Microsoft

Azure Networking Overview

June 2020

Mike Wedderburn-Clarke

Senior Cloud Solution Architect
Financial Services


miwedder@microsoft.com


https://twitter.com/MikeWeddClarke


https://www.linkedin.com/in/mikewedderburnclarke

Microsoft global network

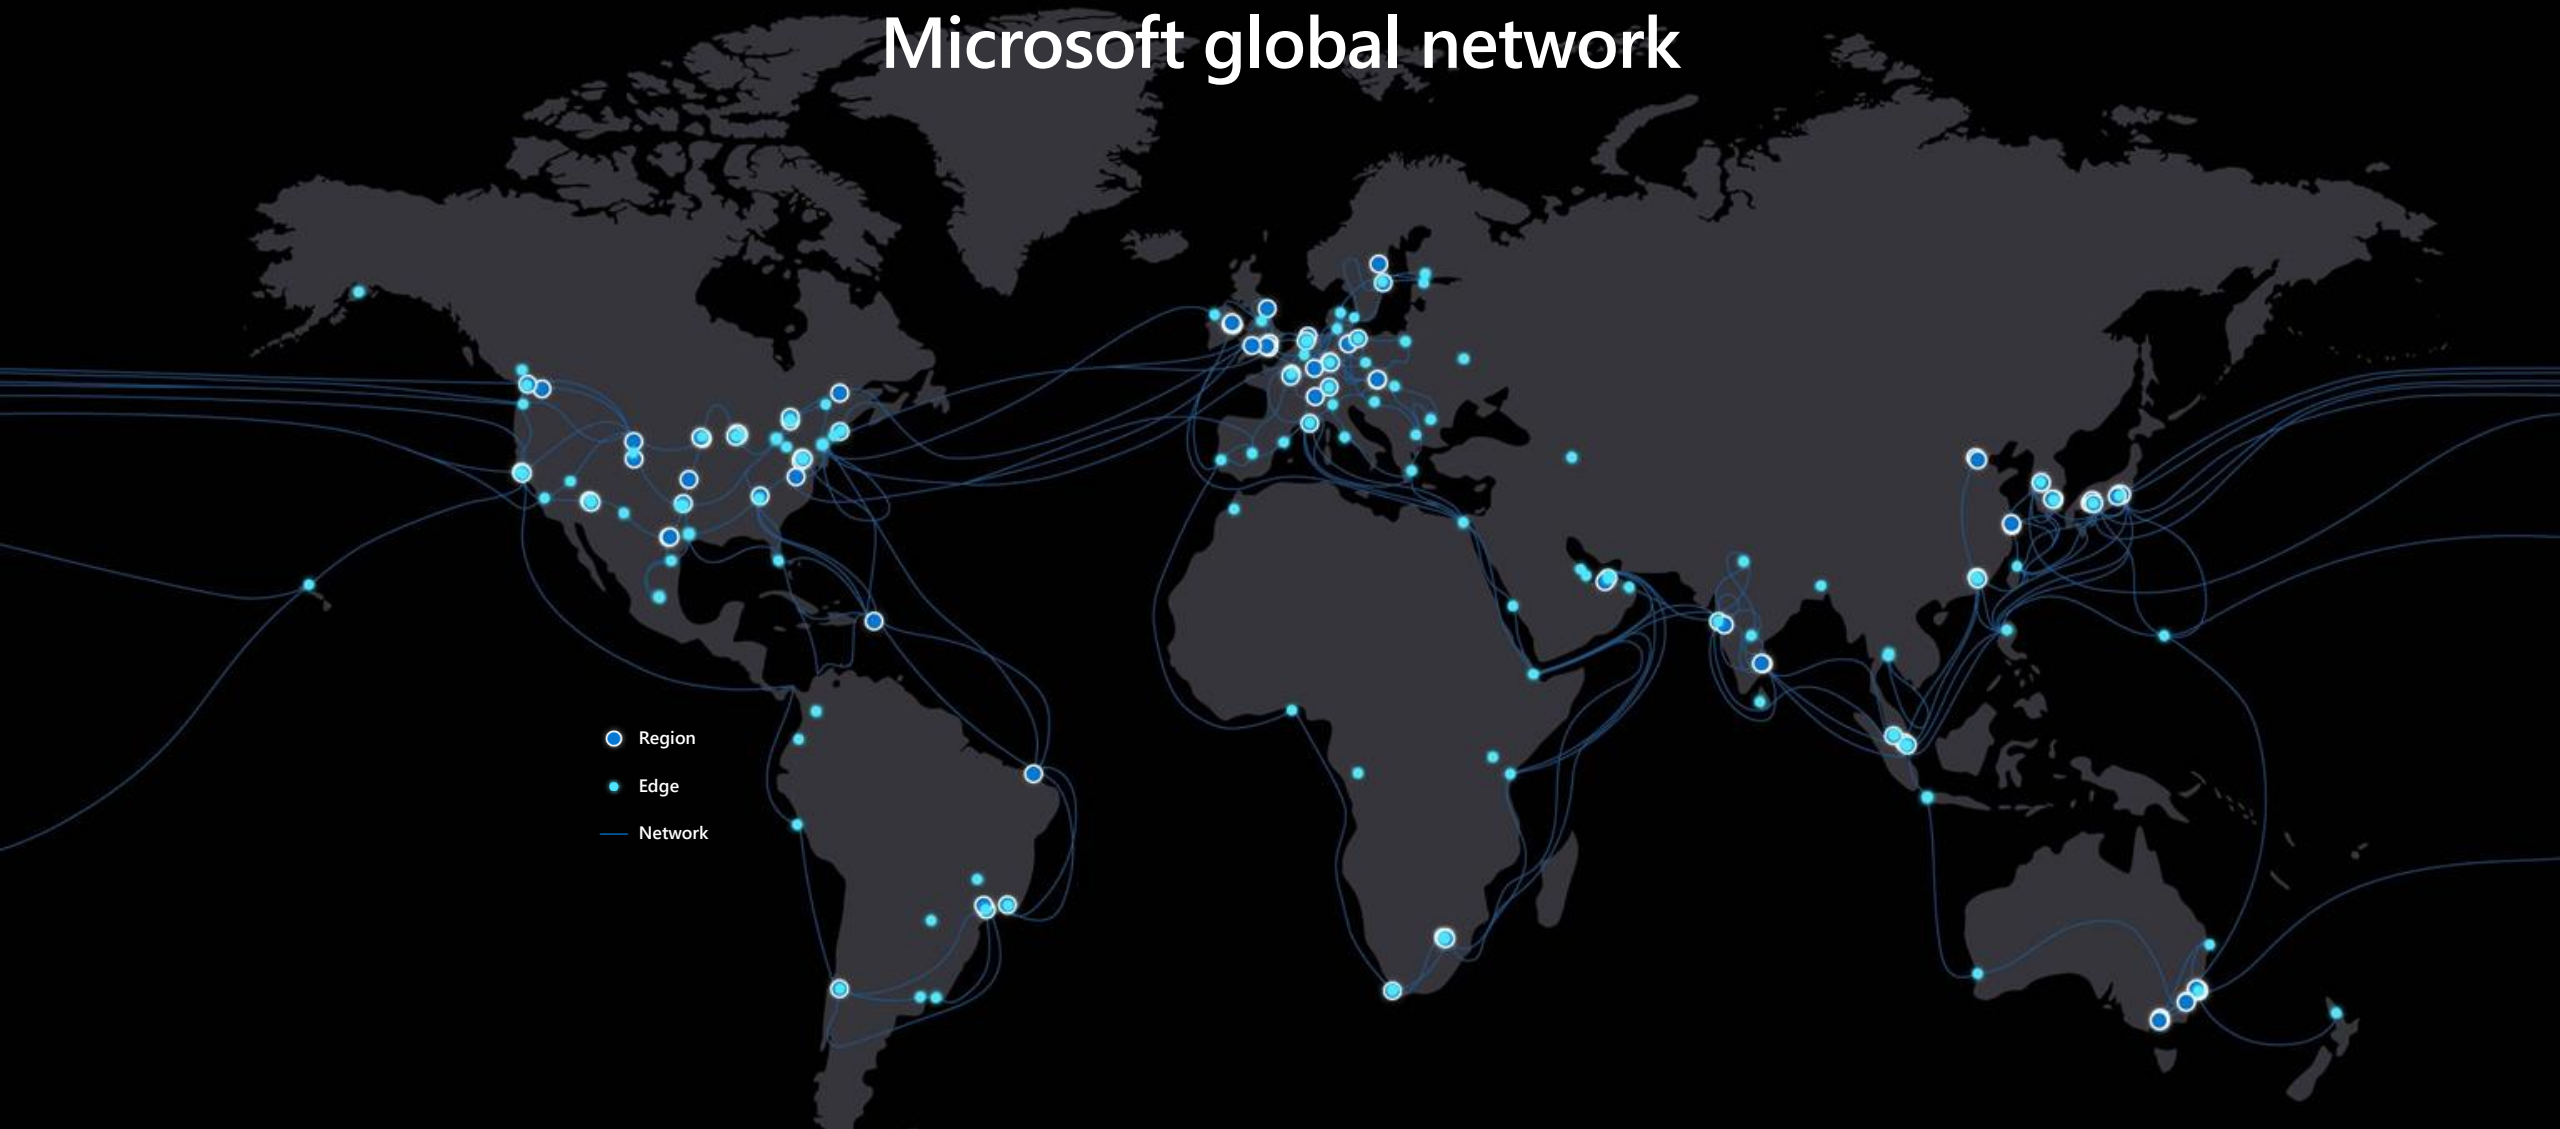

60 Azure regions | 130k+ miles of fiber + subsea cables | 160+ edge sites | 500+ network partners | 20k+ peering connections

"IP traffic stays entirely within our global network and never enters the public Internet"

