east Asia"
2
3 $rgName="MulitIP-MultiNIC-RG"
4
5 $pubIPName4="PIPFORGSLB1"
6
7 $domName4="vpxgslbdns"
8
9 $lbName="MultiIPALB"
10
11 $frontEndConfigName2="FrontEndIP2"
12
13 $backendPoolName1="BackendPoolHttp"
    
```

```

14
15 $lbRuleName2="LBRuleGSLB1"
16
17 $lbRuleName3="LBRuleGSLB2"
18
19 $lbRuleName4="LBRuleDNS"
20
21 $healthProbeName="HealthProbe"

```

Configure ALB with the front-end IP address and rules to allow GSLB and DNS traffic

Step 1. Create a public IP for GSLB site IP

```

1 $pip4=New-AzureRmPublicIpAddress -Name $pubIPName4 -ResourceGroupName
  $rgName -DomainNameLabel $domName4 -Location $locName -
  AllocationMethod Dynamic
2
3
4 Get-AzureRmLoadBalancer -Name \$lbName -ResourceGroupName \$rgName |
  Add-AzureRmLoadBalancerFrontendIpConfig -Name \$frontEndConfigName2
  -PublicIpAddress \$pip4 | Set-AzureRmLoadBalancer

```

Step 2. Create LB rules and update the existing ALB.

```

1 $alb = get-AzureRmLoadBalancer -Name $lbName -ResourceGroupName $rgName
2
3
4 $frontendipconfig2=Get-AzureRmLoadBalancerFrontendIpConfig -
  LoadBalancer $alb -Name $frontEndConfigName2
5
6
7 $backendPool=Get-AzureRmLoadBalancerBackendAddressPoolConfig -
  LoadBalancer $alb -Name $backendPoolName1
8
9
10 $healthprobe=Get-AzureRmLoadBalancerProbeConfig -LoadBalancer $alb -
  Name $healthProbeName
11
12
13 \$alb | Add-AzureRmLoadBalancerRuleConfig -Name \$lbRuleName2 -
  BackendAddressPool \$backendPool -FrontendIPConfiguration \
  $frontendipconfig2 -Protocol \"Tcp\" -FrontendPort 3009 -BackendPort
  3009 -Probe \$healthprobe -EnableFloatingIP | Set-
  AzureRmLoadBalancer
14
15
16 \$alb | Add-AzureRmLoadBalancerRuleConfig -Name \$lbRuleName3 -
  BackendAddressPool \$backendPool -FrontendIPConfiguration \
  $frontendipconfig2 -Protocol \"Tcp\" -FrontendPort 3008 -BackendPort
  3008 -Probe \$healthprobe -EnableFloatingIP | Set-
  AzureRmLoadBalancer

```

```

17
18
19 \ $alb | Add-AzureRmLoadBalancerRuleConfig -Name \$lbRuleName4 -
    BackendAddressPool \$backendPool -FrontendIPConfiguration \
    $frontendipconfig2 -Protocol \"Udp\" -FrontendPort 53 -BackendPort
    53 -Probe \$healthprobe -EnableFloatingIP | Set-AzureRmLoadBalancer

```

Enable GSLB on each high availability pair

Now you've two front-end IP addresses for each ALB: ALB 1 and ALB 2. One IP address is for the LB virtual server and the other for the GSLB site IP.

HA 1 has the following front-end IP addresses:

- FrontEndIPofALB1 (for LB virtual server)
- PIPFORGSLB1 (GSLB IP)

HA 2 has the following front-end IP addresses:

- FrontEndIPofALB2 (for LB virtual server)
- PIPFORGSLB2 (GSLB IP)

The following commands are used for this scenario.

```

1 enable ns feature LB GSLB
2
3 add service dnssvc PIPFORGSLB1 ADNS 53
4
5 add gslb site site1 PIPFORGSLB1 -publicIP PIPFORGSLB1
6
7 add gslb site site2 PIPFORGSLB2 -publicIP PIPFORGSLB2
8
9 add gslb service site1_gslb_http_svc1 FrontEndIPofALB1 HTTP 80 -
    publicIP FrontEndIPofALB1 -publicPort 80 -siteName site1
10
11 add gslb service site2_gslb_http_svc1 FrontEndIPofALB2 HTTP 80 -
    publicIP FrontEndIPofALB2 -publicPort 80 -siteName site2
12
13 add gslb vserver gslb_http_vip1 HTTP
14
15 bind gslb vserver gslb_http_vip1 -serviceName site2_gslb_http_svc1
16
17 bind gslb vserver gslb_http_vip1 -serviceName site1_gslb_http_svc1
18
19 bind gslb vserver gslb_http_vip1 -domainName www.gslbindia.com -TTL 5

```

Related resources:

[Configure GSLB on Citrix ADC VPX instances](#)

[Global Server Load Balancing](#)

Configure address pools intranet IP for a Citrix Gateway appliance

September 18, 2024

In some situations, users who connect with the Citrix Gateway Plug-in need a unique IP address for a Citrix ADC Gateway appliance. When you enable address pools (also known as IP pooling) for a group, the Citrix Gateway appliance can assign a unique IP address alias to each user. You configure address pools by using intranet IP (IIP) addresses.

You can configure address pools on a Citrix Gateway appliance deployed on Azure by following this 2-step procedure:

- Registering the private IP addresses that are used in the address pool, in Azure
- Configuring address pools in the Citrix Gateway appliance

Register a private IP address in the Azure portal

In Azure, you can deploy a Citrix ADC VPX instance with multiple IP addresses. You can add IP addresses to a VPX instance in two ways:


a. While provisioning a VPX instance

For more information about how to add multiple IP addresses while provisioning a VPX instance, see [Configure multiple IP addresses for a Citrix ADC standalone instance](#). To add IP addresses by using PowerShell commands while provisioning a VPX instance, see [Configure multiple IP addresses for a Citrix ADC VPX instance in standalone mode by using PowerShell commands](#).


b. After provisioning a VPX instance

After you've provisioned a VPX instance, follow these steps to register a private IP address in the Azure portal, which you configure as an address pool in the Citrix Gateway appliance.

1. From Azure Resource Manager (ARM), go to the already created Citrix ADC VPX instance > **Network interfaces**. Choose the network interface which is bound to a subnet to which the IIP that you want to register belongs.



2. Click **IP Configurations**, and then click **Add**.



3. Provide the required details as shown in the example below and click **OK**.

Add IP configuration
nsdoc0330vm923

* Name
PrivateIP5 ✓

Type
Primary Secondary

Primary IP configuration already exists

Private IP address settings

Allocation
Dynamic Static

* IP address
192.0.0.8 ✓

Public IP address
Disabled Enabled

OK

Configure address pools in the Citrix Gateway appliance

For more information about how to configure address pools on the Citrix Gateway, see [Configuring Address Pools](#).

Limitation:

You cannot bind a range of IIP addresses to users. Every IIP address that is used in an address pool must be registered.

Configure multiple IP addresses for a Citrix ADC VPX standalone instance by using PowerShell commands

November 21, 2024

In an Azure environment, a Citrix ADC VPX virtual appliance can be deployed with multiple NICs. Each NIC can have multiple IP addresses. This section describes how to deploy a Citrix ADC VPX instance with a single NIC and multiple IP addresses, by using PowerShell commands. You can use the same script for multi-NIC and multi-IP deployment.

Note:

In this document, IP-Config refers to a pair of IP addresses, public IP, and private IP, that is associated with an individual NIC. For more information, see the [Azure terminology](#) section.

Use case

In this use case, a single NIC is connected to a virtual network (VNET). The NIC is associated with three IP configurations, as shown in the following table.


IP Config	Associated with
IPConfig-1	Static public IP address; static private IP address
IPConfig-2	Static public IP address; static private address
IPConfig-3	Static private IP address

Note:

IPConfig-3 is not associated with any public IP address.

Diagram: Topology

Here is the visual representation of the use case.

**Note:**

In a multi-NIC, multi-IP Azure Citrix ADC VPX deployment, the private IP address associated with the primary (first) `IPConfig` of the primary (first) NIC is automatically added as the management NSIP address of the appliance. The remaining private IP addresses associated with `IPConfigs` must be added in the VPX instance as VIPs or SNIPs by using the `add ns ip` command, as determined by your requirements.

Here is the summary of the steps required for configuring multiple IP addresses for a Citrix ADC VPX virtual appliance in standalone mode:

1. Create Resource Group
2. Create Storage Account
3. Create Availability Set
4. Create Network service group
5. Create Virtual Network
6. Create Public IP Address
7. Assign IP Configuration
8. Create NIC
9. Create Citrix ADC VPX Instance

10. Check NIC Configurations
11. Check VPX-side Configurations

Script

Parameters

Following are sample parameters settings for the use case in this document.

```
1 $locName="westcentralus"
2
3 $rgName="Azure-MultiIP"
4
5 $nicName1="VM1-NIC1"
6
7 $vNetName="Azure-MultiIP-vnet"
8
9 $vNetAddressRange="11.6.0.0/16"
10
11 $frontEndSubnetName="frontEndSubnet"
12
13 $frontEndSubnetRange="11.6.1.0/24"
14
15 $prmStorageAccountName="multiipstorage"
16
17 $avSetName="multiip-avSet"
18
19 $vmSize="Standard_DS4_V2" (This parameter creates a VM with up to
    four NICs.)
```

Note:

The minimum requirement for a VPX instance is 2 vCPUs and 2 GB RAM.

```
1 $publisher="Citrix"
2
3 $offer="netscalervpx110-6531" (You can use different offers.)
4
5 $sku="netscalerbyol" (According to your offer, the SKU can be different
    .)
6
7 $version="latest"
8
9 $pubIPName1="PIP1"
10
11 $pubIPName2="PIP2"
12
13 $domName1="multiipvpx1"
14
15 $domName2="multiipvpx2"
```

```
16
17 $vmNamePrefix="VPXMultiIP"
18
19 $osDiskSuffix="osmultiipalbdiskdb1"
20
21 **Network Security Group (NSG)-related information**:
22
23 $nsgName="NSG-MultiIP"
24
25 $rule1Name="Inbound-HTTP"
26
27 $rule2Name="Inbound-HTTPS"
28
29 $rule3Name="Inbound-SSH"
30
31 $IpConfigName1="IPConfig1"
32
33 $IPConfigName2="IPConfig-2"
34
35 $IPConfigName3="IPConfig-3"
```

1. Create Resource Group

```
New-AzureRmResourceGroup -Name $rgName -Location $locName
```

2. Create Storage Account

```
$prmStorageAccount = New-AzureRMStorageAccount -Name $prmStorageAccountName  
-ResourceGroupName $rgName -Type Standard_LRS -Location $locName
```

3. Create Availability Set

```
$avSet = New-AzureRMAvailabilitySet -Name $avSetName -ResourceGroupName  
$rgName -Location $locName
```

4. Create Network Security Group

1. Add rules. You must add a rule to the network security group for any port that serves traffic.

```
$rule1=New-AzureRmNetworkSecurityRuleConfig -Name $rule1Name -  
Description "Allow HTTP"-Access Allow -Protocol Tcp -Direction  
Inbound -Priority 101 -SourceAddressPrefix Internet -SourcePortRange  
* -DestinationAddressPrefix * -DestinationPortRange 80
```

```
$rule2=New-AzureRmNetworkSecurityRuleConfig -Name $rule2Name -
Description "Allow HTTPS"-Access Allow -Protocol Tcp -Direction
Inbound -Priority 110 -SourceAddressPrefix Internet -SourcePortRange
* -DestinationAddressPrefix * -DestinationPortRange 443
$rule3=New-AzureRmNetworkSecurityRuleConfig -Name $rule3Name
-Description "Allow SSH"-Access Allow -Protocol Tcp -Direction
Inbound -Priority 120 -SourceAddressPrefix Internet -SourcePortRange
* -DestinationAddressPrefix * -DestinationPortRange 22
```

2. Create network security group object.

```
$nsg=New-AzureRmNetworkSecurityGroup -ResourceGroupName $rgName
-Location $locName -Name $nsgName -SecurityRules $rule1,$rule2,
$rule3
```

5. Create Virtual Network

1. Add subnets.

```
$frontendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
$frontEndSubnetName -AddressPrefix $frontEndSubnetRange
```

2. Add virtual network object.

```
$vnet=New-AzureRmVirtualNetwork -Name $vNetName -ResourceGroupName
$rgName -Location $locName -AddressPrefix $vNetAddressRange -
Subnet $frontendSubnet
```

3. Retrieve subnets.

```
$subnetName="frontEndSubnet"
$subnet1=$vnet.Subnets|?{ $_.Name -eq $subnetName }
```

6. Create Public IP Address

```
$pip1=New-AzureRmPublicIpAddress -Name $pubIPName1 -ResourceGroupName
$rgName -DomainNameLabel $domName1 -Location $locName -AllocationMethod
Static
$pip2=New-AzureRmPublicIpAddress -Name $pubIPName2 -ResourceGroupName
$rgName -DomainNameLabel $domName2 -Location $locName -AllocationMethod
Static
```

Note:

Check availability of domain names before using.

Allocation method for IP addresses can be dynamic or static.

7. Assign IP Configuration

In this use case, consider the following points before assigning IP addresses:

- IPConfig-1 belongs to subnet1 of VPX1.
- IPConfig-2 belongs to subnet 1 of VPX1.
- IPConfig-3 belongs to subnet 1 of VPX1.

Note:

When you assign multiple IP configurations to a NIC, one configuration must be assigned as primary.

```
1 $IPAddress1="11.6.1.27"
2 $IPConfig1=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName1 -
    Subnet $subnet1 -PrivateIpAddress $IPAddress1 -PublicIpAddress $pip1
    - Primary
3 $IPAddress2="11.6.1.28"
4 $IPConfig2=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName2 -
    Subnet $subnet1 -PrivateIpAddress $IPAddress2 -PublicIpAddress $pip2
5 $IPAddress3="11.6.1.29"
6 $IPConfig3=New-AzureRmNetworkInterfaceIpConfig -Name $IPConfigName3 -
    Subnet $subnet1 -PrivateIpAddress $IPAddress3 -Primary
```

Use a valid IP address that meets your subnet requirements and check its availability.

8. Create NIC

```
$nic1=New-AzureRmNetworkInterface -Name $nicName1 -ResourceGroupName
    $rgName -Location $locName -IpConfiguration $IpConfig1,$IpConfig2,
    $IPConfig3 -NetworkSecurityGroupId $nsg.Id
```

9. Create Citrix ADC VPX Instance

1. Initialize variables.

```
$suffixNumber = 1
$vmName = $vmNamePrefix + $suffixNumber
```

2. Create VM config object.

```
$vmConfig=New-AzureRMVMConfig -VMName $vmName -VMSize $vmSize -
AvailabilitySetId $avSet.Id
```

3. Set credentials, OS, and image.

```
$cred=Get-Credential -Message "Type the name and password for VPX
login."
$vmConfig=Set-AzureRMVMOperatingSystem -VM $vmConfig -Linux -
ComputerName $vmName -Credential $cred
$vmConfig=Set-AzureRMVMSourceImage -VM $vmConfig -PublisherName
$publisher -Offer $offer -Skus $sku -Version $version
```

4. Add NIC.

```
$vmConfig=Add-AzureRMVMNetworkInterface -VM $vmConfig -Id $nic1.
Id -Primary
```

Note:

In a multi-NIC VPX deployment, one NIC must be primary. So, “-Primary” must be appended while adding that NIC to the VPX instance.

5. Specify OS disk and create VM.

```
$osDiskName=$vmName + "-" + $osDiskSuffix1
$osVhdUri=$prmStorageAccount.PrimaryEndpoints.Blob.ToString()+ "
vhds/" + $osDiskName + ".vhd"
$vmConfig=Set-AzureRMVMOSDisk -VM $vmConfig -Name $osDiskName -
VhdUri $osVhdUri -CreateOption fromImage
Set-AzureRmVMPlan -VM $vmConfig -Publisher $publisher -Product
$offer -Name $sku
New-AzureRMVM -VM $vmConfig -ResourceGroupName $rgName -Location
$locName
```

10. Check NIC Configurations

After the VPX instance starts, you can check the IP addresses allocated to `IPConfigs` of the VPX NIC by using the following command.

```
$nic.IPConfig
```


11. Check VPX-side Configurations

When the Citrix ADC VPX instance starts, a private IP address associated with primary `IPconfig` of the primary NIC is added as the NSIP address. The remaining private IP addresses must be added as VIP or SNIP addresses, as determined by your requirements. Use the following command.

```
add nsip <Private IPAddress><netmask> -type VIP/SNIP
```

You've now configured multiple IP addresses for a Citrix ADC VPX instance in standalone mode.

Additional PowerShell scripts for Azure deployment

September 19, 2024

This section provides the PowerShell cmdlets with which you can perform the following configurations in Azure PowerShell:

- Provision a Citrix ADC VPX standalone instance
- Provision a Citrix ADC VPX pair in a high availability setup with an Azure external load balancer
- Provision a Citrix ADC VPX pair in a high availability setup with Azure internal load balancer

Also see the following topics for configurations that you can perform by using PowerShell commands:

- [Configure a high-availability setup with multiple IP addresses and NICs by using PowerShell commands](#)
- [Configure GSLB on Citrix ADC VPX instances](#)
- [Configure GSLB on a NetScaler active-standby high-availability setup](#)
- [Configure multiple IP addresses for a Citrix ADC VPX instance in standalone mode by using PowerShell commands](#)

Provision a Citrix ADC VPX standalone instance

1. Create a resource group

The resource group can include all the resources for the solution, or only those resources that you want to manage as a group. The location specified here is the default location for resources in that resource group. Make sure all commands to create a load balancer use the same resource group.

```
$rgName="<resource group name>"  
$locName="<location name, such as West US>"  
New-AzureRmResourceGroup -Name $rgName -Location $locName
```

For example:

```
1 $rgName = "ARM-VPX"
2 $locName = "West US"
3 New-AzureRmResourceGroup -Name $rgName -Location $locName
```

2. Create a storage account

Choose a unique name for your storage account that contains only lowercase letters and numbers.

```
$saName="<storage account name>"
$saType="<storage account type>", specify one: Standard_LRS, Standard_GRS
, Standard_RAGRS, or Premium_LRS
New-AzureRmStorageAccount -Name $saName -ResourceGroupName
$rgName -Type $saType -Location $locName
```

For example:

```
1 $saName="vpxstorage"
2 $saType="Standard_LRS"
3 New-AzureRmStorageAccount -Name $saName -ResourceGroupName $rgName
- Type $saType -Location $locName
```

3. Create an availability set

Availability set helps to keep your virtual machines available during downtime, such as during maintenance. A load balancer configured with an availability set ensures that your application is always available.

