Book of Network Services - Flow Virtual Networking

Flow Virtual Networking allows you to create completely isolated virtual networks that are separated from the physical network. You can do multi-tenant network provisioning with overlapping IP addresses, self-service network provisioning, and IP address preservation.

Supported Configurations

Core Use Cases:

- Multi-tenant networking
- Network isolation
- Overlapping IP addresses
- Self-service network creation
- VM IP mobility
- Hybrid Cloud connectivity

Management interfaces(s):

- Prism Central (PC)

Supported Environment(s):

- On-Premises:
  - AHV

Upgrades:

- Major Feature Upgrades Depend on
  - Prism Central
  - AHV

- Minor Feature Upgrades Included in LCM
  - Advanced Network Controller
  - Network Gateway (VPN)

Implementation Constructs

Flow Virtual Networking introduces a number of new constructs to provide a complete networking solution.

- VPCs
- Overlay Subnets
- Routes
- Policies
- External Networks
  - NAT
  - Routed (NoNAT)

- Network Gateway
  - Layer 3 VPN
  - Layer 2 Extended Subnets with VXLAN

VPCs (Virtual Private Clouds)

The VPC or Virtual Private Cloud is the basic unit of Flow Virtual Networking. Each VPC is an isolated network namespace with a virtual router instance to connect all of the subnets inside the VPC. This is what allows the IP addresses inside of one VPC to overlap with any other VPC, or even with the physical network. A VPC can expand to include any cluster managed by the same Prism Central, but generally a VPC should exist only within a single AHV cluster, or within clusters in the same availability zone. More on that when we get to External Networks.

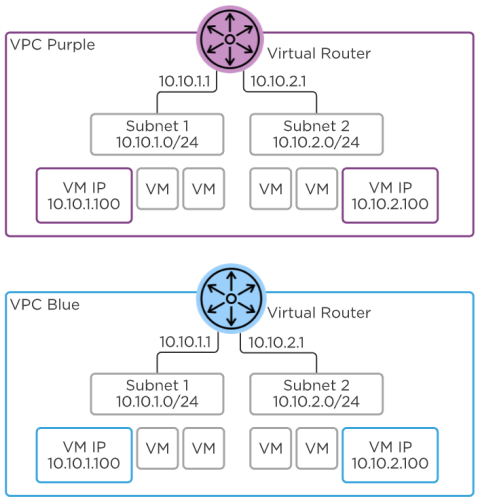

Flow Virtual Networking - VPC

Overlay Subnets

Each VPC can have one or more subnets and they're all connected to the same VPC virtual router. Behind the scenes, a VPC leverages Geneve encapsulation to tunnel traffic between AHV hosts as needed. This means the subnets inside the VPC don't need to be created or even present on the top-of-rack switches for VMs on different hosts to communicate. When two VMs in a VPC on two different hosts send traffic to each other, packets are encapsulated in Geneve on the first host and sent to the other host where they're decapsulated and sent to the destination VM.

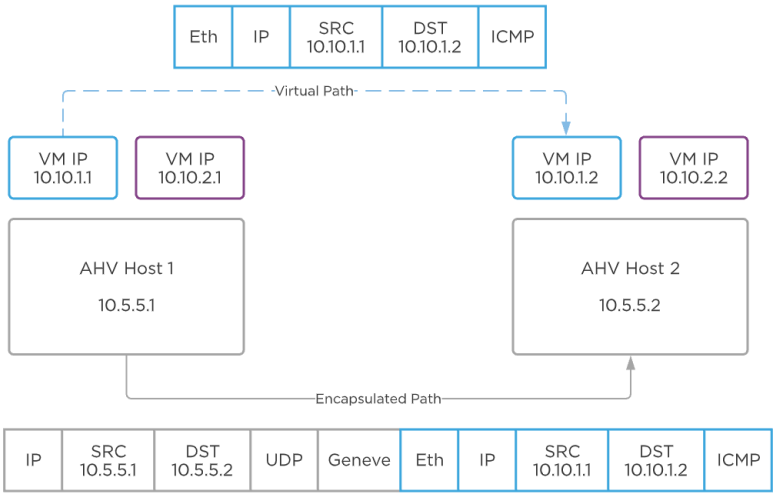

Flow Virtual Networking - Geneve Encapsulation

When you select a NIC for a VM, you can place that NIC into an overlay subnet, or a traditional VLAN backed subnet. When you choose an overlay subnet, that is also choosing the VPC.