Classic network vs. Hyper-scale network architecture

Large L2 Domains

HW-based service modules

Simple Tree Design

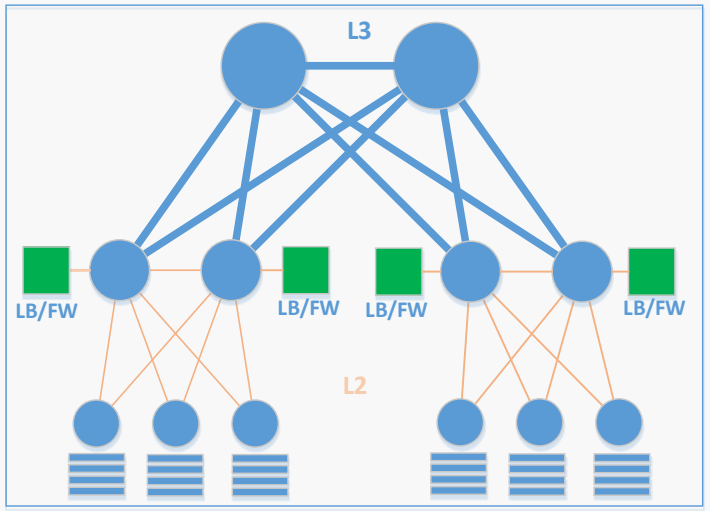

L3 at all layers

Services in software

Clos-based design

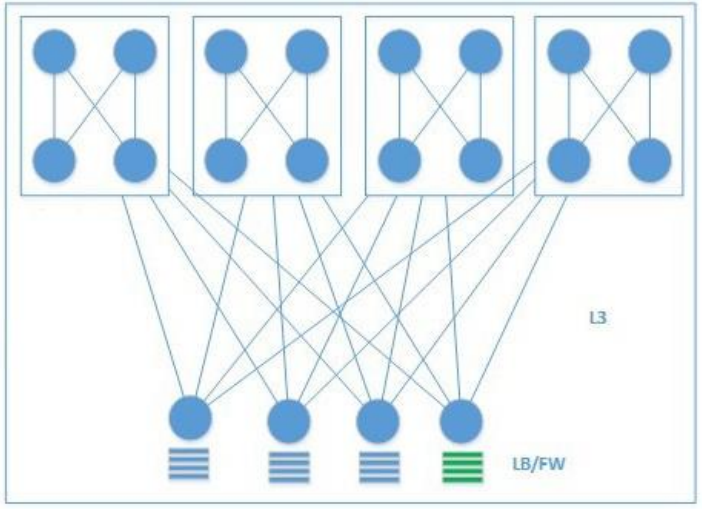

Low due to diversity and manual provisioning process	**Agility**	Automated network provisioning, integrated process
Low due to complex hardware and lack of automated operations	**Efficiency**	Simplify requirements, optimize design, and unify infrastructure
Low due to high complexity and human error	**Availability**	Resilient design, automated monitoring and remediation, minimum human involvement