```
$avName="<availability set name>"
New-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
$rgName -Location $locName
```

4. Create a virtual network

Add a new virtual network with at least one subnet, if the subnet was not created previously.

```
$FrontendAddressPrefix="10.0.1.0/24"
$BackendAddressPrefix="10.0.2.0/24"
$VnetAddressPrefix="10.0.0.0/16"
$frontendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
frontendSubnet -AddressPrefix $FrontendAddressPrefix
$backendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
backendSubnet -AddressPrefix $BackendAddressPrefix
New-AzureRmVirtualNetwork -Name TestNet -ResourceGroupName
$rgName -Location $locName -AddressPrefix $VnetAddressPrefix -
Subnet $frontendSubnet,$backendSubnet
```

For example:

```

1 $frontendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
  frontendSubnet -AddressPrefix $FrontendAddressPrefix
2
3 $backendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
  backendSubnet -AddressPrefix $BackendAddressPrefix
4
5 New-AzureRmVirtualNetwork -Name TestNet -ResourceGroupName $rgName
  -Location $locName -AddressPrefix $vnetAddressPrefix -Subnet
  $frontendSubnet,$backendSubnet

```

5. Create a NIC

Create a NIC and associate the NIC with the Citrix ADC VPX instance. The front end Subnet created in the above procedure is indexed at 0 and the back end Subnet is indexed at 1. Now create NIC in one of the three following ways:

a) NIC with Public IP address

```

$nicName="<name of the NIC of the VM>"
$pip = New-AzureRmPublicIpAddress -Name $nicName -ResourceGroupName
  $rgName -Location $locName -AllocationMethod Dynamic
$nic = New-AzureRmNetworkInterface -Name $nicName -ResourceGroupName
  $rgName -Location $locName -SubnetId $vnet.Subnets[$subnetIndex
].Id -PublicIpAddressId $pip.Id

```

b) NIC with Public IP and DNS label

```

$nicName="<name of the NIC of the VM>"
$domName="<domain name label>"
$pip = New-AzureRmPublicIpAddress -Name $nicName -ResourceGroupName
  $rgName -DomainNameLabel $domName -Location $locName -AllocationMethod
  Dynamic

```

Before assigning \$domName, check it is available or not by using command:

```

Test-AzureRmDnsAvailability -DomainQualifiedName $domName -
  Location $locName
$nic = New-AzureRmNetworkInterface -Name $nicName -ResourceGroupName
  $rgName -Location $locName -SubnetId $vnet.Subnets[$subnetIndex
].Id -PublicIpAddressId $pip.Id

```

For example:

```

1 $nicName="frontendNIC"
2

```

```

3 $domName="vpxazure"
4
5 $pip = New-AzureRmPublicIpAddress -Name $nicName -
    ResourceGroupName $rgName -DomainNameLabel $domName -Location
    $locName -AllocationMethod Dynamic
6
7 $nic = New-AzureRmNetworkInterface -Name $nicName -
    ResourceGroupName $rgName -Location $locName -SubnetId $vnet.
    Subnets\[0\].Id -PublicIpAddressId $pip.Id

```

c) NIC with Dynamic Public Address and Static Private IP address

Make sure that the private (static) IP address you add to the VM must be the same range as that of the subnet specified.

```
$nicName="<name of the NIC of the VM>"
```

```
$staticIP="<available static IP address on the subnet>"
```

```
$pip = New-AzureRmPublicIpAddress -Name $nicName -ResourceGroupName
    $rgName -Location $locName -AllocationMethod Dynamic
```

```
$nic = New-AzureRmNetworkInterface -Name $nicName -ResourceGroupName
    $rgName -Location $locName -SubnetId $vnet.Subnets[$subnetIndex
    ].Id -PublicIpAddressId $pip.Id -PrivateIpAddress $staticIP
```

6. Create a virtual object

```
$vmName="<VM name>"
```

```
$vmSize="<VM size string>"
```

```
$avSet=Get-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
    $rgName
```

```
$vm=New-AzureRmVMConfig -VMName $vmName -VMSize $vmSize -AvailabilitySetId
    $avset.Id
```

7. Get the Citrix ADC VPX image

```
$pubName="<Image publisher name>"
```

```
$offerName="<Image offer name>"
```

```
$skuName="<Image SKU name>"
```

```
$cred=Get-Credential -Message "Type the name and password of the
local administrator account."
```

Provide your credentials that is used to log in into VPX

```
$vm=Set-AzureRmVMOperatingSystem -VM $vm -Linux -ComputerName
    $vmName -Credential $cred -Verbose
```

```
$vm=Set-AzureRmVMSourceImage -VM $vm -PublisherName $pubName -
Offer $offerName -Skus $skuName -Version "latest"
```

```
$vm=Add-AzureRmVMNetworkInterface -VM $vm -Id $nic.Id
```

For example:

```
$pubName="citrix"
```

The following command is used for displaying all offers from Citrix:

```
1 Get-AzureRMVMImageOffer -Location $locName -Publisher $pubName |
   Select Offer
2
3 $offerName="netscalervpx110-6531"
```

The following command is used to know SKU offered by publisher for specific offer name:

```
Get-AzureRMVMImageSku -Location $locName -Publisher $pubName -
Offer $offerName | Select Skus
```

8. Create a virtual machine

```
$diskName="<name identifier for the disk in Azure storage, such
as OSDisk>"
```

For example:

```
1 $diskName="dynamic"
2
3 $pubName="citrix"
4
5 $offerName="netscalervpx110-6531"
6
7 $skuName="netscalerbyol"
8
9 $storageAcc=Get-AzureRmStorageAccount -ResourceGroupName $rgName -
   Name $saName
10
11 $osDiskUri=$storageAcc.PrimaryEndpoints.Blob.ToString() + "vhds/"
   + $diskName + ".vhd"
12
13 $vm=Set-AzureRmVMOSDisk -VM $vm -Name $diskName -VhdUri $osDiskUri
   -CreateOption fromImage
```

When you create VM from Images present in marketplace, use the following command to specify the VM plan:

```
Set-AzureRmVMPlan -VM $vm -Publisher $pubName -Product $offerName
  -Name $skuName
```

```
New-AzureRmVM -ResourceGroupName $rgName -Location $locName -VM
  $vm
```

Provision a Citrix ADC VPX pair in a high availability setup with an Azure external load balancer

Log on to AzureRmAccount using your Azure user credentials.

1. Create a resource group

The location specified here is the default location for resources in that resource group. Make sure that all commands used to create a load balancer use the same resource group.

```
$rgName="<resource group name>"
```

```
$locName="<location name, such as West US>"
```

```
New-AzureRmResourceGroup -Name $rgName -Location $locName
```

For example:

```
1 $rgName = "ARM-LB-NS"
2
3 $locName = "West US"
4
5 New-AzureRmResourceGroup -Name $rgName -Location $locName
```

2. Create a storage account

Choose a unique name for your storage account that contains only lowercase letters and numbers.

```
$saName="<storage account name>"
```

```
$saType="<storage account type>", specify one: Standard_LRS, Standard_GRS, Standard_RAGRS, or Premium_LRS
```

```
New-AzureRmStorageAccount -Name $saName -ResourceGroupName $rgName -Type $saType -Location $locName
```

For example:

```
1 $saName="vpxstorage"
2
3 $saType="Standard_LRS"
4
5 New-AzureRmStorageAccount -Name $saName -ResourceGroupName $rgName -Type $saType -Location $locName
```

3. Create an availability set

A load balancer configured with an availability set ensures that your application is always available.

```
$avName="<availability set name>"
```

```
New-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
$rgName -Location $locName
```

4. Create a virtual network

Add a new virtual network with at least one subnet, if the subnet was not created previously.

```
1 $vnetName = "LBVnet"
2
3 $FrontendAddressPrefix="10.0.1.0/24"
4
5 $BackendAddressPrefix="10.0.2.0/24"
6
7 $vnetAddressPrefix="10.0.0.0/16"
8
9 $frontendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
    frontendSubnet -AddressPrefix $FrontendAddressPrefix
10
11 $backendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
    backendSubnet -AddressPrefix $BackendAddressPrefix
12
13 $vnet=New-AzureRmVirtualNetwork -Name $vnetName -ResourceGroupName
    $rgName -Location $locName -AddressPrefix $vnetAddressPrefix -
    Subnet $frontendSubnet,$backendSubnet
```

Note:

Choose the AddressPrefix parameter value as per your requirement.

Assign front end and back end subnet to the virtual network that you created earlier in this step.

If the front end subnet is the first element of array VNet, subnetId must be \$vnet.Subnets[0].Id.

If the front end subnet is the second element in the array, the subnetId must be \$vnet.Subnets[1].Id, and so on.

5. Configure front end IP address and create back end address pool

Configure a front end IP address for the incoming load balancer network traffic and create a back end address pool to receive the load balanced traffic.

```
1 $pubName="PublicIp1"
2
3 $publicIP1 = New-AzureRmPublicIpAddress -Name $pubName -
    ResourceGroupName $rgName -Location $locName -AllocationMethod
    Static -DomainNameLabel nsvpx
```

Note:

Check for the availability of the value for DomainNameLabel.

```

1 $FIPName = "ELBFIP"
2
3 $frontendIP1 = New-AzureRmLoadBalancerFrontendIpConfig -Name
    $FIPName -PublicIpAddress $publicIP1
4
5 $BEPool = "LB-backend-Pool"
6
7 $beaddresspool1= New-AzureRmLoadBalancerBackendAddressPoolConfig -
    Name $BEPool

```

6. Create a health probe

Create a TCP health probe with port 9000 and interval 5 seconds.

```

1 $healthProbe = New-AzureRmLoadBalancerProbeConfig -Name
    HealthProbe -Protocol Tcp -Port 9000 -IntervalInSeconds 5 -
    ProbeCount 2

```

7. Create a load balancing rule

Create an LB rule for each service that you are load balancing.

For example:

You can use the following example to load balance HTTP service.

```

1 $lbrule1 = New-AzureRmLoadBalancerRuleConfig -Name "HTTP-LB" -
    FrontendIpConfiguration $frontendIP1 -BackendAddressPool
    $beAddressPool1 -Probe $healthProbe -Protocol Tcp -FrontendPort
    80 -BackendPort 80

```

8. Create inbound NAT rules

Create NAT rules for services that you are not load balancing.

For example, when creating an SSH access to a Citrix ADC VPX instance.

Note:

Protocol-FrontEndPort-BackendPort triplet must not be the same for two NAT rules.

```

1 $inboundNATRule1= New-AzureRmLoadBalancerInboundNatRuleConfig -
    Name SSH1 -FrontendIpConfiguration $frontendIP1 -Protocol
    TCP -FrontendPort 22 -BackendPort 22
2
3 $inboundNATRule2= New-AzureRmLoadBalancerInboundNatRuleConfig -
    Name SSH2 -FrontendIpConfiguration $frontendIP1 -Protocol TCP -
    FrontendPort 10022 -BackendPort 22

```

9. Create a load balancer entity

Create the load balancer adding all objects (NAT rules, load balancer rules, probe configurations) together.


```

1 $lbName="ELB"
2
3 $NRPLB = New-AzureRmLoadBalancer -ResourceGroupName $rgName -Name
    $lbName -Location $locName -InboundNatRule $inboundNATRule1,
    $inboundNATRule2 -FrontendIpConfiguration $frontendIP1 -
    LoadBalancingRule $lbrule1 -BackendAddressPool $beAddressPool1
    -Probe $healthProbe

```

10. Create a NIC

Create two NICs and associate each NIC with each VPX instance

a) NIC1 with VPX1

For example:

```

1 $nicName="NIC1"
2
3 $lbName="ELB"
4
5 $bePoolIndex=0
6
7 \* Rule indexes starts from 0.
8
9 $natRuleIndex=0
10
11 $subnetIndex=0
12
13 \* Frontend subnet index
14
15 $lb=Get-AzureRmLoadBalancer -Name $lbName -ResourceGroupName
    $rgName
16
17 $nic1=New-AzureRmNetworkInterface -Name $nicName -
    ResourceGroupName $rgName -Location $locName -Subnet $vnet.
    Subnets\[ $subnetIndex\] -LoadBalancerBackendAddressPool $lb.
    BackendAddressPools\[ $bePoolIndex\] -LoadBalancerInboundNatRule
    $lb.InboundNatRules\[ $natRuleIndex\]

```

b) NIC2 with VPX2

For example:

```

1 $nicName="NIC2"
2
3 $lbName="ELB"
4
5 $bePoolIndex=0
6
7 $natRuleIndex=1
8
9 \* Second Inbound NAT (SSH) rule we need to use
10

```

```

11 ` $subnetIndex=0
12
13 \* Frontend subnet index
14
15 $lb=Get-AzureRmLoadBalancer -Name $lbName -ResourceGroupName
    $rgName
16
17 $nic2=New-AzureRmNetworkInterface -Name $nicName -
    ResourceGroupName $rgName -Location $locName -Subnet $vnet.
    Subnets\[ $subnetIndex\] -LoadBalancerBackendAddressPool $lb.
    BackendAddressPools\[ $bePoolIndex\] -LoadBalancerInboundNatRule
    $lb.InboundNatRules\[ $natRuleIndex\]

```

11. Create Citrix ADC VPX instances

Create two Citrix ADC VPX instances as part of the same resource group and availability set, and attach it to the external load balancer.

a) Citrix ADC VPX instance 1

For example:

```

1  $vmName="VPX1"
2
3  $vmSize="Standard_A3"
4
5  $pubName="citrix"
6
7  $offerName="netscalervpx110-6531"
8
9  $skuName="netscalerbyol"
10
11 $avSet=Get-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
    $rgName
12
13 $vm1=New-AzureRmVMConfig -VMName $vmName -VMSize $vmSize -
    AvailabilitySetId $avset.Id
14
15 $cred=Get-Credential -Message "Type Credentials which will be used
    to login to VPX instance"
16
17 $vm1=Set-AzureRmVMOperatingSystem -VM $vm1 -Linux -ComputerName
    $vmName -Credential $cred -Verbose
18
19 $vm1=Set-AzureRmVMSourceImage -VM $vm1 -PublisherName $pubName -
    Offer $offerName -Skus $skuName -Version "latest"
20
21 $vm1=Add-AzureRmVMNetworkInterface -VM $vm1 -Id $nic1.Id
22
23 $diskName="dynamic"
24
25 $storageAcc=Get-AzureRmStorageAccount -ResourceGroupName $rgName -
    Name $saName
26

```

```

27 $osDiskUri1=$storageAcc.PrimaryEndpoints.Blob.ToString() + "vhds1/
    " + $diskName + ".vhd"
28
29 $vm1=Set-AzureRmVMOSDisk -VM $vm1 -Name $diskName -VhdUri
    $osDiskUri1 -CreateOption fromImage
30
31 Set-AzureRmVMPlan -VM $vm1 -Publisher $pubName -Product $offerName
    -Name $skuName
32
33 New-AzureRmVM -ResourceGroupName $rgName -Location $locName -VM
    $vm1

```

b) Citrix ADC VPX instance 2

For example:

```

1 $vmName="VPX2"
2
3 $vmSize="Standard\_A3"
4
5 $avSet=Get-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
    $rgName
6
7 $vm2=New-AzureRmVMConfig -VMName $vmName -VMSize $vmSize -
    AvailabilitySetId $avset.Id
8
9 $cred=Get-Credential -Message " Type Credentials which will be
    used to login to VPX instance "
10
11 $vm2=Set-AzureRmVMOperatingSystem -VM $vm2 -Linux -ComputerName
    $vmName -Credential $cred -Verbose
12
13 $vm2=Set-AzureRmVMSourceImage -VM $vm2 -PublisherName $pubName -
    Offer $offerName -Skus $skuName -Version "latest"
14
15 $vm2=Add-AzureRmVMNetworkInterface -VM $vm2 -Id $nic2.Id
16
17 $diskName="dynamic"
18
19 $storageAcc=Get-AzureRmStorageAccount -ResourceGroupName $rgName -
    Name $saName
20
21 $osDiskUri1=$storageAcc.PrimaryEndpoints.Blob.ToString() + "vhds2/
    " + $diskName + ".vhd"
22
23 $vm2=Set-AzureRmVMOSDisk -VM $vm2 -Name $diskName -VhdUri
    $osDiskUri1 -CreateOption fromImage
24
25 Set-AzureRmVMPlan -VM $vm2 -Publisher $pubName -Product $offerName
    -Name $skuName
26
27 New-AzureRmVM -ResourceGroupName $rgName -Location $locName -VM
    $vm2

```

12. Configure the virtual machines

When both the Citrix ADC VPX instances start, then connect to both Citrix ADC VPX instances using the SSH protocol to configure the virtual machines.

a) Active-Active: Run the same set of configuration commands on the command line of both the Citrix ADC VPX instances.

b) Active-Passive: Run this command on the command line of both the Citrix ADC VPX instances.

```
add ha node #nodeID <nsip of other Citrix ADC VPX>
```

In Active-Passive mode, run configuration commands on the primary node only.

Provision a Citrix ADC VPX pair in a high availability setup with Azure internal load balancer

Log on to AzureRmAccount using your Azure user credentials.

1. Create a resource group

The location specified here is the default location for resources in that resource group. Make sure all commands to create a load balancer use the same resource group.

```
$rgName="\<resource group name\>"
```

```
$locName="\<location name, such as West US\>"
```

```
New-AzureRmResourceGroup -Name $rgName -Location $locName
```

For example:

```
1 $rgName = "ARM-LB-NS"
2
3 $locName = "West US"
4
5 New-AzureRmResourceGroup -Name $rgName -Location $locName
```

2. Create a storage account

Choose a unique name for your storage account that contains only lowercase letters and numbers.

```
$saName="\<storage account name>"
```

```
$saType="\<storage account type>", specify one: Standard_LRS, Standard_GRS, Standard_RAGRS, or Premium_LRS
```

```
New-AzureRmStorageAccount -Name $saName -ResourceGroupName $rgName -Type $saType -Location $locName
```

For example:

```

1 $saName="vpxstorage"
2
3 $saType="Standard_LRS"
4
5 New-AzureRmStorageAccount -Name $saName -ResourceGroupName $rgName
   -Type $saType -Location $locName

```

3. Create an availability set

A load balancer configured with an availability set ensures that your application is always available.

```
$avName="<availability set name>"
```

```
New-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
$rgName -Location $locName
```

4. Create a virtual network

Add a new virtual network with at least one subnet, if the subnet was not created previously.