Pro tip

Each VM can be placed inside only a single VPC. You can't connect a VM to both a VPC and a VLAN at the same time, or to two different VPCs at the same time.

Routes

Every VPC contains a single virtual router and there are a few types of routes:

- External Networks
- Direct Connections
- Remote Connections

External networks should be the default destination of the 0.0.0.0/0 network prefix for the whole VPC. You can choose an alternate network prefix route for each external network in use. For completely isolated VPCs, you may choose not to set a default route.

Directly connected routes are created for each subnet inside the VPC. Flow Virtual Networking assigns the first IP address of each subnet as the default gateway for that subnet. The default gateway and network prefix are determined by the subnet configuration and cannot be altered directly. Traffic between two VMs on the same host and same VPC, but in two different subnets, will be routed locally in that host.

Remote connections such as VPN connections and External Networks can be set as the next hop destination for a network prefix.

Pro tip

Each network prefix inside a VPC routing table must be unique. Don't program two different next hop destinations with the same destination prefix.

Policies

The virtual router acts as a control point for traffic inside a VPC. You can apply simple stateless policies here, and any traffic that flows through the router will be evaluated by the policies. Traffic from one VM to another inside the same subnet won't go through a policy.

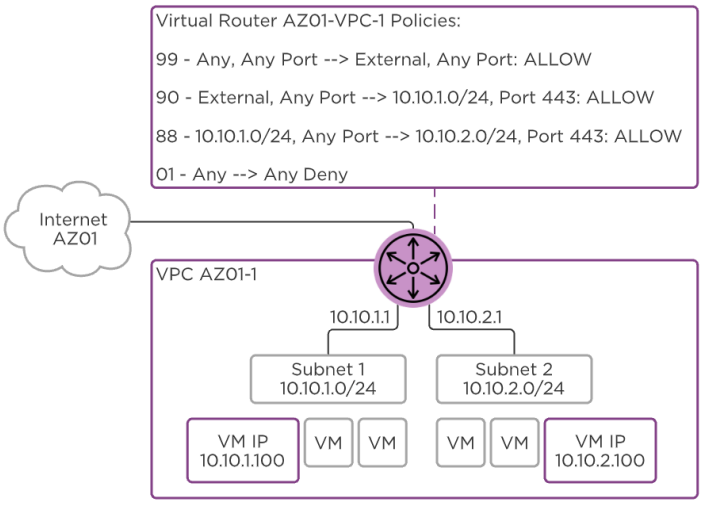

Flow Virtual Networking - Policies

Inside a VPC, policies are evaluated in priority order from highest (1,000) to lowest (10).

You can match traffic based on the following values:

- Source or Destination IPv4 Address
  - Any
  - External to the VPC (Any traffic entering the VPC)
  - Custom IP Prefix
- Protocol
  - Any
  - Protocol Number
  - TCP
  - UDP
  - ICMP
- Source or Destination Port Number

Once traffic is matched a policy can take the following actions:

- Permit
- Deny
- Reroute
  - Redirect traffic to another /32 IPv4 address in another subnet

The reroute policy is incredibly helpful to take action like routing all inbound traffic through a load balancer or firewall VM running inside another subnet in the VPC. This has the added value of requiring only a single Network Function VM (NFVM) for all traffic inside the VPC, rather than a traditional service chain that requires an NFVM per AHV host.

Pro tip

Stateless policies require separate rules defined in both the forward and reverse direction if a Permit rule is overriding a Drop rule. Otherwise, return traffic would be deined by the Drop rule. Use similar priorities to group these matching forward and reverse entries.