Security layers

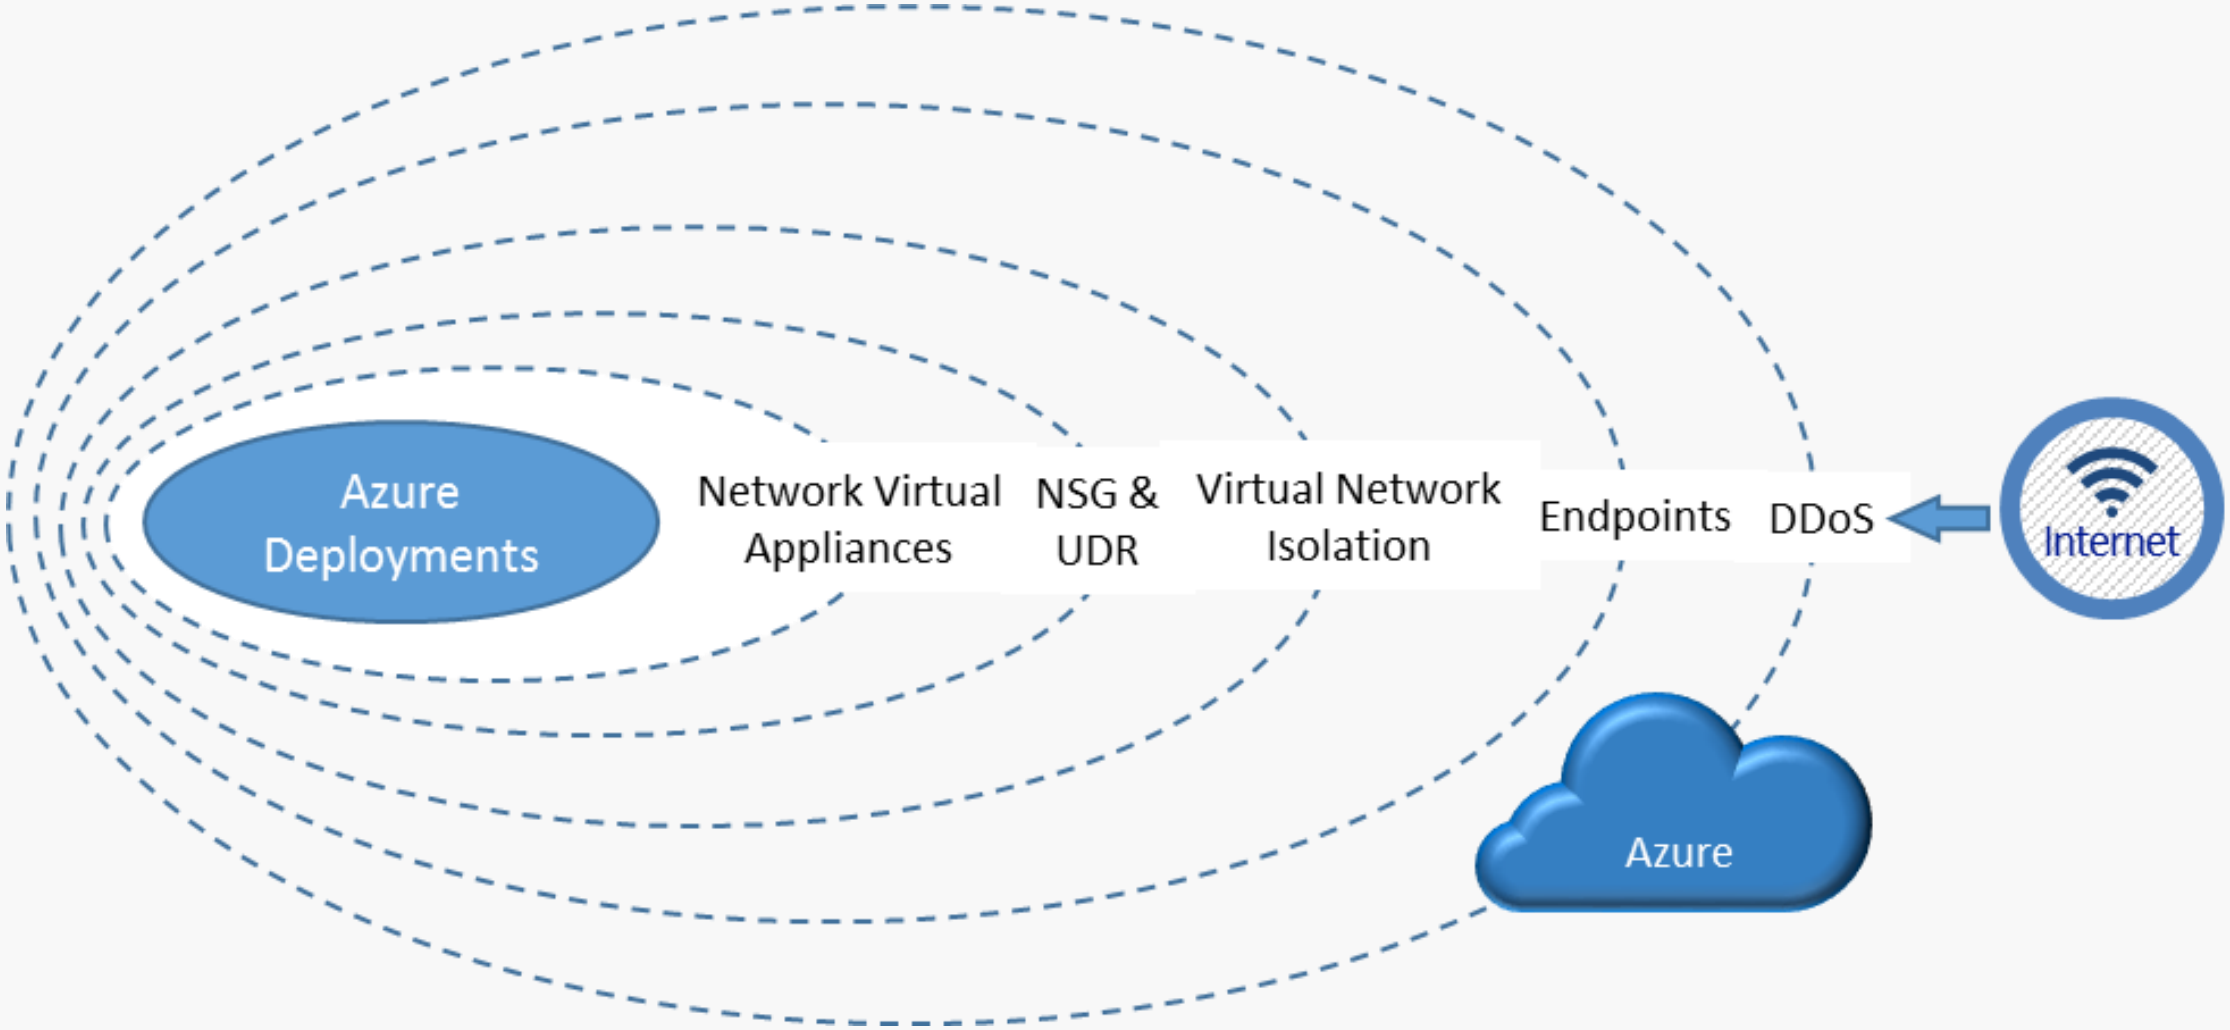

Virtual Network

Isolated, logical network that provides connectivity for Azure Virtual Machines

User-defined address space (can be one or more IP ranges)