```

1 $vnetName = "LBVnet"
2
3 $vnetAddressPrefix="10.0.0.0/16"
4
5 $FrontendAddressPrefix="10.0.1.0/24"
6
7 $BackendAddressPrefix="10.0.2.0/24"
8
9 $vnet=New-AzureRmVirtualNetwork -Name $vnetName -ResourceGroupName
   $rgName -Location $locName -AddressPrefix $vnetAddressPrefix -
   Subnet $frontendSubnet,$backendSubnet`
10
11 $frontendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
   frontendSubnet -AddressPrefix $FrontendAddressPrefix
12
13 $backendSubnet=New-AzureRmVirtualNetworkSubnetConfig -Name
   backendSubnet -AddressPrefix $BackendAddressPrefix

```

Note:

Choose the AddressPrefix parameter value as per your requirement.

Assign front end and back end subnet to the virtual network that you created earlier in this step.

If the front end subnet is the first element of array VNet, subnetId must be \$vnet.Subnets[0].Id.

If the front end subnet is the second element in the array, the subnetId must be \$vnet.Subnets[1].Id, and so on.

5. Create a backend address pool

```
$beaddresspool= New-AzureRmLoadBalancerBackendAddressPoolConfig -
Name "LB-backend"
```

6. Create NAT rules

Create NAT rules for services that you are not load balancing.

```
1 $inboundNATRule1= New-AzureRmLoadBalancerInboundNatRuleConfig -
  Name "Inboundnatrule1" -FrontendIpConfiguration $frontendIP -
  Protocol TCP -FrontendPort 3441 -BackendPort 3389
2
3 $inboundNATRule2= New-AzureRmLoadBalancerInboundNatRuleConfig -
  Name "RDP2" -FrontendIpConfiguration $frontendIP -Protocol TCP -
  -FrontendPort 3442 -BackendPort 3389
```

Use front end and back end ports as per your requirement.

7. Create a health probe

Create a TCP health probe with port 9000 and interval 5 seconds.

```
1 $healthProbe = New-AzureRmLoadBalancerProbeConfig -Name "
  HealthProbe" " -Protocol tcp -Port 9000 -IntervalInSeconds 5 -
  ProbeCount 2
```

8. Create a load balancing rule

Create an LB rule for each service that you are load balancing.

For example:

You can use the following example to load balance HTTP service.

```
1 $lbrule = New-AzureRmLoadBalancerRuleConfig -Name "lbrule1" -
  FrontendIpConfiguration $frontendIP -BackendAddressPool
  $beAddressPool -Probe $healthProbe -Protocol Tcp -FrontendPort
  80 -BackendPort 80
```

Use front end and back end ports as per your requirement.

9. Create a load balancer entity

Create the load balancer adding all objects (NAT rules, load balancer rules, probe configurations) together.

```
1 $NRPLB = New-AzureRmLoadBalancer -ResourceGroupName $rgname -Name
  "InternalLB" -Location $locName -FrontendIpConfiguration
  $frontendIP -InboundNatRule $inboundNATRule1,$inboundNatRule2 -
  LoadBalancingRule $lbrule -BackendAddressPool $beAddressPool -
  Probe $healthProbe
```

10. Create a NIC

Create two NICs and associate each NIC with each Citrix ADC VPX instance

```

1 $backendnic1= New-AzureRmNetworkInterface -ResourceGroupName
  $rgName -Name lb-nic1-be -Location $locName -PrivateIpAddress
  10.0.2.6 -Subnet $backendSubnet -LoadBalancerBackendAddressPool
  $nrplb.BackendAddressPools\[0\] -LoadBalancerInboundNatRule
  $nrplb.InboundNatRules\[0\]

```

This NIC is for Citrix ADC VPX 1. The Private IP must be in same subnet as that of subnet added.

```

1 $backendnic2= New-AzureRmNetworkInterface -ResourceGroupName
  $rgName -Name lb-nic2-be -Location $locName -PrivateIpAddress
  10.0.2.7 -Subnet $backendSubnet -LoadBalancerBackendAddressPool
  $nrplb.BackendAddressPools\[0\] -LoadBalancerInboundNatRule
  $nrplb.InboundNatRules\[1\].

```

This NIC is for Citrix ADC VPX 2. The parameter `Private IP Address` can have any private IP as per your requirement.

11. Create Citrix ADC VPX instances

Create two VPX instances part of the same resource group and availability set, and attach it to the internal load balancer.

a) Citrix ADC VPX instance 1

For example:

```

1 $vmName="VPX1"
2
3 $vmSize="Standard\_A3"
4
5 $avSet=Get-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
  $rgName
6
7 $vm1=New-AzureRmVMConfig -VMName $vmName -VMSize $vmSize -
  AvailabilitySetId $avset.Id
8
9 $cred=Get-Credential -Message "Type Credentials which will be used
  to login to VPX instance"
10
11 $vm1=Set-AzureRmVMOperatingSystem -VM $vm1 -Linux -ComputerName
  $vmName -Credential $cred -Verbose
12
13 $vm1=Set-AzureRmVMSourceImage -VM $vm1 -PublisherName $pubName -
  Offer $offerName -Skus $skuName -Version "latest"
14
15 $vm1=Add-AzureRmVMNetworkInterface -VM $vm1 -Id $backendnic1.Id
16
17 $diskName="dynamic"
18
19 $storageAcc=Get-AzureRmStorageAccount -ResourceGroupName $rgName -
  Name $saName
20

```

```

21 $osDiskUri1=$storageAcc.PrimaryEndpoints.Blob.ToString() + "vhds1/"
    " + $diskName + ".vhd"
22
23 $vm1=Set-AzureRmVMOSDisk -VM $vm1 -Name $diskName -VhdUri
    $osDiskUri1 -CreateOption fromImage
24
25 Set-AzureRmVMPlan -VM $vm1 -Publisher $pubName -Product $offerName
    -Name $skuName
26
27 New-AzureRmVM -ResourceGroupName $rgName -Location $locName -VM
    $vm1

```

b) Citrix ADC VPX instance 2

For example:

```

1 $vmName="VPX2"
2
3 $vmSize="Standard\_A3"
4
5 $avSet=Get-AzureRmAvailabilitySet -Name $avName -ResourceGroupName
    $rgName
6
7 $vm2=New-AzureRmVMConfig -VMName $vmName -VMSize $vmSize -
    AvailabilitySetId $avset.Id
8
9 $cred=Get-Credential -Message " Type Credentials which will be
    used to login to VPX instance "
10
11 $vm2=Set-AzureRmVMOperatingSystem -VM $vm2 -Linux -ComputerName
    $vmName -Credential $cred -Verbose
12
13 $vm2=Set-AzureRmVMSourceImage -VM $vm2 -PublisherName $pubName -
    Offer $offerName -Skus $skuName -Version "latest"
14
15 $vm2=Add-AzureRmVMNetworkInterface -VM $vm2 -Id $backendnic2.Id
16
17 $diskName="dynamic"
18
19 $storageAcc=Get-AzureRmStorageAccount -ResourceGroupName $rgName -
    Name $saName
20
21 $osDiskUri1=$storageAcc.PrimaryEndpoints.Blob.ToString() + "vhds2/"
    " + $diskName + ".vhd"
22
23 $vm2=Set-AzureRmVMOSDisk -VM $vm2 -Name $diskName -VhdUri
    $osDiskUri1 -CreateOption fromImage
24
25 Set-AzureRmVMPlan -VM $vm2 -Publisher $pubName -Product $offerName
    -Name $skuName
26
27 New-AzureRmVM -ResourceGroupName $rgName -Location $locName -VM
    $vm2

```


12. Configure the virtual machines

When both the Citrix ADC VPX instances start, then connect to both Citrix ADC VPX instances using the SSH protocol to configure the virtual machines.

a) Active-Active: Run the same set of configuration commands on the command line of both the Citrix ADC VPX instances.

b) Active-Passive: Run this command on the command line of both the Citrix ADC VPX instances.

```
add ha node #nodeID <nsip of other Citrix ADC VPX>
```

In Active-Passive mode, run configuration commands on the primary node only.

Azure FAQs

June 20, 2024

- **Is the upgrade procedure of Citrix ADC VPX instance installed from Azure Marketplace different from the on-premises upgrade procedure?**

No. You can upgrade your Citrix ADC VPX instance in the Microsoft Azure cloud to Citrix ADC VPX release 11.1 or later, using standard Citrix ADC VPX upgrade procedures. You can upgrade either using GUI or CLI procedures. For any new installations, use the Citrix ADC VPX image for Microsoft Azure cloud.

To download the Citrix ADC VPX upgrade builds, go to **Citrix Downloads** > [Citrix ADC Firmware](#).

- **How to correct MAC moves and interface mutes observed on Citrix ADC VPX instances hosted on Azure?**

In Azure Multi-NIC environment, by default, all data interfaces might show MAC moves and interface mutes. To avoid MAC moves and interface mutes on Azure environments, Citrix recommends you to create a VLAN per data interface (without tag) of the ADC VPX instance and bind the primary IP of the NIC in Azure.

For more information, see [CTX224626](#) article.

Deploy a Citrix ADC VPX instance on the Google Cloud Platform

January 27, 2025

You can deploy a Citrix ADC VPX instance on the Google Cloud Platform (GCP). A VPX instance in GCP enables you to take advantage of GCP cloud computing capabilities and use Citrix load balancing and traffic management features for your business needs. You can deploy VPX instances in GCP as standalone instances. Both single NIC and multi NIC configurations are supported.

Supported features

All Premium, Advanced, and Standard features are supported on the GCP based on the license/version type used.

Limitation

- IPv6 isn't supported.

Hardware requirements

VPX instance in GCP must have minimum of 2 vCPUs and 4 GB RAM.


Points to note

Consider the following GCP-specific points before you begin your deployment.

- After creating the instance, you cannot add or remove any network interfaces.
- For a multi-NIC deployment, create separate VPC networks for each NIC. One NIC can be associated with only one network.
- For a single-NIC instance, the GCP console creates a network by default.
- Minimum 4 vCPUs are required for an instance with more than two network interfaces.
- If IP forwarding is required, you must enable IP forwarding while creating the instance and configuring the NIC.

Scenario: Deploy a multi-NIC, multi-IP standalone VPX instance

This scenario illustrates how to deploy a Citrix VPX standalone instance in GCP. In this scenario, you create a standalone VPX instance with multiple NICs. The instance communicates with back-end servers (the server farm).



Create three NICs to serve the following purposes.

NIC	Purpose	Associated with VPC network
NIC 0	Serves management traffic (Citrix ADC IP)	Management network
NIC 1	Serves client-side traffic (VIP)	Client network
NIC 2	Communicates with back-end servers (SNIP)	Back-end server network

Also, set up the required communication routes between the instance and the back-end servers, and between the instance and the external hosts on the public internet.

Summary of deployment steps

1. Create three VPC networks for three different NICs.
2. Create firewall rules for ports 22, 80, and 443
3. Create an instance with three NICs

Note:

Create an instance in the same region where you've created the VPC networks.

Step 1. Create VPC networks.

Create three VPC networks that are associated with management NIC, client NIC, and server NIC. To create a VPC network, log on to **Google console > Networking > VPC network > Create VPC Network**. Complete the required fields, as shown in the screen capture, and click **Create**.

netscaler-vpx-platform-eng

← Create a VPC network

Name ?
vpxmgmt

Description (Optional)
management vpc

Subnets

Subnets let you create your own private cloud topology within Google Cloud. Click Automatic to create a subnet in each region, or click Custom to manually define the subnets. [Learn more](#)

Subnet creation mode
 Custom Automatic

New subnet 🗑️ ⬆️

Name ?
vpxmgmtsubnet

[Add a description](#)

Region ?
asia-east1

IP address range ?
192.168.30.0/24

[Create secondary IP range](#)

Private Google access ?
 On
 Off

Flow logs
 On
 Off

Dynamic routing mode ?
 Regional
Cloud Routers will learn routes only in the region in which they were created
 Global
Global routing lets you dynamically learn routes to and from all regions with a single VPN or interconnect and Cloud Router

Similarly, create VPC networks for client and server-side NICs.

Note:

All three VPC networks must be in the same region, which is asia-east1 in this scenario.

Step 2. Create firewall rules for ports 22, 80, and 443.

Create rules for SSH (port 22), HTTP (port 80), and HTTPS (port 443) for each VPC networks. For more information about firewall rules, see [Firewall Rules Overview](#).

netscaler-vpx-platform-eng

←

Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. [Learn more](#)

Name ?

Description (Optional)

Logs
 Turning on firewall logs can generate a large number of logs which can increase costs in Stackdriver. [Learn more](#)

On
 Off

Network ?

Priority ?
 Priority can be 0 - 65535 [Check priority of other firewall rules](#)

Direction of traffic ?

Ingress
 Egress

Action on match ?

Allow
 Deny

Targets ?

Source filter ?

Source IP ranges ?

Second source filter ?

Protocols and ports ?

Allow all
 Specified protocols and ports

tcp :

udp :

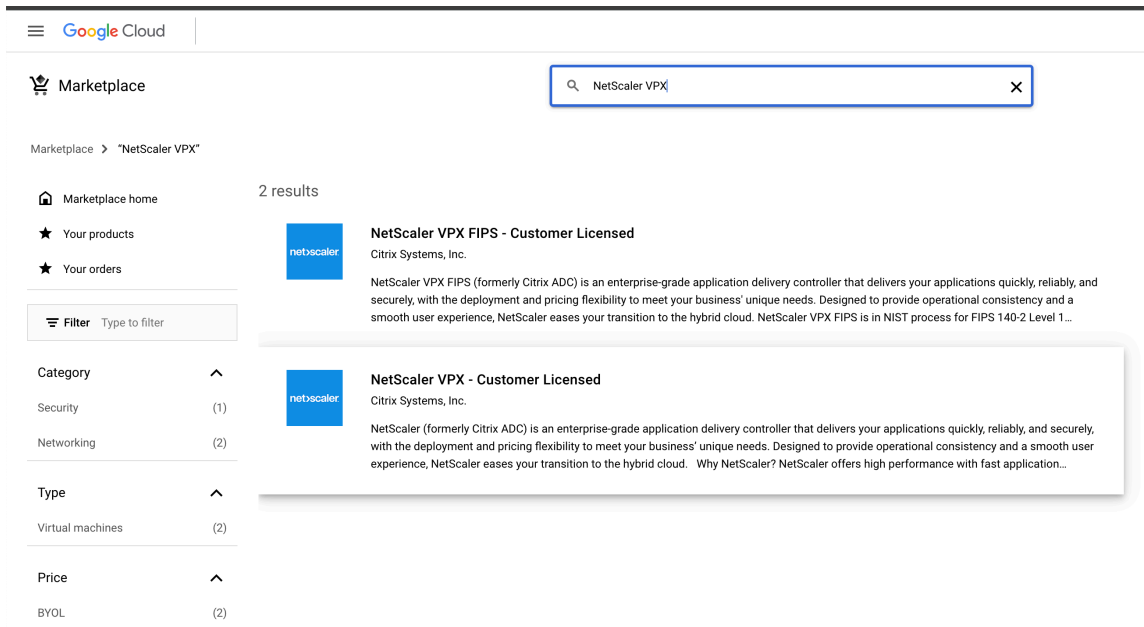
Other protocols

[↕ Disable rule](#)

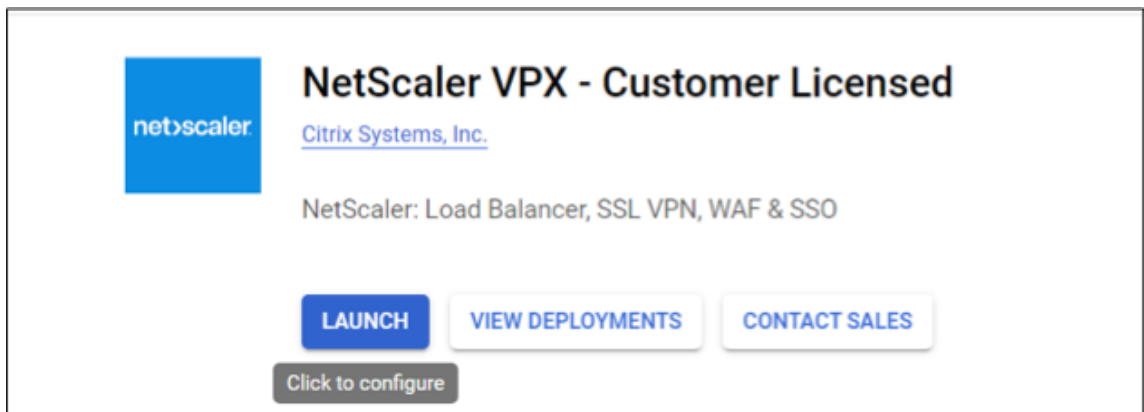
Create
Cancel

Step 3. Create the VPX instance.

1. Log on to the GCP console.
2. Navigate to the [GCP Marketplace](#).
3. Select a subscription based on your requirements.



4. Click **Launch** on the selected subscription.



5. Complete the deployment form and click **Deploy**.

Note:

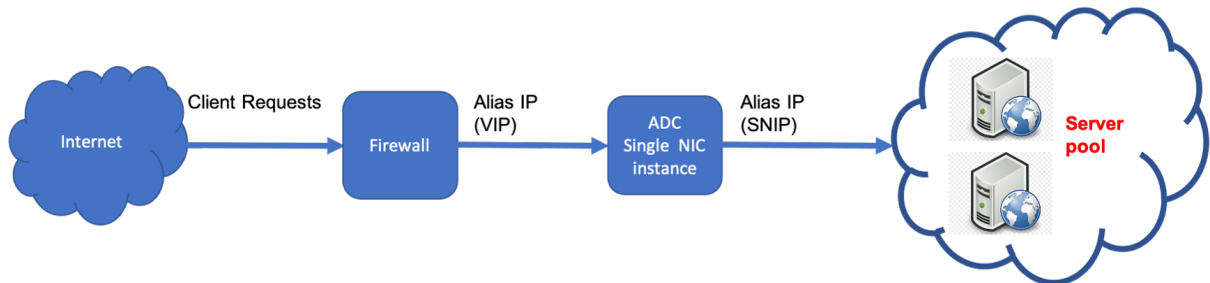
Use the VPC Networks created in **Step 1**.

6. The deployed instance appears under **Compute Engine > VM instances**.

Use the GCP SSH or the serial console to configure and manage the VPX instance.

Scenario: Deploy a single-NIC, standalone VPX instance

This scenario illustrates how to deploy a Citrix VPX standalone instance with a single NIC in GCP. The alias IP addresses are used to achieve this deployment.



Create a single NIC (NIC0) to serve the following purposes:

- Handle management traffic (Citrix ADC IP) in the management network.
- Handle client-side traffic (VIP) in the client network.
- Communicate with back-end servers (SNIP) in the back-end server network.

Set up the required communication routes between the following:

- Instance and the back-end servers.
- Instance and the external hosts on the public internet.

Summary of deployment steps

1. Create a VPC network for NIC0.
2. Create firewall rules for ports 22, 80, and 443.
3. Create an instance with a single NIC.
4. Add Alias IP addresses to VPX.
5. Add VIP and SNIP on VPX.
6. Add a load balancing virtual server.
7. Add a service or service group on the instance.
8. Bind the service or service group to the load balancing virtual server on the instance.

Note:

Create an instance in the same region where you've created the VPC networks.

Step 1. Create one VPC network.

Create one VPC network to associate with NIC0.

To create a VPC network, do these steps:

1. Log on to **GCP console > Networking > VPC network > Create VPC Network**
2. Complete the required fields, and click **Create**.

The screenshot displays two overlapping configuration windows in the Google Cloud Platform console. The top window, titled 'Create a VPC network', shows the 'Name' field set to 'vpxmgmt' and the 'Description' field set to 'management vpc'. Below this, the 'Subnets' section is visible, with 'Subnet creation mode' set to 'Custom'. The bottom window, titled 'New subnet', shows the 'Name' field set to 'vpxmgmtsubnet', the 'Region' set to 'asia-east1', and the 'IP address range' set to '192.168.30.0/24'. The 'Private Google access' option is set to 'On', and 'Flow logs' are set to 'Off'. At the bottom of the 'New subnet' window, the 'Dynamic routing mode' is set to 'Regional'. Both windows have 'Done' and 'Cancel' buttons.

Step 2. Create firewall rules for ports 22, 80, and 443.

Create rules for SSH (port 22), HTTP (port 80), and HTTPS (port 443) for the VPC network. For more information about firewall rules, see [Firewall Rules Overview](#).

netscaler-vpx-platform-eng

← Create a firewall rule

Firewall rules control incoming or outgoing traffic to an instance. By default, incoming traffic from outside your network is blocked. [Learn more](#)

Name

Description (Optional)

Logs
Turning on Firewall logs can generate a large number of logs which can increase costs in Stackdriver. [Learn more](#)
 On
 Off

Network

Priority
Priority can be 0 - 65535 Check priority of other firewall rules

Direction of traffic
 Ingress
 Egress

Action on match
 Allow
 Deny

Targets

Source filter

Source IP ranges

Second source filter

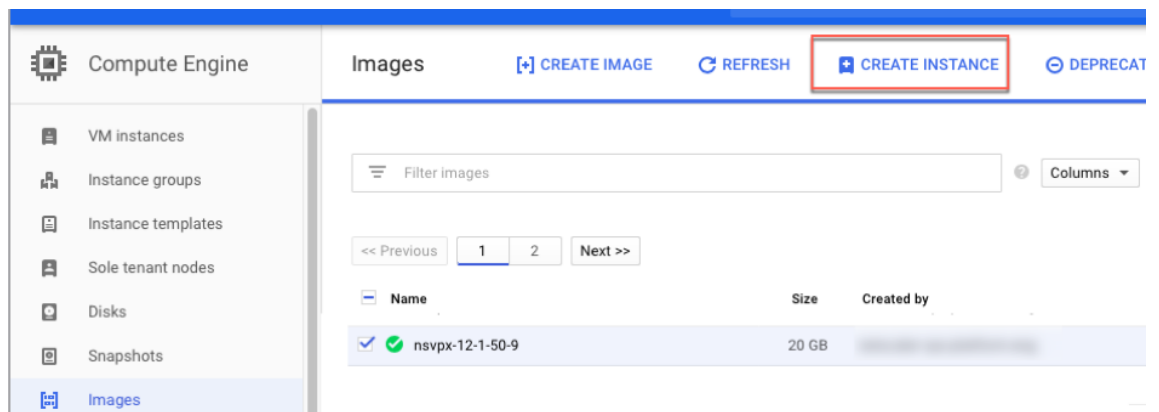
Protocols and ports
 Allow all
 Specified protocols and ports
 tcp:
 udp:
 Other protocols

⌵ Disable rule


Step 3. Create an instance with single NIC.

To create an instance with single NIC, do these steps:

1. Log on to the **GCP console**.
2. Under **Compute**, hover over **Compute Engine**, and select **Images**.
3. Select the image, and click **Create Instance**.



4. Select an instance type with two vCPUs (minimum requirement for ADC).



5. Click the **Networking** tab from the **Management, security, disks, networking** window.
6. Under **Network interfaces**, click the **Edit** icon to edit the default NIC.
7. In the **Network interfaces** window, under **Network**, select the VPC network that you created.
8. You can create a static external IP address. Under the **External IP addresses**, click **Create IP address**.
9. In the **Reserve a static address** window, add a name and description and click **Reserve**.
10. Click **Create** to create the VPX instance.
The new instance appears under VM instances.

Step 4. Add alias IP addresses to the VPX instance.


Assign two alias IP addresses to the VPX instance to use as VIP and SNIP addresses.

Note:
Do not use the primary internal IP address of the VPX instance to configure the VIP or SNIP.

To create an alias IP address, perform these steps:

1. Navigate to the VM instance and click **Edit**.
2. In the **Network interface** window, edit the NIC0 interface.

3. In the **Alias IP range** field, enter the alias IP addresses.



4. Click **Done**, and then **Save**.

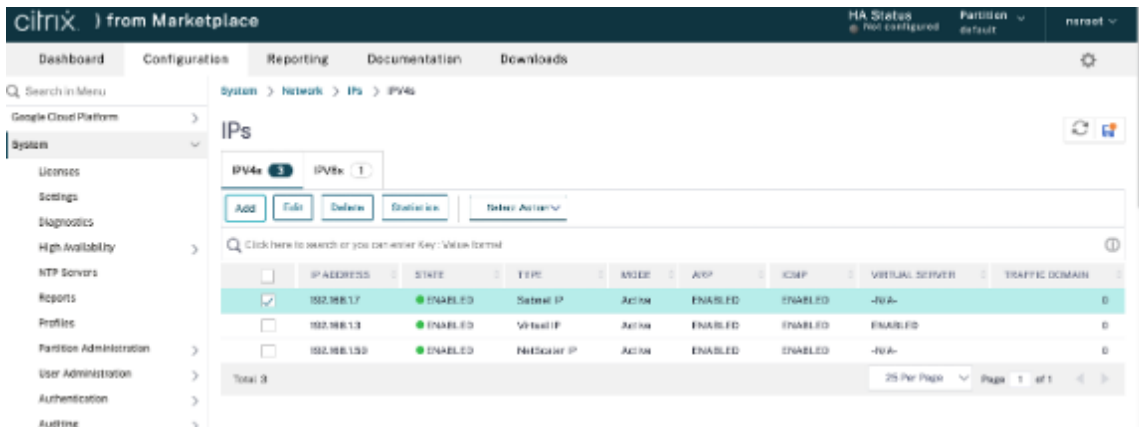
5. Verify the alias IP addresses in the **VM instance details** page.



Step 5. Add VIP and SNIP on the VPX instance.

On the VPX instance, add client alias IP address and server alias IP address.

1. On the Citrix ADC GUI, navigate to **System > Network > IPs > IPv4s**, and click **Add**.



2. To create a client alias IP (VIP) address:

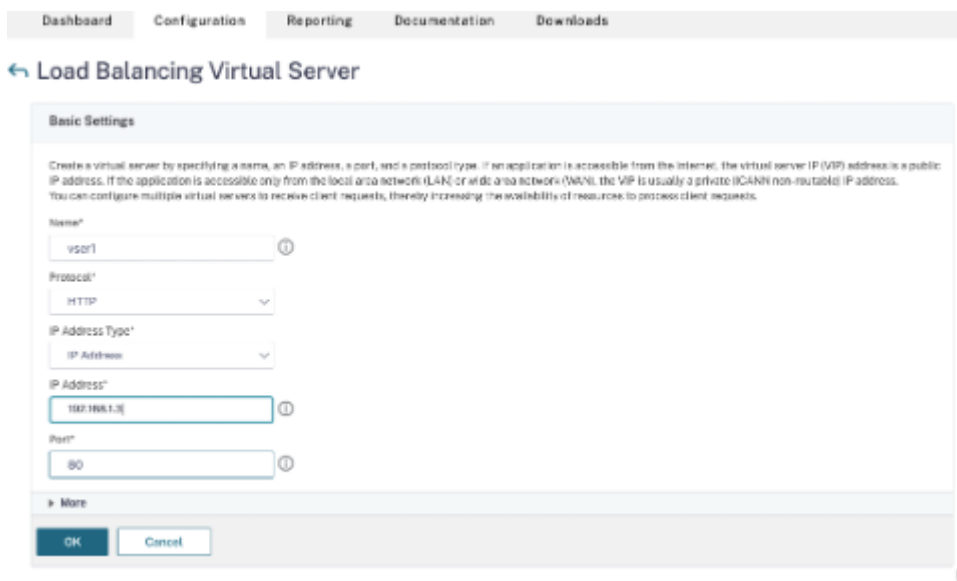
- Enter the client-alias IP address and netmask configured for the VPC subnet in the VM instance.
- In the **IP Type** field, select **Virtual IP** from the drop-down menu.
- Click **Create**.

3. To create a server alias IP (SNIP) address:

- Enter the server-alias IP address and netmask configured for the VPC subnet in the VM instance.
- In the **IP Type** field, select **Subnet IP** from the drop-down menu.
- Click **Create**.

Step 6. Add load balancing virtual server.

1. On the Citrix ADC GUI, navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers**, and click **Add**.
2. Add the required values for Name, Protocol, IP Address Type (IP Address), IP Address (client alias IP), and Port.
3. Click **OK** to create the load balancing virtual server.



Step 7. Add a service or service group on the VPX instance.

1. From the Citrix ADC GUI, navigate to **Configuration > Traffic Management > Load Balancing > Services**, and click **Add**.
2. Add the required values for Service Name, IP Address, Protocol, and Port, and click **OK**.

Step 8. Bind the service/service group to the Load Balancing Virtual Server on the instance.

1. From the GUI, navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers**.
2. Select the load balancing virtual server configured in **Step 6**, and click **Edit**.
3. In the **Service and Service Groups** window, click **No Load Balancing Virtual Server Service Binding**.
4. Select the service configured in **Step 7**, and click **Bind**.

Points to note after you’ve deployed the VPX instance on GCP

- Log on to the VPX with user name `nsroot` and instance ID as password. At the prompt, change the password and save the configuration.
- For collecting a technical support bundle, run the command `shell /netscaler/showtech_cloud.pl` instead of the customary `show techsupport`.
- After deleting a Citrix ADC VM from GCP console, delete the associated Citrix ADC internal target instance also. To do so, go to gcloud CLI and type the following command:

```
1 gcloud compute -q target-instances delete <instance-name>-adcinternal --zone <zone>
```

Note:

`<instance-name>-adcinternal` is the name of the target instance that must be deleted.

Citrix ADC VPX licensing

A Citrix ADC VPX instance on GCP requires a valid license. The licensing option available for Citrix ADC VPX instances running on GCP is as follows:

Bring your own license (BYOL): To use the BYOL option, follow these steps:

- Use the licensing portal on the Citrix website to generate a valid license.
- Upload the generated license to the instance.
- **Citrix ADC VPX Check-in and Check-out license:** This licensing model allows you to check out a license from a pool of available licenses and check it back in when no longer needed. For more information and detailed instructions, see the [Citrix ADC VPX Check-in and Check-out License](#).

Note:

Subscription-based licensing is no longer supported for Citrix ADC VPX instances on GCP.

Supported Citrix ADC VPX offerings on GCP

The following table lists the supported Citrix ADC VPX offerings on GCP.

Supported VPX offerings

Citrix ADC VPX - Customer Licensed

Citrix ADC VPX FIPS - Customer Licensed

Supported GCP machine type families

Machine type family	Minimum machine type
General Purpose Machines	e2-medium, e2-standard-2, e2-highmem-2, n1-standard-2, n1-highmem-2, n2-standard-2, n2-highmem-2, n2d-standard-2, n2d-highmem-2
Compute-Optimized Machines	c2-standard-4, c2d-standard-2, c2d-highmem-2

GDM templates to deploy a Citrix ADC VPX instance

You can use a Citrix ADC VPX Google Deployment Manager (GDM) template to deploy a VPX instance on GCP. For details, see [Citrix ADC GDM Templates](#).

Citrix ADC marketplace images

You can use the images in GDM templates to bring up the Citrix ADC appliance.

The following table lists the images that are available on the GCP marketplace.

Release	Image name	Image location
13.0	citrix-adc-vpx-10-enterprise-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-10-enterprise-13-0-83-29
13.0	citrix-adc-vpx-10-platinum-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-10-platinum-13-0-83-29
13.0	citrix-adc-vpx-10-standard-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-10-standard-13-0-83-29
13.0	citrix-adc-vpx-200-enterprise-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-200-enterprise-13-0-83-29
13.0	citrix-adc-vpx-200-platinum-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-200-platinum-13-0-83-29
13.0	citrix-adc-vpx-200-standard-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-200-standard-13-0-83-29
13.0	citrix-adc-vpx-1000-advanced-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-1000-advanced-13-0-83-29

Release	Image name	Image location
13.0	citrix-adc-vpx-1000-premium-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-1000-premium-13-0-83-29
13.0	citrix-adc-vpx-1000-standard-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-1000-standard-13-0-83-29
13.0	citrix-adc-vpx-3000-enterprise-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-3000-enterprise-13-0-83-29
13.0	citrix-adc-vpx-3000-platinum-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-3000-platinum-13-0-83-29
13.0	citrix-adc-vpx-3000-standard-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-3000-standard-13-0-83-29
13.0	citrix-adc-vpx-5000-enterprise-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-5000-enterprise-13-0-83-29
13.0	citrix-adc-vpx-5000-platinum-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-5000-platinum-13-0-83-29
13.0	citrix-adc-vpx-5000-standard-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-5000-standard-13-0-83-29
13.0	citrix-adc-vpx-byol-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-byol-13-0-83-29
13.0	citrix-adc-vpx-express-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-express-13-0-83-29

Release	Image name	Image location
13.0	citrix-adc-vpx-waf-1000-13-0-83-29	projects/citrix-master-project/global/images/citrix-adc-vpx-waf-1000-13-0-83-29

Resources

- [Creating Instances with Multiple Network Interfaces](#)
- [Creating and Starting a VM Instance](#)

Related information

- [Deploy a VPX high-availability pair on Google Cloud Platform](#)

Deploy a VPX high-availability pair on Google Cloud Platform

September 9, 2024

You can configure two Citrix ADC VPX instances on Google Cloud Platform (GCP) as a high availability (HA) active-passive pair. When you configure one instance as the primary node and the other as the secondary node, the primary node accepts connections and manages servers. The secondary node monitors the primary. If for any reason, if the primary node is unable to accept connections, the secondary node takes over.

For more information on HA, see [High Availability](#).

The nodes must be in the same region; however, they can be either in the same zone or different zones. For more information, see [Regions and Zones](#).

Each VPX instance requires at least three IP subnets (Google VPC networks):

- A management subnet
- A client-facing subnet (VIP)
- A back-end facing subnet (SNIP, MIP, and so on)

Citrix recommends three network interfaces for a standard VPX instance.

You can deploy a VPX high-availability pair in the following methods:

- [Using external static IP address](#)

- [Using private IP address](#)

GDM templates to deploy a VPX high-availability pair on GCP

You can use a Citrix ADC Google Deployment Manager (GDM) template to deploy a VPX high-availability pair on GCP. For details, see [Citrix ADC GDM Templates](#).

Forwarding rules support for VPX high-availability pair on GCP

You can deploy a VPX high-availability pair on the GCP using forwarding rules.

For more information on forwarding rules, see [Forwarding rules overview](#).

Prerequisites

- Forwarding rules must be in the same region as the VPX instances.
- Target instances must be in the same zone as the VPX instance.
- Number of target instances for both primary and secondary nodes must match.

Example:

You have a high-availability pair in the `us-east1` region with primary VPX in `us-east1-b` zone and secondary VPX in `us-east1-c` zone. A forwarding rule is configured for the primary VPX with the target instance in `us-east1-b` zone. Configure a target instance for secondary VPX in `us-east1-c` zone to update the forwarding rule on failover.

Limitations

Only forwarding rules that are configured with target instances at the back end are supported in VPX high-availability deployment.

Deploy a VPX high-availability pair with external static IP address on the Google Cloud Platform

September 13, 2024

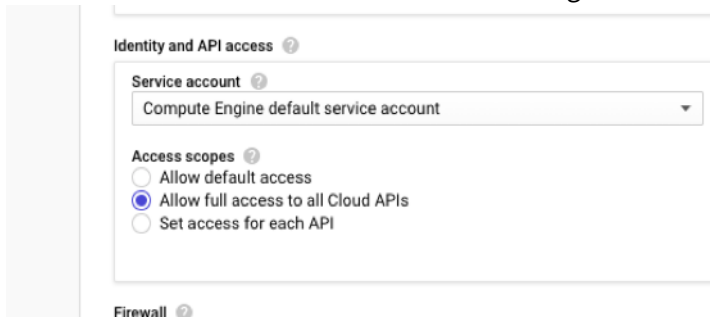
You can deploy a VPX high-availability pair on GCP using an external static IP address. The client IP address of the primary node must be bound to an external static IP address. Upon failover, the external static IP address is moved to the secondary node for traffic to resume.

A static external IP address is an external IP address that is reserved for your project until you decide to release it. If you use an IP address to access a service, you can reserve that IP address so that only your project can use it. For more information, see [Reserving a Static External IP Address](#).

For more information on HA, see [High Availability](#).

Before you start

- Read the Limitation, Hardware requirements, Points to note mentioned in [Deploy a Citrix ADC VPX instance on Google Cloud Platform](#). This information applies to HA deployments also.
- Enable **Cloud Resource Manager API** for your GCP project.
- Allow full access to all Cloud APIs while creating the instances.



- Ensure that the IAM role associated with your GCP service account has the following IAM permissions:

```

1  REQUIRED_INSTANCE_IAM_PERMS = [
2
3  "compute.addresses.use",
4  "compute.forwardingRules.list",
5  "compute.forwardingRules.setTarget",
6  "compute.instances.setMetadata",
7  "compute.instances.addAccessConfig",
8  "compute.instances.deleteAccessConfig",
9  "compute.instances.get",
10 "Compute.instances.list",
11 "compute.networks.useExternalIp",
12 "compute.subnetworks.useExternalIp",
13 "compute.targetInstances.list",
14 "compute.targetInstances.use",
15 "compute.zones.list",
16 ]

```

- If you have configured alias IP addresses on an interface other than the management interface, ensure that your GCP service account has the following extra IAM permissions:

```

1  "compute.instances.updateNetworkInterface"

```

- If you have configured GCP forwarding rules on the primary node, read the limitations and requirements mentioned in [Forwarding rules support for VPX high-availability pair on GCP](#) to update them to new primary on failover.

How to deploy a VPX HA pair on Google Cloud Platform

Here's a summary of the HA deployment steps:

1. Create VPC networks in the same region. For example, Asia-east.
2. Create two VPX instances (primary and secondary nodes) on the same region. They can be in the same zone or different zones. For example Asia east-1a and Asia east-1b.
3. Configure HA settings on both instances by using the Citrix ADC GUI or ADC CLI commands.

Step 1. Create VPC networks

Create VPC networks based on your requirements. Citrix recommends you to create three VPC networks for associating with management NIC, client NIC, and server NIC.

To create a VPC network, perform these steps:

1. Log on the **Google console > Networking > VPC network > Create VPC Network**.
2. Complete the required fields, and click **Create**.


For more information, see the **Create VPC Networks** section in [Deploy a Citrix ADC VPX instance on Google Cloud Platform](#).

Step 2. Create two VPX instances

Create two VPX instances by following the steps given in [Scenario: deploy a multi-NIC, multi-IP stand-alone VPX instance](#).

Important:

Assign a static external IP address to client IP address (VIP) of the primary node. You can use an existing reserved IP address or create a new one. To create a static external IP address, navigate to **Network interface > External IP**, click **Create IP address**.



After the failover, when the old primary becomes the new secondary, the static external IP address moves from the old primary and is attached to the new primary. For more information, see the Google cloud document [Reserving a Static External IP Address](#).

After you've configured the VPX instances, you can configure the VIP and SNIP addresses. For more information, see [Configuring Citrix ADC-owned IP addresses](#).

Step 3. Configure high availability

After you've created the instances on Google Cloud Platform, you can configure HA by using the Citrix ADC GUI for CLI.

Configure HA by using the GUI Step 1. Set up high availability in INC mode on both the instances.

On the **primary node**, perform the following steps:

1. Log on to the instance with user name `nsroot` and instance ID of the node from GCP console as the password.
2. Navigate to **Configuration > System > High Availability > Nodes**, and click **Add**.
3. In the **Remote Node IP address** field, enter the private IP address of the management NIC of the secondary node.
4. Select the **Turn on INC (Independent Network Configuration) mode on self node** check box.
5. Click **Create**.

On the **secondary node**, perform the following steps:

1. Log on to the instance with user name `nsroot` and instance ID of the node from GCP console as the password.
2. Navigate to **Configuration > System > High Availability > Nodes**, and click **Add**.
3. In the **Remote Node IP address** field, enter the private IP address of the management NIC of the primary node.
4. Select the **Turn on INC (Independent Network Configuration) mode on self node** check box.
5. Click **Create**.

Before you proceed further, ensure that the Synchronization state of the secondary node is shown as **SUCCESS** in the **Nodes** page.

System / High Availability / Nodes

Nodes 2

Add Edit Delete Statistics Select Action

<input type="checkbox"/>	ID	IP ADDRESS	HOST NAME	MASTER STATE	NODE STATE	INC	SYNCHRONIZATION STATE	SYNCHRONIZATION FAILURE REASON
<input type="checkbox"/>	0	192.168.1.3		Primary	● UP	ENABLED	ENABLED	-NA-
<input type="checkbox"/>	1	192.168.1.66		Secondary	● UP	ENABLED	SUCCESS	-NA-

Total 2 25 Per Page Page 1 of 1

Note:

Now, the secondary node has the same log-on credentials as the primary node.

Step 2. Add Virtual IP address and Subnet IP address on both the nodes.

On the **primary node**, perform the following steps:

1. Navigate to **System > Network > IPs > IPv4s**, and click **Add**.
2. Add a primary VIP address by following these steps:
 - a) Enter the internal IP address of the client-facing interface of the primary instance and net-mask configured for the client subnet in the VM instance.

- b) In the **IP Type** field, select **Virtual IP** from the drop-down menu.
 - c) Click **Create**.
3. Add a primary SNIP address by following these steps:
- a) Enter the internal IP address of the server-facing interface of the primary instance and netmask configured for the server subnet in the primary instance.
 - b) In the **IP Type** field, select **Subnet IP** from the drop-down menu.
 - c) Click **Create**.
4. Add a secondary VIP address by following these steps:
- a) Enter the internal IP address of the client-facing interface of the secondary instance and netmask configured for the client subnet in the VM instance.
 - b) In the **IP Type** field, select **Virtual IP** from the drop-down menu.
 - c) Click **Create**.

IPs

IPv4s 4 IPv6s 1

Add Edit Delete Statistics Select Action

Click here to search or you can enter Key : Value format

<input type="checkbox"/>	IP ADDRESS	STATE	TYPE	MODE	ARP	ICMP	VIRTUAL SERVER	TRAFFIC DOMAIN
<input type="checkbox"/>	192.168.2.54	ENABLED	Virtual IP	Active	ENABLED	ENABLED	ENABLED	0
<input type="checkbox"/>	192.168.3.7	ENABLED	Subnet IP	Active	ENABLED	ENABLED	-N/A-	0
<input type="checkbox"/>	192.168.2.37	ENABLED	Virtual IP	Active	ENABLED	ENABLED	ENABLED	0
<input type="checkbox"/>	192.168.1.3	ENABLED	NetScaler IP	Active	ENABLED	ENABLED	-N/A-	0

Total 4

25 Per Page Page 1 of 1

On the **secondary node**, perform the following steps:

1. Navigate to **System > Network > IPs > IPv4s**, and click **Add**.
2. Add a secondary VIP address by following these steps:
 - a) Enter the internal IP address of the client-facing interface of the secondary instance and netmask configured for the client subnet in the VM instance.
 - b) In the **IP Type** field, select **Virtual IP** from the drop-down menu.
3. Add a secondary SNIP address by following these steps:
 - a) Enter the internal IP address of the server-facing interface of the secondary instance and netmask configured for the server subnet in the secondary instance.
 - b) In the **IP Type** field, select **Subnet IP** from the drop-down menu.
 - c) Click **Create**.

	IP ADDRESS	STATE	TYPE	MODE	ARP	ICMP	VIRTUAL SERVER	TRAFFIC DOMAIN
Secondary SNIP	192.168.3.76	ENABLED	Subnet IP	Active	ENABLED	ENABLED	-N/A-	0
Secondary VIP	192.168.2.54	ENABLED	Virtual IP	Passive	ENABLED	ENABLED	ENABLED	0
	192.168.1.66	ENABLED	NetScaler IP	Active	ENABLED	ENABLED	-N/A-	0

Step 3. Add IP set and bind IP set to the secondary VIP on both the instances.

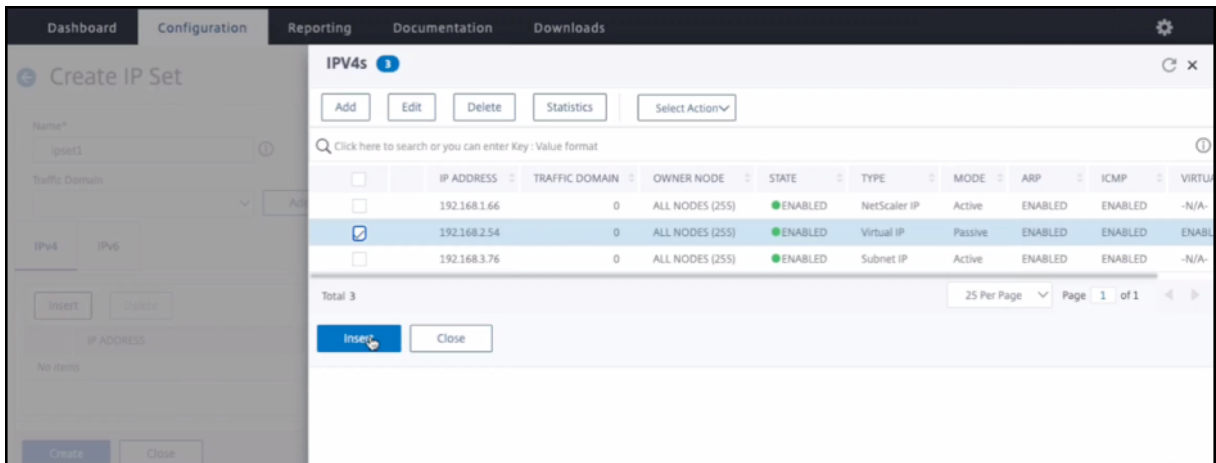
On the **primary node**, perform the following steps:

1. Navigate to **System > Network > IP Sets > Add**.
2. Add an IP set name and click **Insert**.
3. From the **IPv4s** page, select the virtual IP (secondary VIP) and click **Insert**.
4. Click **Create** to create the IP set.

	IP ADDRESS	TRAFFIC DOMAIN	OWNER NODE	STATE	TYPE	MODE	ARP	ICMP	VIRTUAL SERVER
	192.168.1.3	0	ALL NODES (255)	ENABLED	NetScaler IP	Active	ENABLED	ENABLED	-N/A-
	192.168.2.37	0	ALL NODES (255)	ENABLED	Virtual IP	Active	ENABLED	ENABLED	ENABLED
	192.168.3.7	0	ALL NODES (255)	ENABLED	Subnet IP	Active	ENABLED	ENABLED	-N/A-
	192.168.2.54	0	ALL NODES (255)	ENABLED	Virtual IP	Active	ENABLED	ENABLED	ENABLED

On the **secondary node**, perform the following steps:

1. Navigate to **System > Network > IP Sets > Add**.
2. Add an IP set name and click **Insert**.
3. From the **IPv4s** page, select the virtual IP (secondary VIP) and click **Insert**.
4. Click **Create** to create the IP set.




Note:

IP set name must be same on both the instances.

Step 4. Add a load balancing virtual server on the primary instance.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers > Add.**
2. Add the required values for Name, Protocol, IP Address Type (IP Address), IP address (primary VIP), and Port.



3. Click **More**. Navigate to **IP Range IP Set Settings**, select **IPset** from the drop-down menu, and provide the IPset created in **Step 3**.
4. Click **OK** to create the load balancing virtual server.

Step 5. Add a service or service group on the primary node.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Services > Add.**
2. Add the required values for Service Name, IP Address, Protocol and Port, and click **OK**.

Step 6. Bind the service or service group to the load balancing virtual server on the primary node.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers**.
2. Select the load balancing virtual server configured in **Step 4**, and click **Edit**.
3. In the **Service and Service Groups** tab, click **No Load Balancing Virtual Server Service Binding**.
4. Select the service configured in the **Step 5**, and click **Bind**.

Save the configuration. After a forced failover, the secondary becomes the new primary. The external static IP of the old primary VIP moves to the new secondary VIP.

Configure high availability using CLI Step 1. Set up high availability in INC mode in both the instances.

On the primary node, type the following command.

```
1 add ha node 1 <sec_ip> -inc ENABLED
```

On the secondary node, type the following command.

```
1 add ha node 1 <prim_ip> -inc ENABLED
```

`sec_ip` refers to the internal IP address of the management NIC of the secondary node.

`prim_ip` refers to the internal IP address of the management NIC of the primary node.

Step 2. Add Virtual and Subnet IPs on both the nodes.

On the primary node, type the following command.

```
1 add ns ip <primary_vip> <subnet> -type VIP
2
3 add ns ip <secondary_vip> <subnet> -type VIP
4
5 add ns ip <primary_snip> <subnet> -type SNIP
```

`primary_vip` refers to the internal IP address of the client-facing interface of the primary instance.

`secondary_vip` refers to the internal IP address of the client-facing interface of the secondary instance.

`primary_snip` refers to the internal IP address of the server-facing interface of the primary instance.

On the secondary node, type the following command.

```
1 add ns ip <secondary_vip> <subnet> -type VIP
2
3 add ns ip <secondary_snip> <subnet> -type SNIP
```

`secondary_vip` refers to the internal IP address of the client-facing interface of the secondary instance.

`secondary_snip` refers to the internal IP address of the server-facing interface of the secondary instance.

Step 3. Add IP set and bind IP set to secondary VIP on both the instances.

On the primary node, type the following command:

```
1 add ipset <ipsetname>
2
3 bind ipset <ipsetname> <secondary VIP>
```

On the secondary node, type the following command:

```
1 add ipset <ipsetname>
2
3 bind ipset <ipsetname> <secondary VIP>
```

Note:

IP set name must be same on both the instances.

Step 4. Add a virtual server on the primary instance.

Type the following command:

```
1 add <server_type> vserver <vserver_name> <protocol> <primary_vip> <port>
  > -ipset <ipset_name>
```

Step 5. Add a service or service group on the primary instance.

Type the following command:

```
1 add service <service_name> <service_ip_address> <protocol> <port>
```

Step 6. Bind the service/service group to the load balancing virtual server on the primary instance.

Type the following command:

```
1 bind <server_type> vserver <vserver_name> <service_name>
```

Note:

To save your configuration, type the command `save config`. Otherwise, the configurations are lost after you restart the instances.

Step 7. Verify the configuration.

Ensure that the external IP address attached to the primary client NIC moves to the secondary on a failover.

1. Make a cURL request to the external IP address and make sure that it is reachable.
2. On the primary instance, perform failover:

From GUI, navigate to **Configuration > System > High Availability > Action > Force Failover**.

From CLI, type the following command:

```
1 force ha failover -f
```

On the GCP console, goto the Secondary instance. The external IP address must have moved to the client NIC of secondary after failover.

3. Issue a cURL request to the external IP and ensure it is reachable again.

Deploy a VPX high-availability pair with private IP address on Google Cloud Platform

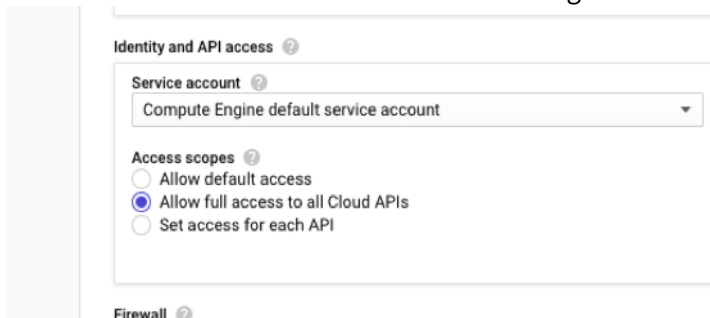
September 13, 2024

You can deploy a VPX high-availability pair on GCP using private IP address. The client IP (VIP) must be configured as alias IP address on the primary node. Upon failover, the Client IP address is moved to the secondary node, for the traffic to resume.

For more information on high availability, see [High Availability](#).

Before you start

- Read the Limitation, Hardware requirements, Points to note mentioned in [Deploy a Citrix ADC VPX instance on Google Cloud Platform](#). This information applies to high availability deployments also.
- Enable **Cloud Resource Manager API** for your GCP project.
- Allow full access to all Cloud APIs while creating the instances.



- Ensure that your GCP service account has the following IAM permissions:

```
1  REQUIRED_INSTANCE_IAM_PERMS = [  
2  "compute.forwardingRules.list" ,  
3  "compute.forwardingRules.setTarget" ,  
4  "compute.instances.setMetadata" ,  
5  "compute.instances.get",  
6  "compute.instances.list",  
7  "compute.instances.updateNetworkInterface",  
8  "compute.targetInstances.list" ,  
9  "compute.targetInstances.use" ,  
10 "compute.zones.list",  
11 ]
```

- If you have configured external IP addresses on an interface other than the management interface, ensure that your GCP service account has the following additional IAM permissions:

```
1  REQUIRED_INSTANCE_IAM_PERMS = [  
2  "compute.addresses.use"  
3  "compute.instances.addAccessConfig",  
4  "compute.instances.deleteAccessConfig",  
5  "compute.networks.useExternalIp",  
6  "compute.subnetworks.useExternalIp",  
7  ]
```

- If your VMs do not have internet access, you must enable **Private Google Access** on the management subnet.

Add a subnet

Name ⓘ
Name is permanent
management-subnet

Add a description

VPC Network
automationmgmtnetwork

Region ⓘ
us-east1

Reserve for Internal HTTP(S) Load Balancing ⓘ
 On
 Off

IP address range ⓘ
192.168.2.0/24

Create secondary IP range

Private Google access ⓘ
 On
 Off

Flow logs
Turning on VPC flow logs doesn't affect performance, but some systems generate a large number of logs, which can increase costs in Stackdriver. [Learn more](#)
 On
 Off

CANCEL **ADD**

- If you have configured GCP forwarding rules on the primary node, read the limitations and requirements mentioned in [Forwarding rules support for VPX high-availability pair on GCP](#) to update them to new primary on failover.

How to deploy a VPX high availability pair on Google Cloud Platform

Here is a summary of the high availability deployment steps:

1. Create VPC networks in the same region. For example, Asia-east.
2. Create two VPX instances (primary and secondary nodes) on the same region. They can be in the same zone or different zones. For example Asia east-1a and Asia east-1b.
3. Configure high availability settings on both instances by using the Citrix ADC GUI or ADC CLI commands.

Step 1. Create VPC networks

Create VPC networks based on your requirements. Citrix recommends you to create three VPC networks for associating with management NIC, client NIC, and server NIC.

To create a VPC network, perform these steps:

1. Log on the **Google console > Networking > VPC network > Create VPC Network**.
2. Complete the required fields, and click **Create**.

For more information, see the **Create VPC Networks** section in [Deploy a Citrix ADC VPX instance on Google Cloud Platform](#).

Step 2. Create two VPX instances

Create two VPX instances by following the steps given in [Scenario: deploy a multi-NIC, multi-IP stand-alone VPX instance](#).

Important:

Assign a client alias IP address to the primary node. Do not use the internal IP address of the VPX instance to configure the VIP.

To create a client alias IP address, perform these steps:

1. Navigate to the VM instance and click **Edit**.
2. In the **Network Interface** window, edit the client interface.
3. In the **Alias IP range** field, enter the client alias IP address.

VM instance details

Creation time
Jan 16, 2020, 4:00:22 PM

Network interfaces

nic0: automationmgmtnetwork mgmtsubnet

Network interface

Network
automationclientnetwork

Subnetwork
clientsubnet

Internal IP
192.168.2.65

Internal IP type
Ephemeral

Alias IP ranges

Subnet range
Primary (192.168.2.0/24)

Alias IP range
Example: 10.0.1.0/24 or /32

+ Add IP range

Hide alias IP ranges

External IP
None

Done Cancel

nic2: automationservernetwork serversubnet

Network interfaces		Subnetwork	Primary internal IP	Alias IP ranges	External IP	Network Tier	IP forwarding	Network details
nic0	automationmgmtnetwork	mgmtsubnet	192.168.1.62	—	adc-ha-instance1-ip1 (35.185.108.124)	Premium	Off	View details
nic1	automationclientnetwork	clientsubnet	192.168.2.8	192.168.2.7/32	None			View details
nic2	automationservernetwork	serversubnet	192.168.3.8	—	None			View details

After the failover, when the old primary becomes the new secondary, the alias IP addresses move from the old primary and is attached to the new primary.

After you have configured the VPX instances, you can configure the Virtual (VIP) and Subnet IP (SNIP) addresses. For more information, see [Configuring Citrix ADC-owned IP addresses](#).

Step 3. Configure high availability

After you've created the instances on Google Cloud Platform, you can configure high availability by using the Citrix ADC GUI or CLI.

Configure high availability by using the GUI

Step 1. Set up high availability in INC Enabled mode on both the nodes.

On the **primary node**, perform the following steps:

1. Log on to the instance with user name `nsroot` and instance ID of the node from GCP console as the password.
2. Navigate to **Configuration > System > High Availability > Nodes**, and click **Add**.
3. In the **Remote Node IP address** field, enter the private IP address of the management NIC of the secondary node.
4. Select the **Turn on INC (Independent Network Configuration) mode on self node** check box.
5. Click **Create**.

On the **secondary node**, perform the following steps:

1. Log on to the instance with user name `nsroot` and instance ID of the node from GCP console as the password.
2. Navigate to **Configuration > System > High Availability > Nodes**, and click **Add**.
3. In the **Remote Node IP address** field, enter the private IP address of the management NIC of the primary node.
4. Select the **Turn on INC (Independent Network Configuration) mode on self node** check box.
5. Click **Create**.

Before you proceed further, ensure that the Synchronization state of the secondary node is shown as **SUCCESS** in the **Nodes** page.

	ID	IP ADDRESS	HOST NAME	MASTER STATE	NODE STATE	INC	SYNCHRONIZATION STATE	SYNCHRONIZATION FAILURE RE
<input type="checkbox"/>	0	192.168.1.62		Primary	UP	ENABLED	ENABLED	-NA-
<input type="checkbox"/>	1	192.168.1.6		Secondary	UP	ENABLED	SUCCESS	-NA-

Note:

After the secondary node is synchronized with the primary node, the secondary node has the same log-on credentials as the primary node.

Step 2. Add Virtual IP address and Subnet IP address on both the nodes.

On the primary node, perform the following steps:

1. Navigate to **System > Network > IPs > IPv4s**, and click **Add**.
2. To create a client alias IP (VIP) address:

- a) Enter the Alias IP address and netmask configured for the client subnet in the VM instance.
 - b) In the **IP Type** field, select **Virtual IP** from the drop-down menu.
 - c) Click **Create**.
3. To create a server IP (SNIP) address:
- a) Enter the internal IP address of the server-facing interface of the primary instance and netmask configured for the server subnet.
 - b) In the **IP Type** field, select **Subnet IP** from the drop-down menu.
 - c) Click **Create**.

System > Network > IPs > IPv4s

IPs

IPv4s (3) IPv6s (1)

Add Edit Delete Statistics Select Action

Click here to search or you can enter Key : Value format

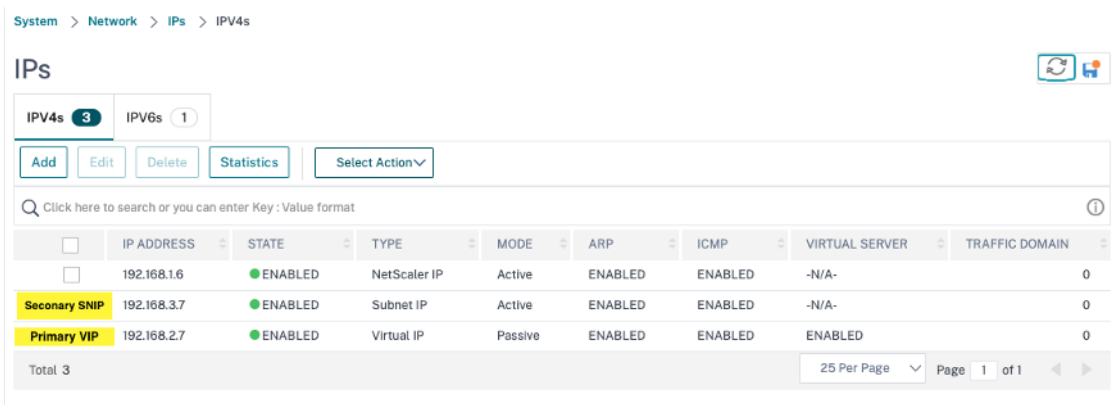
	IP ADDRESS	STATE	TYPE	MODE	ARP	ICMP	VIRTUAL SERVER	TRAFFIC DOMAIN
<input checked="" type="checkbox"/>	192.168.2.7	ENABLED	Virtual IP	Active	ENABLED	ENABLED	ENABLED	0
<input type="checkbox"/>	192.168.1.62	ENABLED	NetScaler IP	Active	ENABLED	ENABLED	-N/A-	0
<input checked="" type="checkbox"/>	192.168.3.8	ENABLED	Subnet IP	Active	ENABLED	ENABLED	-N/A-	0

Total 3

25 Per Page Page 1 of 1

On the secondary node, perform the following steps:

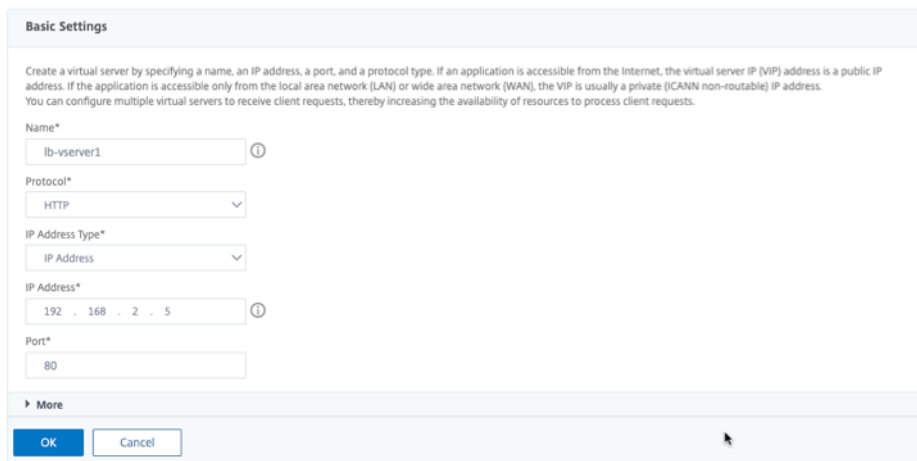
1. Navigate to **System > Network > IPs > IPv4s**, and click **Add**.
2. To create a client alias IP (VIP) address:
 - a) Enter the Alias IP address and netmask configured for the client subnet on the primary VM instance.
 - b) In the **IP Type** field, select **Subnet IP** from the drop-down menu.
 - c) Click **Create**.
3. To create a server IP (SNIP) address:
 - a) Enter the internal IP address of the server-facing interface of the secondary instance and netmask configured for the server subnet.
 - b) In the **IP Type** field, select **Subnet IP** from the drop-down menu.
 - c) Click **Create**.



Step 3. Add a load balancing virtual server on the primary node.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers > Add**.
2. Add the required values for Name, Protocol, IP Address Type (IP Address), IP Address (primary client alias IP address) and Port, and click **OK**.

Load Balancing Virtual Server



Step 4. Add a service or service group on the primary node.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Services > Add**.
2. Add the required values for Service Name, IP Address, Protocol and Port, and click **OK**.

Step 5. Bind the service or service group to the load balancing virtual server on the primary node.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers**.
2. Select the load balancing virtual server configured in **Step 3**, and click **Edit**.
3. In the **Service and Service Groups** tab, click **No Load Balancing Virtual Server Service Binding**.
4. Select the service configured in the **Step 4**, and click **Bind**.

Step 5. Save the configuration.

After a forced failover, the secondary becomes the new primary. The client alias IP (VIP) and the server alias IP (SNIP) from the old primary moves to the new primary.

Configure high availability by using the CLI

Step 1. Set up high availability in **INC Enabled** mode in both the instances by using the Citrix ADC CLI.

On the primary node, type the following command.

```
1 add ha node 1 <sec_ip> -inc ENABLED
```

On the secondary node, type the following command.

```
1 add ha node 1 <prim_ip> -inc ENABLED
```

The `sec_ip` refers to the internal IP address of the management NIC of the secondary node.

The `prim_ip` refers to the internal IP address of the management NIC of the primary node.

Step 2. Add VIP and SNIP on both nodes.

Type the following commands on the primary node:

```
1 add ns ip <primary_client_alias_ip> <subnet> -type VIP
```

Note:

Enter the Alias IP address and netmask configured for the client subnet in the VM instance.

```
1 add ns ip <primary_snip> <subnet> -type SNIP
```

The `primary_snip` refers to the internal IP address of the server-facing interface of the primary instance.

Type the following commands on the secondary node:

```
1 add ns ip <primary_client_alias_ip> <subnet> -type VIP
```

Note:

Enter the Alias IP address and netmask configured for the client subnet on the primary VM instance.

```
1 add ns ip <secondary_snip> <subnet> -type SNIP
```

The `secondary_snip` refers to the internal IP address of the server-facing interface of the secondary instance.

Note:

Enter the IP address and netmask configured for the server subnet in the VM instance.

Step 3. Add a virtual server on the primary node.

Type the following command:

```
1 add <server_type> vserver <vserver_name> <protocol> <
  primary_client_alias_ip> <port>
```

Step 4. Add a service or service group on the primary node.

Type the following command:

```
1 add service <service_name> <service_ip_address> <protocol> <port>
```

Step 5. Bind the service or service group to the load balancing virtual server on the primary node.

Type the following command:

```
1 bind <server_type> vserver <vserver_name> <service_name>
```

Note:

To save your configuration, type the command `save config`. Otherwise, the configurations are lost after you restart the instances.

Add back-end GCP Autoscaling service

September 12, 2024

Efficient hosting of applications in a cloud requires easy and cost-effective management of resources, depending on the application demand. To meet the increasing demand, you have to scale network resources upward. When demand subsides, you need to scale down to avoid the unnecessary cost of underutilized resources. To minimize the cost of running the application, you have to constantly monitor traffic, memory and CPU use, and so on. However, monitoring traffic manually is cumbersome. For the application environment to scale up or down dynamically, you must automate the processes of monitoring traffic and of scaling resources up and down whenever necessary.

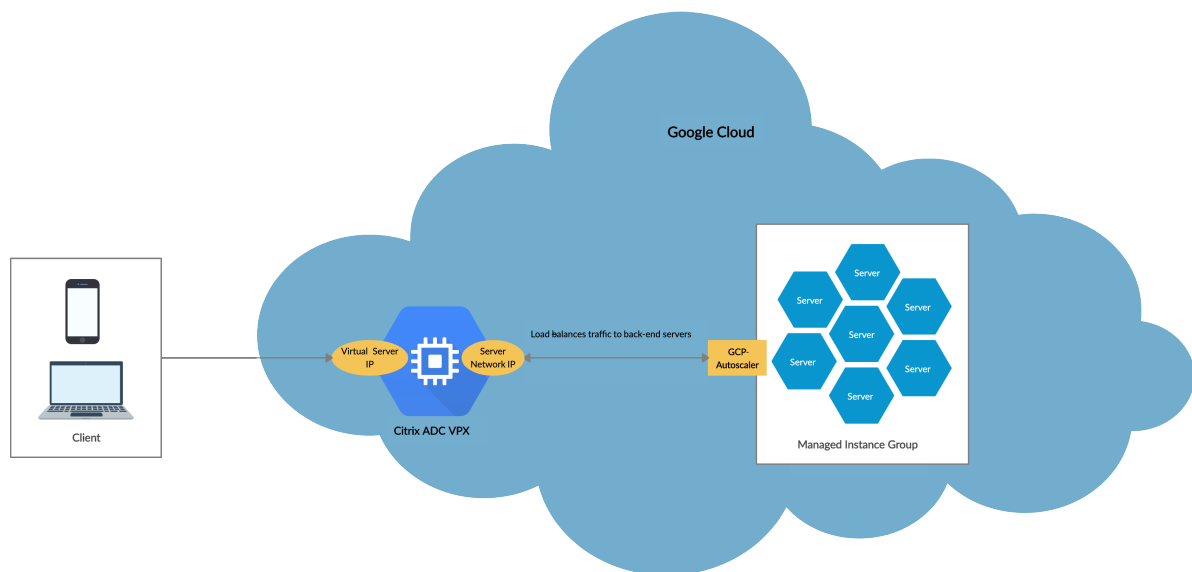
Integrated with the GCP Autoscaling service, the Citrix ADC VPX instance provides the following advantages:

- **Load balance and management:** Auto configures servers to scale up and scale down, depending on demand. The VPX instance auto detects managed instance groups in the back-end sub-

net and allows you to select the managed instance groups to balance the load. The virtual and subnet IP addresses are auto configured on the VPX instance.

- **High availability:** Detects managed instance groups that span multiple zones and load-balance servers.
- **Better network availability:** The VPX instance supports:
 - Back-end servers on same placement groups
 - Back-end servers on different zones


This diagram illustrates how the GCP Autoscaling service works in a Citrix ADC VPX instance acting as the load balancing virtual server.



Before you begin

Before you start using Autoscaling with your Citrix ADC VPX instance, you must complete the following tasks.

- Create a Citrix ADC VPX instance on GCP according to your requirement.
 - For more information about how to create a Citrix ADC VPX instance, see [Deploy a Citrix ADC VPX instance on the Google Cloud Platform](#).
 - For more information about how to deploy VPX instances in HA mode, see [Deploy a VPX high-availability pair on the Google Cloud Platform](#).
- Enable **Cloud Resource Manager API** for your GCP project.
- Allow full access to all Cloud APIs while creating the instances.



- Ensure your GCP service account has the following IAM permissions:

```

1  REQUIRED_INSTANCE_IAM_PERMS = [
2
3  "compute.instances.get",
4  "compute.zones.list",
5  "compute.instanceGroupManagers.list",
6  "compute.instanceGroupManagers.get"
7  ]

```

- To set up Autoscaling, ensure the following are configured:
 - Instance template
 - Managed Instance group
 - Autoscaling policy

Add the GCP Autoscaling service to a Citrix ADC VPX instance

You can add the Autoscaling service to a VPX instance with a single click by using the GUI. Complete these steps to add the Autoscaling service to the VPX instance:

1. Log on to the VPX instance by using your credentials for `nsroot`.
2. When you log on to the Citrix ADC VPX instance for the first time, you see the default Cloud Profile page. Select the GCP managed instance group from the drop-down menu and click **Create** to create a cloud profile.

Citrix ADC VPX Express (Freemium)

Dashboard Configuration Reporting Documentation Downloads

← Create Cloud Profile

Name
DemoCloudProfile

Virtual Server IP Address*
192.168.2.24

Load Balancing Server Protocol
HTTP

Load Balancing Server Port
80

Auto Scale Group*
ansible-mig-defaultuser-1585300924-

Auto Scale Group Protocol
HTTP

Auto Scale Group Port
80

Select this option to drain the connections gracefully. Else the connections will be dropped in the event of scale down.
 Graceful

Create Close

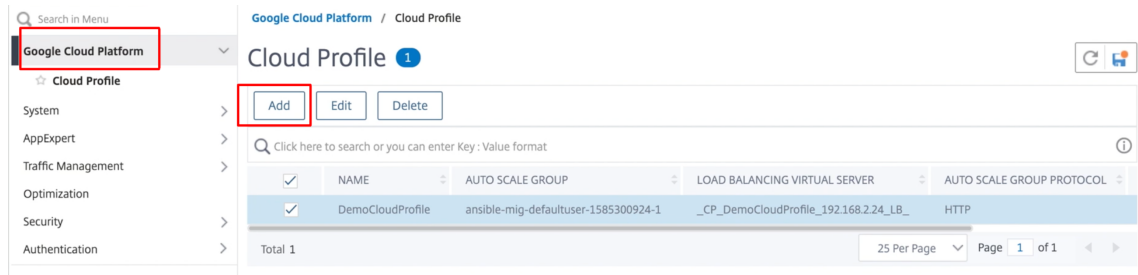
- The **Virtual Server IP Address** field is auto-populated from all the IP addresses associated with the instances.
- The **Autoscale Group** is prepopulated from the managed instance group configured on your GCP account.
- When selecting the **Autoscale Group Protocol** and **Autoscale Group Port**, ensure that your servers listen on the configured protocol and ports. Bind the correct monitor in the service group. By default, the TCP monitor is used.
- Clear the **Graceful** check box because it is not supported.

Note:

For SSL Protocol type Autoscaling, after you create the Cloud Profile, the load balance virtual server or service group is down because of a missing certificate. You can bind the certificate to the virtual server or service group manually.

3. After the first time logon if you want to create Cloud Profile, on the GUI go to **System > Google**

Cloud Platform > Cloud Profile and click **Add**.



The **Create Cloud Profile** configuration page appears.

← Create Cloud Profile

Name
DemoCloudProfile

Virtual Server IP Address*
192.168.2.24

Load Balancing Server Protocol
HTTP

Load Balancing Server Port
80

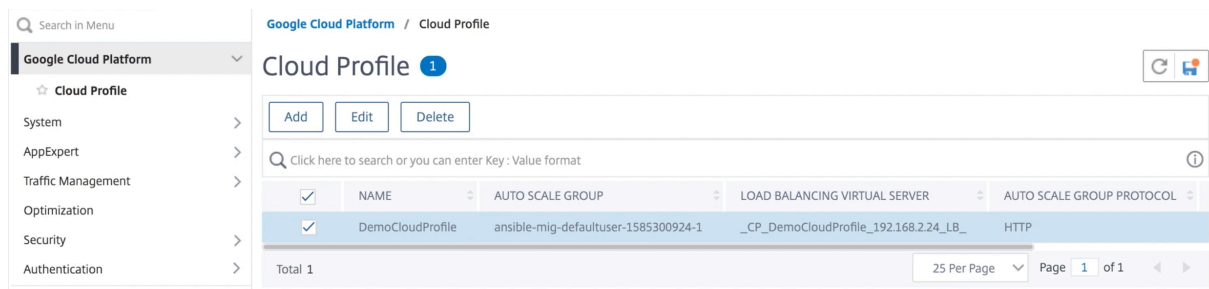
Auto Scale Group*
ansible-mig-defaultuser-1585300924-1

Auto Scale Group Protocol
HTTP

Auto Scale Group Port
80

Select this option to drain the connections gracefully. Else the connections will be dropped in the event of scale down.
 Graceful

Cloud Profile creates a Citrix ADC load-balancing virtual server and a service group with members as the servers of the managed instance group. Your back-end servers must be reachable through the SNIP configured on the VPX instance.



VIP scaling support for Citrix ADC VPX instance on GCP

September 6, 2024

A Citrix ADC appliance resides between the clients and the servers, so that client requests and server responses pass through it. In a typical installation, virtual servers configured on the appliance provide connection points that clients use to access the applications behind the appliance. The number of public virtual IP (VIP) addresses needed for a deployment varies on a case-by-case basis.

The GCP architecture restricts each interface on the instance to be connected to a different VPC. A VPC on GCP is a collection of subnets, and each subnet can span across zones of a region. In addition, GCP imposes the following limitation:

- There is a 1:1 mapping of number of public IP addresses to number of NICs. Only one public IP address can be assigned to a NIC.
- A maximum of only 8 NICs can be attached on a higher capacity instance type.

For example, an n1-standard-2 instance can have only 2 NICs, and the Public VIPs that can be added is limited to 2. For more information, see [VPC resource quotas](#).

To achieve higher scales of public virtual IP addresses on a Citrix ADC VPX instance, you can configure the VIP addresses as part of the metadata of the instance. The ADC VPX instance internally uses forwarding rules provided by the GCP to achieve VIP scaling. The ADC VPX instance also provides high availability to the VIPs configured.

After you configure VIP addresses as part of the metadata, you can configure an LB virtual server using the same IP that is used to create the forwarding rules. Thus, we can use forwarding rules to mitigate the limitations we have w.r.t scale in using public VIP addresses on an ADC VPX instance on GCP.

For more information on forwarding rules, see [Forwarding rules overview](#).

For more information on HA, see [High Availability](#).

Points to note

- Google charges some additional cost for each virtual IP forwarding rule. The actual cost depends on the number of entries created. The associated cost can be found from the Google pricing documents.
- Forwarding rules are applicable only for public VIPs. You can use alias IP addresses when the deployment needs private IP addresses as VIPs.
- You can create forwarding rules only for the protocols, which need the LB virtual server. VIPs can be created, updated, or deleted on the fly. You can also add a new load balancing virtual server with the same VIP address but with a different protocol.

Before you start

- Citrix ADC VPX instance must be deployed on GCP.
- External IP address must be reserved. For more information, see [Reserving a static external IP address](#).
- Ensure that your GCP service account has the following IAM permissions:

```
1  REQUIRED_IAM_PERMS = [  
2  "compute.addresses.list",  
3  "compute.addresses.get",  
4  "compute.addresses.use",  
5  "compute.forwardingRules.create",  
6  "compute.forwardingRules.delete",  
7  "compute.forwardingRules.get",  
8  "compute.forwardingRules.list",  
9  "compute.instances.use",  
10 "compute.subnetworks.use",  
11 "compute.targetInstances.create",  
12 "compute.targetInstances.get",  
13 "compute.targetInstances.use",  
14 ]
```

Configure external IP addresses for VIP scaling on Citrix ADC VPX instance

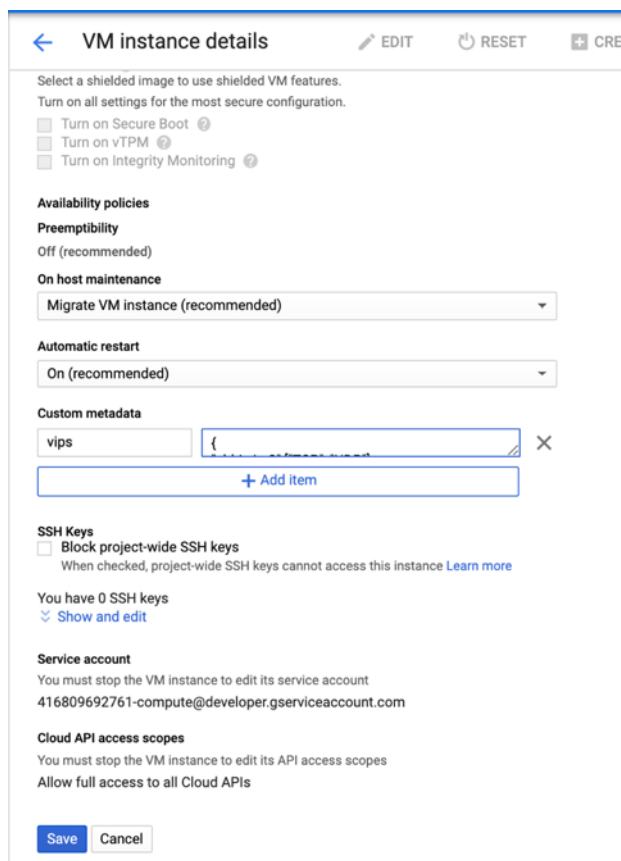
1. In the Google Cloud Console, navigate to the **VM Instances** page.
2. Create a new VM instance or use an existing instance.
3. Click the instance name. On the **VM instance details** page, click **Edit**.
4. Update the **Custom metadata** by entering the following:
 - Key = vips

- Value = Provide a value in the following JSON format:

```
{
  "Name of external reserved IP": [list of protocols],
}
```

GCP supports the following protocols:

- AH
- ESP
- ICMP
- SCT
- TCP
- UDP



For more information, see [Custom metadata](#).

Example for Custom metadata:

```
{
  "external-ip1-name":["TCP", "UDP"],
  "external-ip2-name":["ICMP", "AH"]
}
```

In this example, the ADC VPX instance internally creates one forwarding rule for each IP, protocol pair. The metadata entries are mapped to the forwarding rules. This example helps you understand how many forwarding rules are created for a metadata entry.

Four forwarding rules are created as follows:

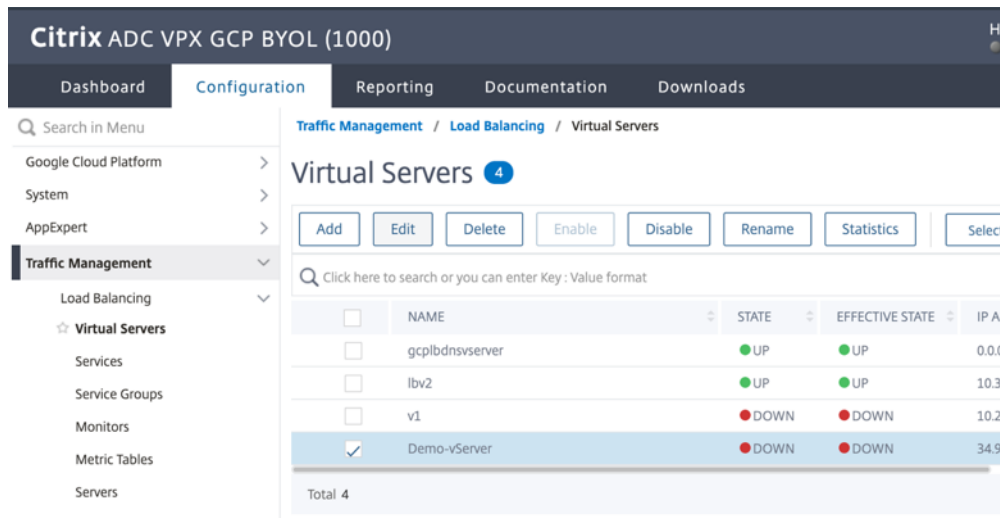
- a) external-ip1-name and TCP
- b) external-ip1-name and UDP
- c) external-ip2-name and ICMP
- d) external-ip2-name and AH

5. Click **Save**.

Setting up a load balancing virtual server with external IP address on a Citrix ADC VPX instance

Step 1. Add a load balancing virtual server.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers > Add**.



2. Add the required values for Name, Protocol, IP Address Type (IP Address), IP Address (External IP address of the forwarding rule that is added as VIP on ADC) and Port, and click **OK**.

Dashboard Configuration Reporting Documentation

← Load Balancing Virtual Server

Basic Settings

Create a virtual server by specifying a name, an IP address, a port, and a protocol type. If an application address is a public IP address. If the application is accessible only from the local area network (LAN), use a private (ICANN non-routable) IP address.
You can configure multiple virtual servers to receive client requests, thereby increasing the available capacity.

Name*
Demo-vServer ⓘ

Protocol*
HTTP ▾

IP Address Type*
IP Address ▾

IP Address*
34 . 93 . 61 . 42 ⓘ

Port*
80

▶ More

OK Cancel

Step 2. Add a service or service group.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Services > Add**.
2. Add the required values for Service Name, IP Address, Protocol and Port, and click **OK**.

← **Load Balancing Service**

Basic Settings

Service Name* ⓘ

New Server Existing Server

IP Address* ⓘ

Protocol* ▾

Port*

▶ More

Step 3. Bind the service or service group to the load balancing virtual server.

1. Navigate to **Configuration > Traffic Management > Load Balancing > Virtual Servers**.
2. Select the load balancing virtual server configured in **Step 1**, and click **Edit**.
3. In the **Service and Service Groups** page, click **No Load Balancing Virtual Server Service Binding**.

← **Load Balancing Virtual Server**

Load Balancing Virtual Server | [Export as a Template](#)

Basic Settings

Name	Demo-vServer	Listen Priority	-
Protocol	HTTP	Listen Policy Expression	NONE
State	DOWN	Redirection Mode	IP
IP Address	34.93.61.42	Range	1
Port	80	IPset	-
Traffic Domain	0	RHI State	PASSIVE
		AppFlow Logging	ENABLED
		Retain Connections on Cluster	NO
		TCP Probe Port	-

Services and Service Groups

- No Load Balancing Virtual Server Service Binding >
- No Load Balancing Virtual Server ServiceGroup Binding >

4. Select the service configured in the **Step 3**, and click **Bind**.

Service Binding

Service Binding

Select Service* > ⓘ

Binding Details

Weight

5. Save the configuration.

Troubleshoot a VPX instance on GCP

September 6, 2024

Google Cloud Platform (GCP) provides console access to a Citrix ADC VPX instance. You can debug only if the network is connected. To view an instance's System Log, access the console and check **System Log files**.

Citrix supports fee based Citrix ADC VPX instances (utility license with hourly fee) on GCP. To file a support case, find your GCP account number and support PIN code, and call Citrix support. You are asked to provide your name and email address. To find the support PIN, log on to the VPX GUI and navigate to the **System** page.

Here is an example of a system page showing the support PIN.

The screenshot shows the 'System' page in the Citrix ADC VPX GUI. The left sidebar contains a search bar and a menu with 'Google Cloud Platform' highlighted. The main content area is titled 'System' and includes tabs for 'System Information', 'System Sessions (1)', and 'System Network'. Below the tabs are buttons for 'System Upgrade', 'Reboot', 'Migration', 'Statistics', 'Call Home', and 'Citrix ADM Service Connect'. The 'System Information' section displays the following details:

Citrix ADC IP Address	10.160.15.230
Netmask	255.255.240.0
Node	Standalone
Technical Support PIN	4051153
Time Zone	Coordinated Universal Time
System Time	Sat, 11 Jul 2020 01:56:22 UTC
Last Config Changed Time	Sat, 11 Jul 2020 01:53:09 UTC
Last Config Saved Time	Sat, 11 Jul 2020 01:53:12 UTC

The 'Technical Support PIN' value, 4051153, is highlighted with a red box in the original image.

Jumbo frames on Citrix ADC VPX instances

September 9, 2024

Citrix ADC VPX appliances support receiving and transmitting jumbo frames containing up to 9216 bytes of IP data. Jumbo frames can transfer large files more efficiently than it is possible with the standard IP MTU size of 1500 bytes.

A Citrix ADC appliance can use jumbo frames in the following deployment scenarios:

- Jumbo to Jumbo. The appliance receives data as jumbo frames and sends it as jumbo frames.

- Non-Jumbo to Jumbo. The appliance receives data as regular frames and sends it as jumbo frames.
- Jumbo to Non-Jumbo. The appliance receives data as jumbo frames and sends it as regular frames.

For more information, see [Configuring Jumbo Frames Support on a Citrix ADC Appliance](#).

Jumbo frames support is available on Citrix ADC VPX appliances running on the following virtualization platforms:

- VMware ESX
- Linux-KVM Platform
- Citrix XenServer
- Amazon Web Services (AWS)

Jumbo frames on VPX appliances work similar to jumbo frames on MPX appliances. For more information on Jumbo Frames and its use cases, see [Configuring Jumbo Frames on MPX appliances](#). The use cases of jumbo frames on MPX appliances also apply to VPX appliances.

Configure jumbo frames for a VPX instance running on VMware ESX

Perform the following tasks to configure jumbo frames on a Citrix ADC VPX appliance running on the VMware ESX server:

1. Set the MTU of the interface or channel of the VPX appliance to a value in the range 1501–9000. Use the CLI or GUI to set the MTU size. The Citrix ADC VPX appliances running on VMware ESX support receiving and transmitting jumbo frames containing up to only 9000 bytes of IP data.
2. Set the same MTU size on the corresponding physical interfaces of the VMware ESX server by using its management applications. For more information about setting the MTU size on the physical interfaces of VMware ESX, see <http://vmware.com/>.

Configure jumbo frames for a VPX instance running on Linux-KVM server

Perform the following tasks to configure jumbo frames on a Citrix ADC VPX appliance running on a Linux-KVM Server:

1. Set the MTU of the interface or channel of the VPX appliance to a value in the range 1501–9216. Use the Citrix ADC VPX CLI or GUI to set the MTU size.
2. Set the same MTU size on the corresponding physical interfaces of a Linux-KVM Server by using its management applications. For more information about setting the MTU size on the physical interfaces of Linux-KVM, see <http://www.linux-kvm.org/>.

Configure jumbo frames for a VPX instance running on Citrix XenServer

Perform the following tasks to configure jumbo frames on a Citrix ADC VPX appliance running on Citrix XenServer:

1. Connect to the XenServer using XenCenter.
2. Shut down all the VPX instances that use the Networks for which the MTU must be changed.
3. On the **Networking** tab, select the network - network 0/1/2.
4. Select **Properties** and edit MTU.

After configuring the jumbo frames on the XenServer, you can configure the jumbo frames on the ADC appliance. For more information, see [Configuring Jumbo Frames Support on a Citrix ADC Appliance](#).

Configure jumbo frames for a VPX instance running on AWS


Host-level configuration is not required for VPX on Azure. To configure Jumbo Frames on VPX, follow the steps given in [Configuring Jumbo Frames Support on a Citrix ADC Appliance](#).

Automate deployment and configurations of Citrix ADC

September 12, 2024

Citrix ADC provides multiple tools to automate your ADC deployments and configurations. This document provides a brief summary of various automation tools and references to various automation resources that you can use to manage ADC configurations.

The following illustration provides an overview of Citrix ADC automation in a hybrid multi cloud (HMC) environment.



Automate Citrix ADC using Citrix ADM

Citrix ADM acts as automation control point to your distributed ADC infrastructure. The Citrix ADM provides comprehensive set of automation capabilities from provisioning ADC appliances to upgrading it. The following are the key automation features of ADM:

- [Provisioning Citrix ADC VPX instances on AWS](#)
- [Provisioning Citrix ADC VPX instances on Azure](#)
- [StyleBooks](#)
- [Configuration jobs](#)
- [Configuration audit](#)
- [ADC upgrades](#)
- [SSL certificate management](#)
- [Integrations - GitHub, ServiceNow, Event notifications integrations](#)

Citrix ADM blogs and videos on automation

- [Application migrations using StyleBooks](#)
- [Integrate ADC configurations with CI/CD using ADM StyleBooks](#)
- [Simplifying public cloud Citrix ADC deployments through ADM](#)
- [10 ways Citrix ADM service supports easier Citrix ADC upgrades](#)

Citrix ADM also provides APIs for its various capabilities that integrate Citrix ADM and Citrix ADC as part of the overall IT automation. For more information, see [Citrix ADM Service APIs](#).

Automate Citrix ADC using Terraform

Terraform is a tool that takes infrastructure as code approach to provision and manage cloud, infrastructure, or service. Citrix ADC terraform resources are available in GitHub for use. Refer GitHub for detailed documentation and usage.

- [Citrix ADC Terraform modules to configure ADC for various use cases such as Load Balancing and GSLB](#)
- [Terraform cloud scripts to deploy ADC in AWS](#)
- [Terraform cloud scripts to deploy ADC in Azure](#)

Videos on Terraform for ADC automation

- [Automate your Citrix ADC deployments with Terraform](#)
- [Provision and configure ADC in HA setup in AWS using Terraform](#)

Automate Citrix ADC using Ansible

Ansible is an open-source software provisioning, configuration management, and application-deployment tool enabling infrastructure as code. Citrix ADC Ansible modules and sample playbooks can be found in GitHub for use. Refer GitHub for detailed documentation and usage.

- [Ansible modules to configure ADC](#)
- [Ansible modules for ADM](#)

Citrix is a certified Ansible Automation Partner. Users having Red Hat Ansible Automation Platform subscription can access Citrix ADC Collections from [Red Hat Automation Hub](#).

Terraform and Ansible automation blogs

- [Terraform and Ansible Automation for app delivery and security](#)

Public cloud templates for ADC deployments

Public cloud templates simplify provisioning of your deployments in public clouds. Different Citrix ADC templates are available for various environments. For usage details, refer to respective GitHub repositories.

AWS CFTs:

- [CFTs to provision Citrix ADC VPX on AWS](#)

Azure Resource Manager (ARM) Templates:

- [ARM templates to provision Citrix ADC VPX on Azure](#)

Google Cloud Deployment Manager (GDM) Templates:

- [GDM templates to provision Citrix ADC VPX on Google](#)

Videos on Templates

- [Deploy Citrix ADC HA in AWS using CloudFormation Template](#)
- [Deploy Citrix ADC HA across Availability Zones using AWS QuickStart](#)
- [Citrix ADC HA deployment in GCP using GDM templates](#)

NITRO APIs

The Citrix ADC NITRO protocol allows you to programmatically configure and monitor the Citrix ADC appliance by using Representational State Transfer (REST) interfaces. Therefore, NITRO applications can be developed in any programming language. For applications that must be developed in Java or .NET or Python, NITRO APIs are exposed through relevant libraries that are packaged as separate Software Development Kits (SDKs).

- [NITRO API documentation](#)
- [Sample ADC use case configuration using NITRO API](#)

FAQs

November 21, 2024

The following section helps you to categorize the FAQs based on Citrix Application Delivery Controller (ADC) VPX.

- Feature and functionality
- Encryption
- Pricing and packaging
- [Citrix ADC VPX Express] (#citrix-ADC-VPX-Express-and-90-day-free-trial)
- Hypervisor
- Capacity planning or sizing

- System requirements
- Other technical FAQs

Feature and functionality

What is Citrix ADC VPX?

Citrix ADC VPX is a virtual ADC appliance that can be hosted on a Hypervisor installed on industry standard servers.

Does Citrix ADC VPX include all the web application optimization functionality as ADC appliances?

Yes. Citrix ADC VPX includes all load balancing, traffic management, application acceleration, application security (including Citrix ADC Gateway and Citrix Application Firewall), and offload functionality. For a complete overview of the Citrix ADC feature and functionality, see [Application delivery your way](#).

Are there any limitations with Citrix Application Firewall when using it on Citrix ADC VPX?

Citrix Application Firewall on Citrix ADC VPX provides the same security protections as it does on Citrix ADC appliances. Performance or throughput of Citrix Application Firewall varies by platform.

Are there any differences between Citrix ADC Gateway on Citrix ADC VPX and Citrix ADC Gateway on Citrix ADC appliances?

Functionally, they are identical. Citrix ADC Gateway on Citrix ADC VPX supports all the Citrix ADC Gateway features available in Citrix ADC software release 9.1. However, because Citrix ADC appliances provide dedicated SSL acceleration hardware, it offers greater SSL VPN scalability than a Citrix ADC VPX instance.

Other than the obvious difference of being able to run on a Hypervisor, how does Citrix ADC VPX differ from Citrix ADC physical appliances?

There are two main areas where customers see differences in behavior. The first is Citrix ADC VPX cannot offer the same performance as many Citrix ADC appliances. The second is that while Citrix ADC appliances incorporate its own L2 networking functionality, Citrix ADC VPX relies upon the Hypervisor for its L2 networking services. Generally, it does not limit how the Citrix ADC VPX can be deployed.

There can be certain L2 functionality that is configured on a physical Citrix ADC appliance must be configured on the underlying Hypervisor.

How does Citrix ADC VPX play a role in the Application Delivery market?

Citrix ADC VPX changes the game in the application delivery market in the following ways:

- By making a Citrix ADC appliance even more affordable, Citrix ADC VPX enables any IT organization to deploy a Citrix ADC appliance. It is not just for their most mission-critical web applications, but for all of their Web applications.
- Citrix ADC VPX allows customers to further converge networking and virtualization within their data centers. Citrix ADC VPX cannot only be used to optimize web applications hosted on virtualized servers. It also enables web application delivery itself to become a virtualized service that can be easily and rapidly deployed anywhere. IT organizations use the standard data center processes for tasks such as provisioning, automation, and charge-back for the web application delivery infrastructure.
- Citrix ADC VPX opens up new deployment architectures that are not practical if only physical appliances are used. Citrix ADC VPX and Citrix ADC MPX appliances can be used basis, tailored to the individual needs of each respective application to handle processor-intensive actions such as compression and application firewall inspection. At the data center edge, Citrix ADC MPX appliances handle high-volume network-wide tasks such as initial traffic distribution, SSL encryption or decryption, denial of service (DoS) attack prevention, and global load balancing. Pairing high-performance Citrix ADC MPX appliances with easy-to-deploy Citrix ADC VPX virtual appliance brings unparalleled flexibility and customization capabilities to modern, large-scale, data center environments while also reducing overall data center costs.

How does Citrix ADC VPX fit into our Citrix delivery center strategy?

With the availability of Citrix ADC VPX, the entire Citrix delivery center offering is available as a virtualized offering. The entire Citrix delivery center benefits from the powerful management, provisioning, monitoring, and reporting capabilities available in Citrix XenCenter. This can be deployed rapidly into almost any environment, and managed centrally from anywhere. With one integrated, virtualized application delivery infrastructure, organizations can deliver desktops, client-server applications, and Web applications.

Encryption

Does Citrix ADC VPX support SSL offload?

Yes. However, Citrix ADC VPX does all SSL processing in software, so Citrix ADC VPX does not offer the same SSL performance as Citrix ADC appliances. Citrix ADC VPX can support up to 750 new SSL transactions per second.

Does third-party SSL cards installed on the server hosting Citrix ADC VPX accelerate SSL encryption or decryption?

No. Supporting third-party SSL cards cannot associate the Citrix ADC VPX to specific hardware implementations. It greatly diminishes an organizations ability to flexibly host Citrix ADC VPX anywhere within the data center. Citrix ADC MPX appliances must be used when more SSL throughput than Citrix ADC VPX provides is required.

Does Citrix ADC VPX support the same encryption ciphers as physical Citrix ADC appliances?

VPX supports all encryption ciphers as physical Citrix ADC appliances, except the ECDSA.

What is the SSL transactions throughput of Citrix ADC VPX?

See [Citrix ADC VPX data sheet](#) for information on SSL transactions throughput.

Pricing and packaging

How is Citrix ADC VPX packaged?

Citrix ADC VPX selection is similar to the selection of Citrix ADC appliances. First, the customer selects the Citrix ADC edition based on its functionality requirements. Then, the customer selects the specific Citrix ADC VPX bandwidth tier based on their throughput requirements. Citrix ADC VPX is available in Standard, Advanced, and Premium Editions. Citrix ADC VPX offers from 10 Mbps (VPX 10) to 100 Gbps (VPX 100G). More details can be found in the Citrix ADC VPX data sheet.

Is Citrix ADC VPX priced the same for all Hypervisors?

Yes.

Are the same Citrix ADC SKUs used for VPX on all Hypervisors?

Yes.

Can a Citrix ADC VPX license be moved from one Hypervisor to another (For example from VMware to Hyper-V)?

Yes. Citrix ADC VPX licenses are independent of the underlying Hypervisor. If you decide to move the Citrix ADC VPX virtual machine from one Hypervisor to another, you do not have to get a new license. However, you might need to rehost the existing Citrix ADC VPX license.

Can Citrix ADC VPX instances be upgraded?

Yes. Both the throughput limits and Citrix ADC family edition can be upgraded. Upgrade SKUs for both types of upgrade are available.

If I want to deploy Citrix ADC VPX in a high availability pair, how many licenses do I need?

As with Citrix ADC physical appliances, a Citrix ADC high availability configuration requires two active instances. Therefore, the customer must purchase two licenses.

Citrix ADC VPX Express and 90-day fee trial

Does Citrix ADC VPX Express include all Citrix ADC standard functionality? Does it include Citrix ADC Gateway and load balancing for Citrix Virtual Apps (formerly XenApp) Web Interface and XML broker?

Yes. Citrix ADC VPX Express includes full Citrix ADC Standard functionality. Starting from Citrix ADC release 12.0–56.20, Citrix modified the VPX express behavior.

Does Citrix ADC VPX Express include all Citrix ADC standard functionality? Does it include Citrix ADC Gateway and load balancing for Citrix Virtual Apps Web Interface and XML broker?

Starting from Citrix ADC release 12.0–56.20, VPX Express offers the Citrix ADC Standard Edition feature set, except Gateway functionality. Earlier to the 12.0–56.20 release, VPX expresses includes all features in the standard edition.

Does Citrix ADC VPX Express require a license?

With the new Citrix ADC VPX Express release (12.0–56.20 and onwards), VPX Express is free and requires no license files to install and comes with no commitment. If you have a VPX Express license already, then the prior VPX Express behavior is preserved. If the VPX Express *license file* is removed and the 12.0–56.20 and onwards release is used, the new VPX express behavior takes effect.

Does the Citrix ADC VPX Express license expire?

With the new VPX express, no. There is no license and no expiry date. If you have a VPX express license already, the license expires one year after download.

Does Citrix ADC VPX Express include the five free Citrix ADC Gateway concurrent licenses?

Yes, if you own a VPX express license.

Is there a limit to how many Citrix ADC VPX Expresses a customer can download?

Five.

Does Citrix ADC VPX Express support the same encryption ciphers as Citrix ADC MPX appliances?

For general availability, all the same strong encryption ciphers supported on Citrix ADC appliances are available on Citrix ADC VPX and Citrix ADC VPX Express. It is subjected to the same import or export regulations.

Can I file technical support cases for Citrix ADC VPX Express?

No. A retail Citrix ADC VPX license such as, VPX-10, VPX-200, VPX-1000, VPX- 3000 is required to file technical support cases. However, Citrix ADC VPX Express users are free to use both the Citrix ADC VPX Knowledge Center, and request help from the community using the Z discussion forums.

Can Citrix ADC VPX Express be upgraded to a retail version?

Yes. Simply purchase the retail Citrix ADC VPX license you need, and then apply the corresponding license to the Citrix ADC VPX Express instance.

Hypervisor

What VMware versions does Citrix ADC VPX support?

Citrix ADC VPX supports both VMware ESX and ESXi for versions 3.5 or later. For more information, see [Support matrix and usage guidelines](#)

For VMware, how many virtual network interfaces can you allocate to a VPX?

You can allocate up to 10 virtual network interfaces to a Citrix ADC VPX.

From vSphere, how can we access the Citrix ADC VPX command line?

The VMware vSphere client provides built-in access to the Citrix ADC VPX command line through a console tab. Also, you can use any SSH or Telnet client to access the command line. You can use the NSIP address of the Citrix ADC VPX in the SSH or Telnet client.

How can you access the Citrix ADC VPX GUI?

To access the Citrix ADC VPX GUI, type the NSIP of the Citrix ADC VPX, for example, <http://NSIP address> in the address field of any browser.

Can two Citrix ADC VPX instances installed on the same VMware ESX be configured in a high availability setup?

Yes, but it is not recommended. A hardware failure would affect both Citrix ADC VPX instances.

Can two Citrix ADC VPX instances running on two different VMware ESX systems be configured in a high availability setup?

Yes. It is recommended in a high availability setup.

For the VMware, are interface related events supported on Citrix ADC VPX?

No. Interface related events are not supported.

For the VMware, are tagged VLANs supported on Citrix ADC VPX?

Yes. Citrix ADC tagged VLANs are supported on Citrix ADC VPX from release 11.0 and higher. For more information, see [Citrix documentation](#).

For VMware, are link aggregation and LACP supported on Citrix ADC VPX?

No. Link Aggregation and LACP are not supported for Citrix ADC VPX. Link aggregation must be configured at the VMware level.

How do we access Citrix ADC VPX documentation?

The documentation is available from the Citrix ADC VPX GUI. After logging in, select the **Documentation** tab.

Capacity planning or sizing

What performance can I expect with Citrix ADC VPX?

Citrix ADC VPX offers good performance. See [Citrix ADC VPX data sheet](#) for a specific performance level achievable using Citrix ADC VPX.

Given that server CPU power varies, how can we estimate the maximum performance of a Citrix ADC instance?

Using a faster CPU can result in higher performance (up to the maximum allowed by the license), while using a slower CPU can certainly limit the performance.

Are Citrix ADC VPX bandwidth or throughput limits for inbound only traffic, or both inbound and outbound traffic?

Citrix ADC VPX bandwidth limits are enforced for traffic inbound to the Citrix ADC only, regardless of whether the request traffic or response traffic. It indicates that a Citrix ADC VPX-1000 (for example) can process both 1 Gbps of inbound traffic and 1 Gbps of outbound traffic simultaneously. Inbound and outbound traffic is not the same as request and response traffic. To the Citrix ADC, both traffic coming from endpoints (request traffic) and traffic coming from origin servers (response traffic) is “inbound” (that is, coming into the Citrix ADC).

Can multiple instances of Citrix ADC VPX be run on the same server?

Yes. However, ensure that the physical server has enough CPU and I/O capacity to support the total workload running on the host, or Citrix ADC VPX performance can be impacted.

If more than one instance of Citrix ADC VPX is running on a physical server, what is the minimum hardware requirement per Citrix ADC VPX instance?

Each Citrix ADC VPX instance must be allocated 2 GB of physical RAM, 20 GB of hard disk space, and 2 vCPUs.

Note:

The Citrix ADC VPX is a latency-sensitive, high-performance virtual appliance. To deliver its expected performance, the appliance requires vCPU reservation, memory reservation, vCPU pinning on the host. Also, hyper threading must be disabled on the host. If the host does not meet these requirements, issues such as high-availability failover, CPU spike within the VPX instance, sluggishness in accessing the VPX CLI, pit boss daemon crash, packet drops, and low throughput occur.

Make sure that every VPX instance meets the predefined conditions.

Can I host Citrix ADC VPX and other applications on the same server?

Yes. For example, Citrix ADC VPX, Citrix Virtual Apps Web Interface and Citrix Virtual Apps XML Broker can all be virtualized and can run on the same server. For best performance, ensure that the physical host has enough CPU and I/O capacity to support all the running workloads.

Will adding CPU cores to a single Citrix ADC VPX instance increase the performance of that instance?

Yes, adding CPU cores can improve NetScaler VPX performance, provided the NetScaler VPX instance is licensed for the extra vCPUs. NetScaler VPX can support up to 20 vCPUs (for 41 Gbps - 100 Gbps performance), depending on the configuration and performance tier. More vCPUs can help increase throughput, especially in high-performance scenarios. However, the impact on performance also depends on factors like the network drivers (for example, PCI passthrough or SR-IOV) and the specific workload. For information on number of vCPUs supported for different VPX performance tiers, see [NetScaler VPX data sheet](#).

Why Citrix ADC VPX looks like consuming more than 90% of the CPU even though it is idle?

It is normal behavior and Citrix ADC appliances exhibit the same behavior. To see the true extent of Citrix ADC VPX CPU utilization, use the `stat CPU` command in the Citrix ADC CLI, or view Citrix ADC VPX CPU utilization from the Citrix ADC GUI. The Citrix ADC packet processing engine is always “looking for work,” even when there is no work to be done. Therefore, it does everything to take control of the CPU and not release it. On a server installed with Citrix ADC VPX and nothing else, results in looking like (from the Hypervisor perspective) that Citrix ADC VPX is consuming the entire CPU. Looking at the CPU utilization from “inside Citrix ADC” (by using the CLI or the GUI) provides a picture of Citrix ADC VPX CPU capacity being used.

System requirements**What are the minimum hardware requirements for Citrix ADC VPX?**

The following table explains the minimum hardware requirements for NetScaler VPX.

Type	Requirements
Processor	Dual core server with Intel Xeon.
Memory	Minimum 2 GB. However, 4 GB is recommended.
Disk	Minimum 20 GB hard drive.
Hypervisor	Citrix Hypervisor 5.6 or later, VMware ESX/ESXi 3.5 or later, or Windows Server 2008 R2 with Hyper-V
Network Connectivity	100 Mbps minimum, but 1 Gbps is recommended.
NIC	A NIC compatible with the Hypervisor you are using.

Note:

For critical deployments, 4 GB memory is preferred for NetScaler VPX. With 2 GB memory, NetScaler VPX operates in a very memory-constrained environment. This might lead to scale, performance, or stability related issues.

For more information on system requirements, see [NetScaler VPX data sheet](#).

Note:

AMD processors are not supported.

What is Intel VT-x?

These features, sometimes referred to as “hardware assist” or “virtualization assist,” trap sensitive or privileged CPU instructions run by the guest OS out to the Hypervisor. This simplifies hosting guest OSs (BSD for a Citrix ADC VPX) on the Hypervisor.

How common are VT-x?

Many servers have virtualization assistance features (such as VT-x or AMD-V) disabled by default in the BIOS settings. Before concluding that you cannot run NetScaler VPX, check the BIOS configuration. If virtualization support is disabled, you may need to enable it in the BIOS to ensure your server can properly run virtualized applications like NetScaler VPX.

Is there a hardware compatibility list (HCL) for Citrix ADC VPX?

As long as the server supports Intel VT-x, Citrix ADC VPX must run on any server compatible with the underlying Hypervisor. See the Hypervisor HCL for a comprehensive list of supported platforms.

What version of Citrix ADC OS is Citrix ADC VPX based on?

Citrix ADC VPX is based on Citrix ADC 9.1 or later releases.

Since Citrix ADC VPX runs on BSD, can it be run natively on a server with BSD Unix installed?

No. Citrix ADC VPX requires the Hypervisor to run. Detailed Hypervisor supports can be found in [Citrix ADC VPX data sheet](#).

Other technical FAQs

Does link aggregation on a physical server with multiple NIC’s work?

LACP is not supported. For the Citrix Hypervisor, Static link aggregation is supported and has limits of four channels and seven virtual interfaces. For VMware, static link aggregation is not supported within Citrix ADC VPX, but can be configured at the VMware level.

Is MAC based forwarding (MBF) supported on VPX? Is there any change from the Citrix ADC appliance implementation?

MBF is supported and it behaves the same way as with the Citrix ADC appliance. The Hypervisor basically switches all the packets received from Citrix ADC VPX to the outside and conversely.

How is the Citrix ADC VPX upgrade process carried out?

Upgrades are performed the same way as for Citrix ADC appliances: download a kernel file and use install ns or the upgrade utility in the GUI.

What is the size of the /var partition when using the default image for VPX? How to increase the disk space?

The size of the root disk is limited to 20 GB to keep the disk image small.

If you want to increase the `/var/core/` or the `/var/crash/` directory space, attach an extra disk. To increase the `/var` size, currently, you must attach an extra disk and create a symbolic link to `/var`, after copying the critical contents to the new disk.

What can we expect to regard the NetScaler VPX build numbering and interoperability with other builds?

Citrix ADC VPX has similar build numbering as the 9.1. Cl (classic) and 9.1. Nc (nCore) release, for instance 9.1_97.3.vpx, 9.1_97.3.nc, and 9.1_97.3.cl.

Can the Citrix ADC VPX be a part of a high availability setup with a Citrix ADC appliance?

Not a supported configuration.

Are all the interfaces visible in Citrix ADC VPX directly related to the number of interfaces on the Hypervisor?

No. You can add up to seven interfaces (10 for VMware) through the Citrix ADC VPX configuration utility with only one physical NIC on the Hypervisor.

Can Citrix Hypervisor XenMotion or VMware VMotion or Hyper-V live migration be used to move active instances of Citrix ADC VPX?

Citrix ADC VPX does not support XenMotion or Hyper-V live migration. VMotion is supported from Citrix ADC 12.1 release onwards.



© 2025 Cloud Software Group, Inc. All rights reserved. This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of Cloud Software Group, Inc. This and other products of Cloud Software Group may be covered by registered patents. For details, please refer to the Virtual Patent Marking document located at <https://www.cloud.com/legal>. Citrix, the Citrix logo, NetScaler, and the NetScaler logo and other marks appearing herein are either registered trademarks or trademarks of Cloud Software Group, Inc. and/or its subsidiaries in the United States and/or other countries. Other marks are the property of their respective owner(s) and are mentioned for identification purposes only. Please refer to Cloud SG's Trademark Guidelines and Third Party Trademark Notices (<https://www.cloud.com/legal>) for more information.