External Networks

An External Network is the primary way traffic enters and exits a VPC. External Networks are created in Prism Central and exist on only a single Prism Element cluster. This network defines the VLAN, the default gateway, the IP address pool, and the NAT type for all the VPCs using it. One External Network can be used by many VPCs.

There are two types of External Networks:

- NAT
- Routed (NoNAT)

Starting in PC.2022.1 and AOS 6.1, you can select a maximum of two External Networks for a VPC. A VPC can have at most one NAT External Network and at most one Routed (NoNAT) External Network.

NAT

A NAT (Network Address Translation) External Network hides the IP addresses of VMs in the VPC behind either a Floating IP or the VPC SNAT (Source NAT) address. Each VPC has an SNAT IP address selected randomly from the External Network IP pool, and traffic exiting the VPC is rewritten with this source address.

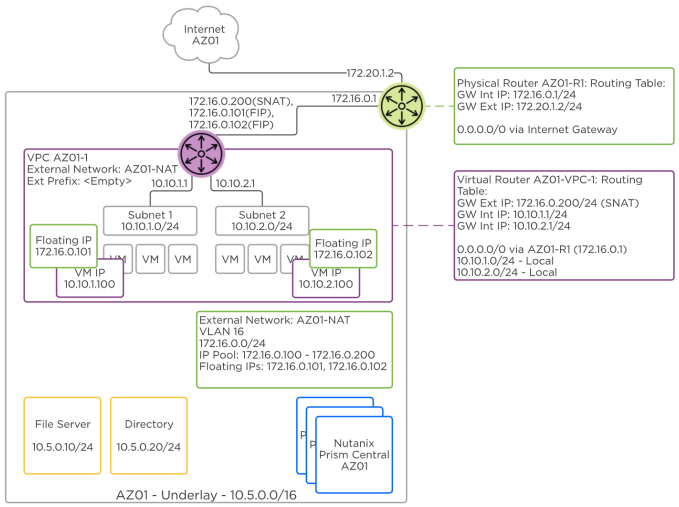

Flow Virtual Networking - NAT External Network

Floating IP addresses are also selected from the External Network IP pool and are assigned to VMs in a VPC to allow ingress traffic. When a Floating IP is assigned to a VM, the egress traffic is rewritten with the Floating IP instead of the VPC SNAT IP. This is useful for advertising public services outside the VPC without revealing the private IP address of the VM.

Routed

Routed, or NoNAT External Networks allow the IP address space of the physical network to be shared inside the VPC through routing. Instead of a VPC SNAT IP address, the VPC router IP is selected randomly from the External Network pool. Share this VPC router IP with the physical network team so they can set this virtual router IP as the next hop for all of the subnets provisioned inside the VPC.

For example, a VPC with an External Network of 10.5.0.0/24 may be assigned a virtual router IP of 10.5.0.200. If the subnets inside the VPC are created in the 10.10.0.0/16 network, the physical network team will create a route to 10.10.0.0/16 via 10.5.0.200. The 10.10.0.0/16 network becomes the externally routable prefix for the VPC.

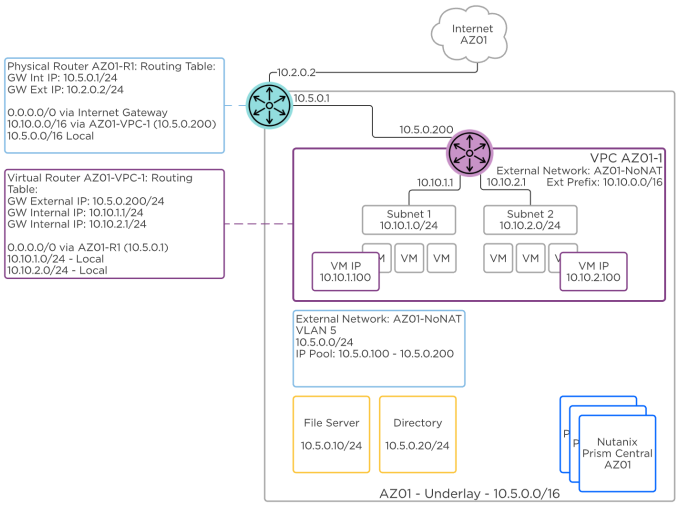

Flow Virtual Networking - Routed External Network

Network Gateways

A Network Gateway acts as a connector between subnets. These subnets can be of many different types and in different locations.