1. Connectivity for VMs in the same VNET
2. Connectivity to external networks/on-prem DC's
3. Internet connectivity

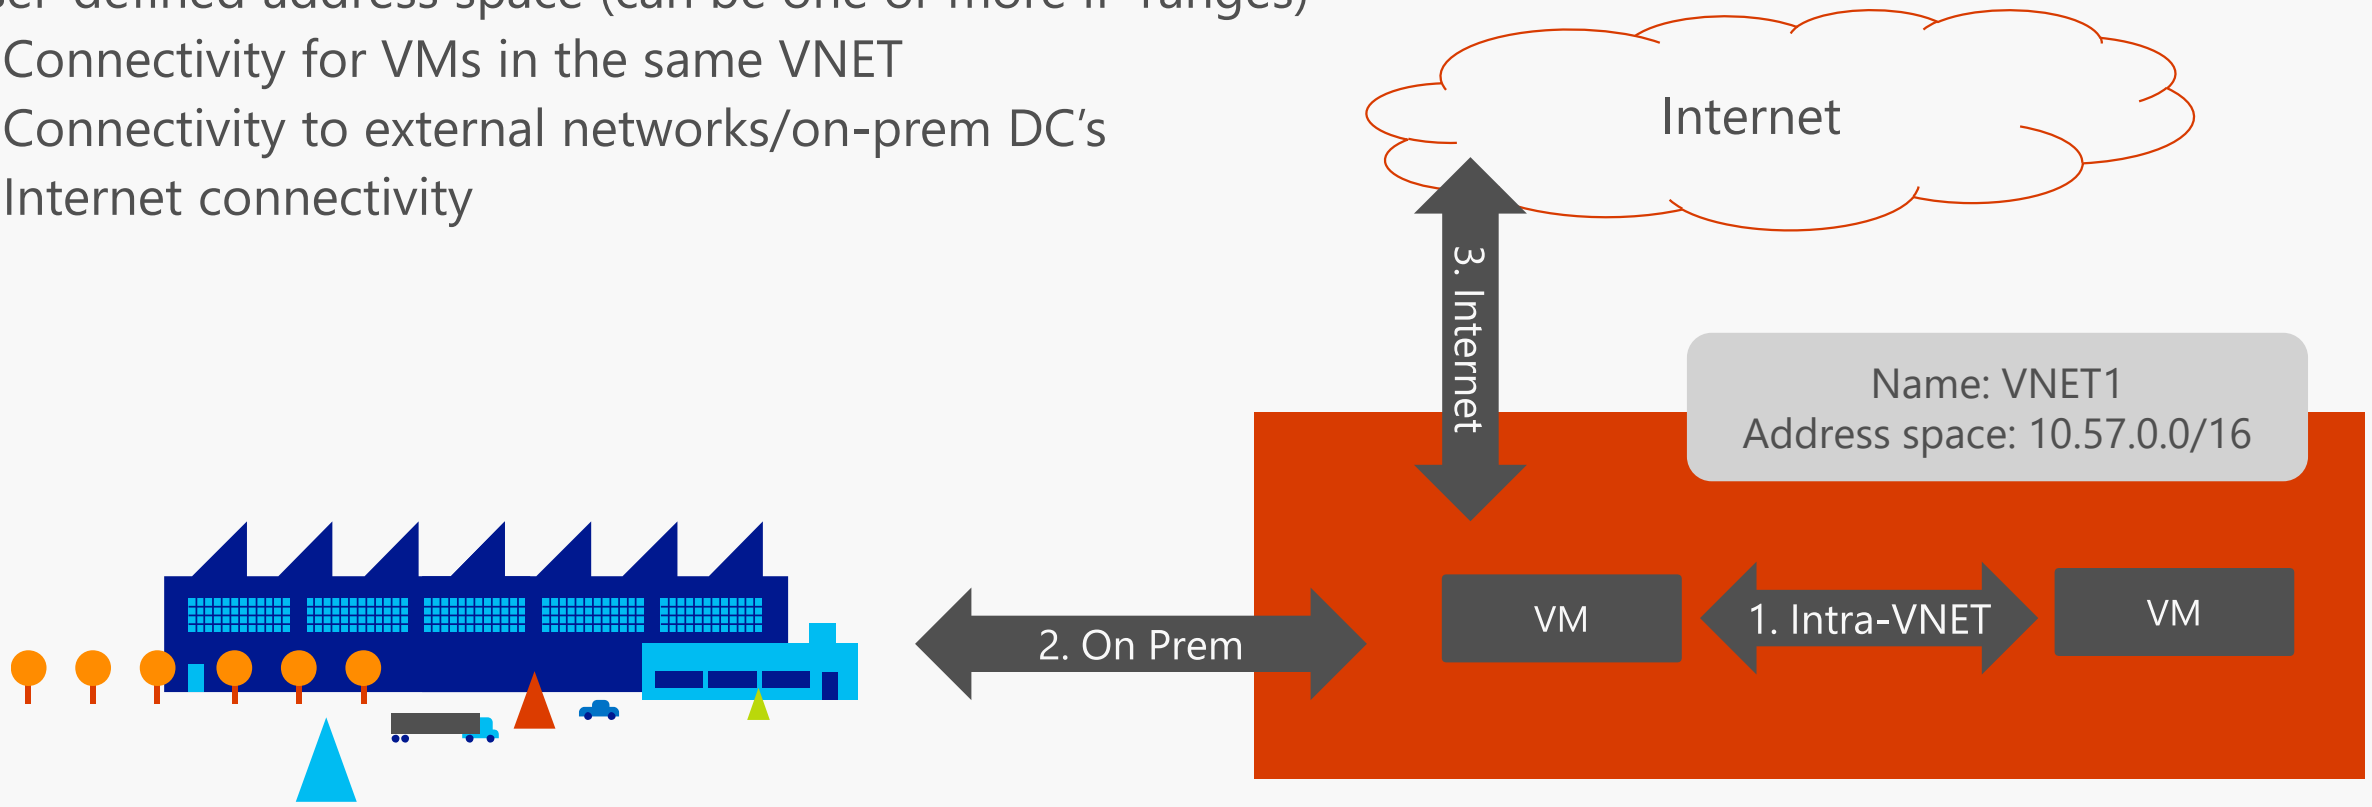

Subnet

IP subnet

Provides full layer-3 semantics and partial layer-2 semantics (DHCP, ARP, no broadcast/multicast)
Subnets can span only one range of contigous IP addresses
VMs can be deployed only to subnets

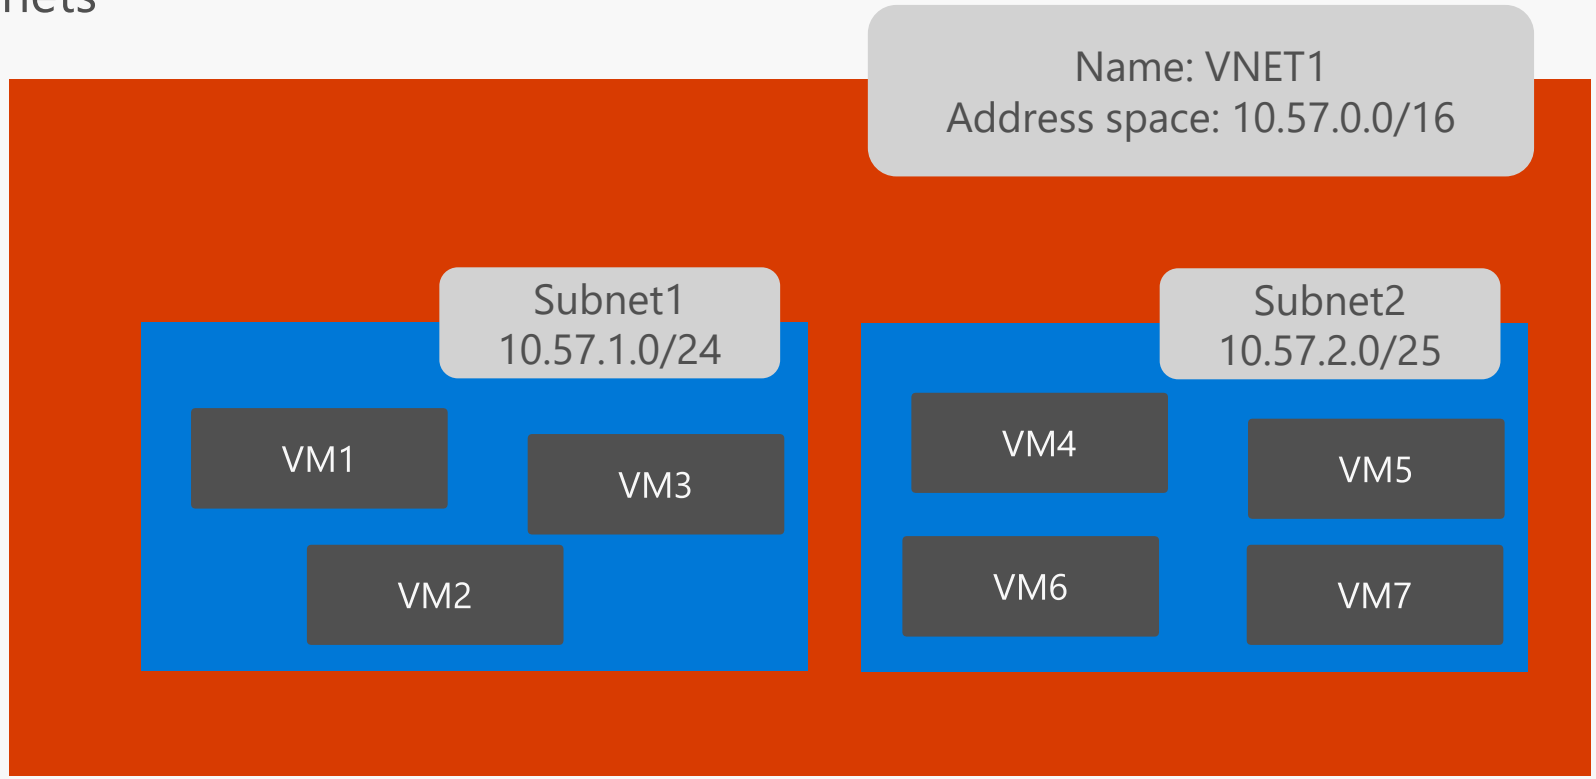

Network Interface

Virtual NIC that connects a VM to a Subnet

One private IP address (private == included in the subnet's IP range)

Private IP address always assigned via Azure DHCP

- Dynamic assignment = DHCP assigns new IP when VM is restarted
- Static assignment =DHCP assigns always the same IP
- IP forwarding = NIC can receive packets with dest IP address different from its private IP
- Multiple NICs
- Multiple IP addr per NIC