- Subnet Types
  - ESXi VLAN
  - AHV VLAN

    - VPC Overlay
    - Physical Network VLAN

- Subnet Locations
    - On-premises VLANs
    - On-premises VPCs
    - Cloud VPCs

The Network Gateway has several methods of connecting subnets.

- Layer 3 VPN
    - Network Gateway to Network Gateway
    - Network Gateway to Physical Firewall or VPN

- Layer 2 VXLAN VTEP
    - Network Gateway to Network Gateway
    - Network Gateway to Physical Router or Switch VTEP

- Layer 2 VXLAN VTEP over VPN
    - Network Gateway to Network Gateway

Layer 3 VPN

In the Layer 3 VPN connection type, two subnets with two separate network prefixes are connected. For example, local subnet 10.10.1.0/24 can be connected to remote subnet 10.10.2.0/24.

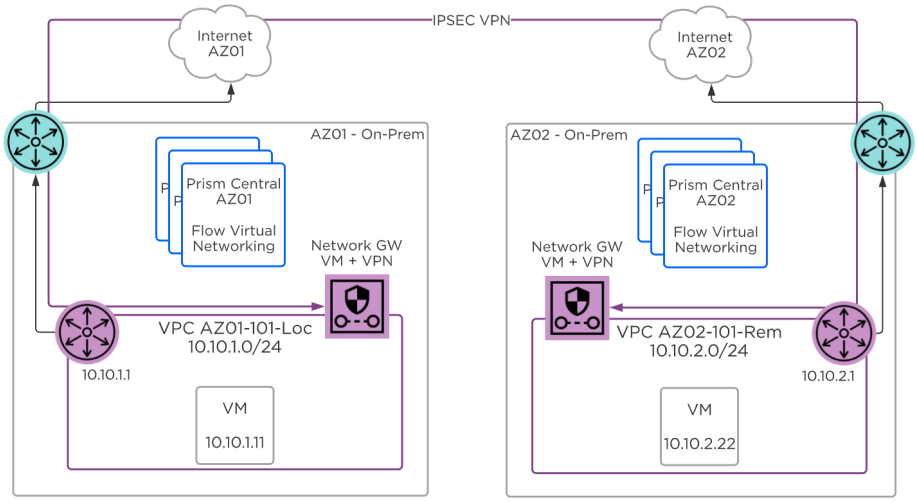

Flow Virtual Networking - Layer 3 VPN

When using two Network Gateways, each Network Gateway is assigned an external IP address from the DHCP pool, and they must be able to communicate over these addresses.

You can also connect the Network Gateway VM to a remote physical firewall or VPN appliance or VM. The local Network Gateway must still be able to communicate with the remote physical or virtual appliance.

Traffic from each subnet to the remote subnet is directed over the created VPN connection using IP routing inside the VPC. Routes can be either static or shared between Network Gateways using BGP. Traffic inside this VPN tunnel is encrypted with IPSEC.

Pro tip

Use VPN with IPSEC encryption when traffic between subnets will go over a public link such as the Internet.

Layer 2 VXLAN VTEP

In the Layer 2 VXLAN VTEP case, two subnets that share the same network prefixes are connected. For example, local subnet 10.10.1.0/24 is connected to a remote subnet that also uses 10.10.1.0/24.