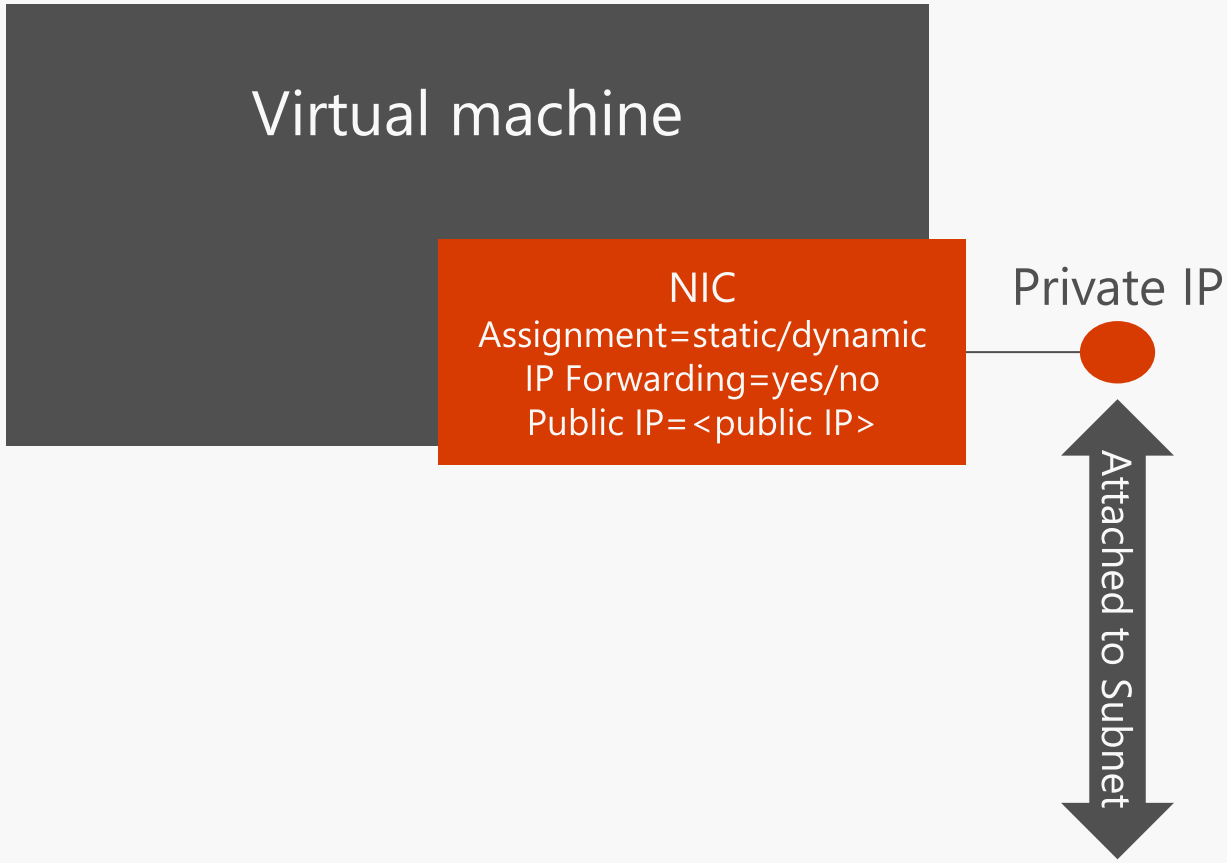

IP addresses come in two types in Azure

Public vs. Private

Public IP Addresses allow Azure resources to communicate with Internet and other Azure public-facing services

- Virtual machines (VM)
- Internet-facing (public) load balancers
- VPN gateways
- Application gateways

Private IP Addresses allows communication between resources in a virtual network, along with those connected through a VPN, without using an Internet-routable IP addresses.