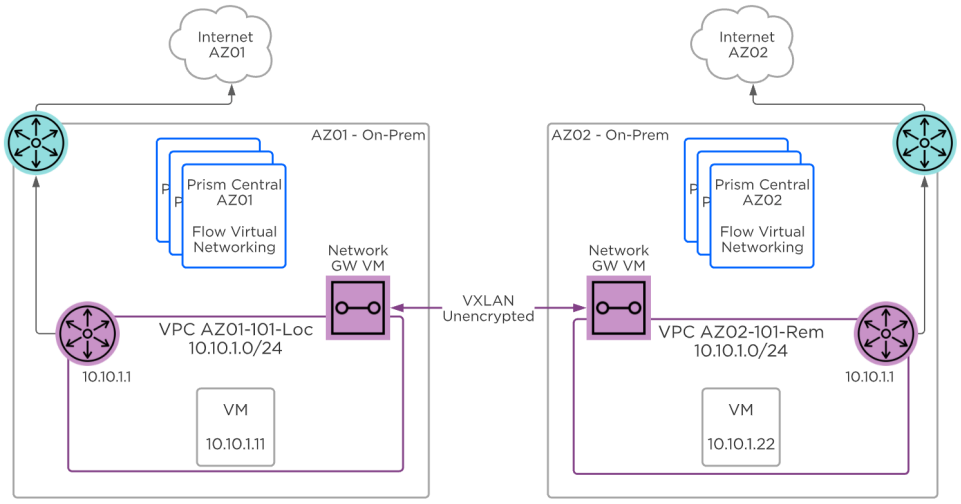

Flow Virtual Networking - Layer 2 VXLAN VTEP

When using two Network Gateway VMs, each Network Gateway is assigned an external IP address, and the two Network Gateway VMs must be able to communicate over these addresses.

The local Network Gateway can also connect to a remote physical VXLAN VTEP terminating device such as a physical switch. The physical device can be any standard VXLAN device from popular vendors such as, but not limited to, Cisco, Arista, and Juniper. Just enter the remote physical device IP address in the VXLAN VTEP connection.

Traffic from the local subnet to the remote subnet is exchanged via layer 2 switching and encapsulated in unencrypted VXLAN. Each Network Gateway maintains a source MAC address table and can forward unicast or flooded packets to the remote subnet.

Pro tip

Use VXLAN VTEP only when traffic is traversing private or secured links, because this traffic is not encrypted.

Layer 2 VXLAN VTEP over VPN

For extra security, the VXLAN connection can be tunneled through an existing VPN connection to add encryption. In this case, the Network Gateway VM provides the VXLAN and the VPN connections, so a Network Gateway VM is required in both the local and remote subnet.

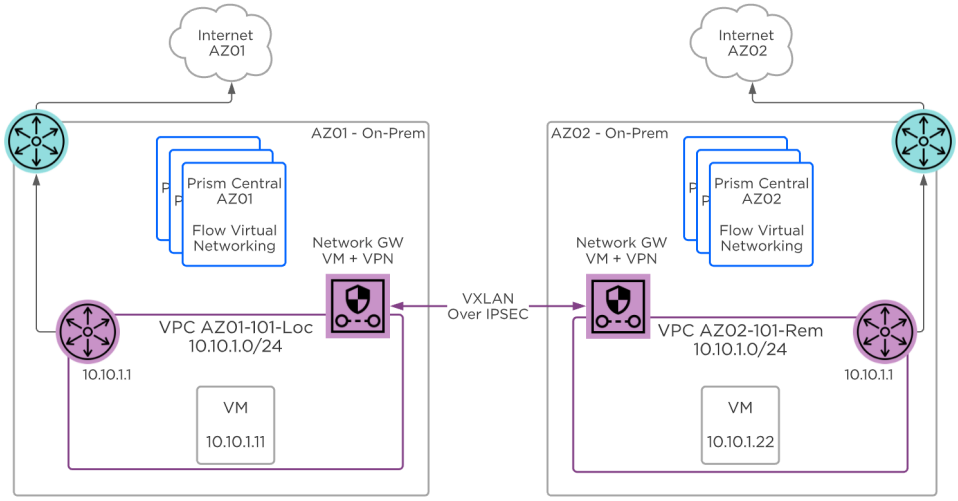

Flow Virtual Networking - Layer 2 VXLAN VTEP over VPN