- VMs
- Internal load balancers (ILBs)
- Application gateways

Load Balancers

External vs. Internal

- **External load balancer.** You can use an external load balancer to provide high availability for IaaS VMs and PaaS role instances accessed from the public Internet.
- **Internal load balancer.** You can use an internal load balancer to provide high availability for IaaS VMs and PaaS role instances accessed from other services in your VNet.

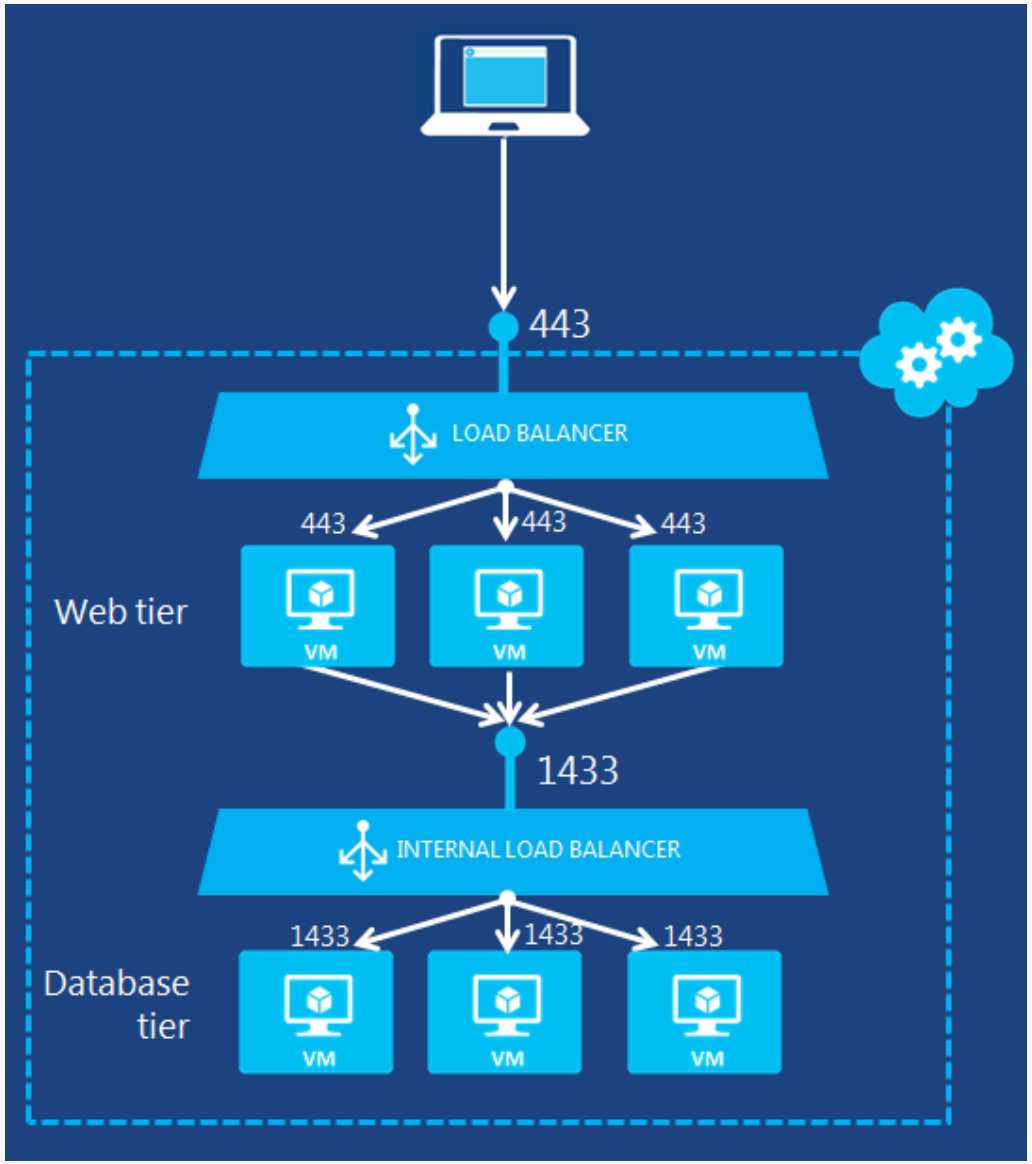

NSG key facts

5-tuple ACL's

Source IP, Destination IP, Source Port, Destination Port, Protocol (TCP, UDP, any)
Actions: allow or deny
Directions: inbound, outbound
Priority: 100-4096 (lower value = higher priority)

Stateful

No need to define rules for «return traffic»

Can be applied to NICs and Subnets (ARM)

Inbound connections: subnet-level NSG evaluated first, NIC-level NSG evaluated next
Outbound connections: NIC-level NSG evaluated first, subnet-level NSG evaluated next
Deny wins

Troubleshooting NSGs

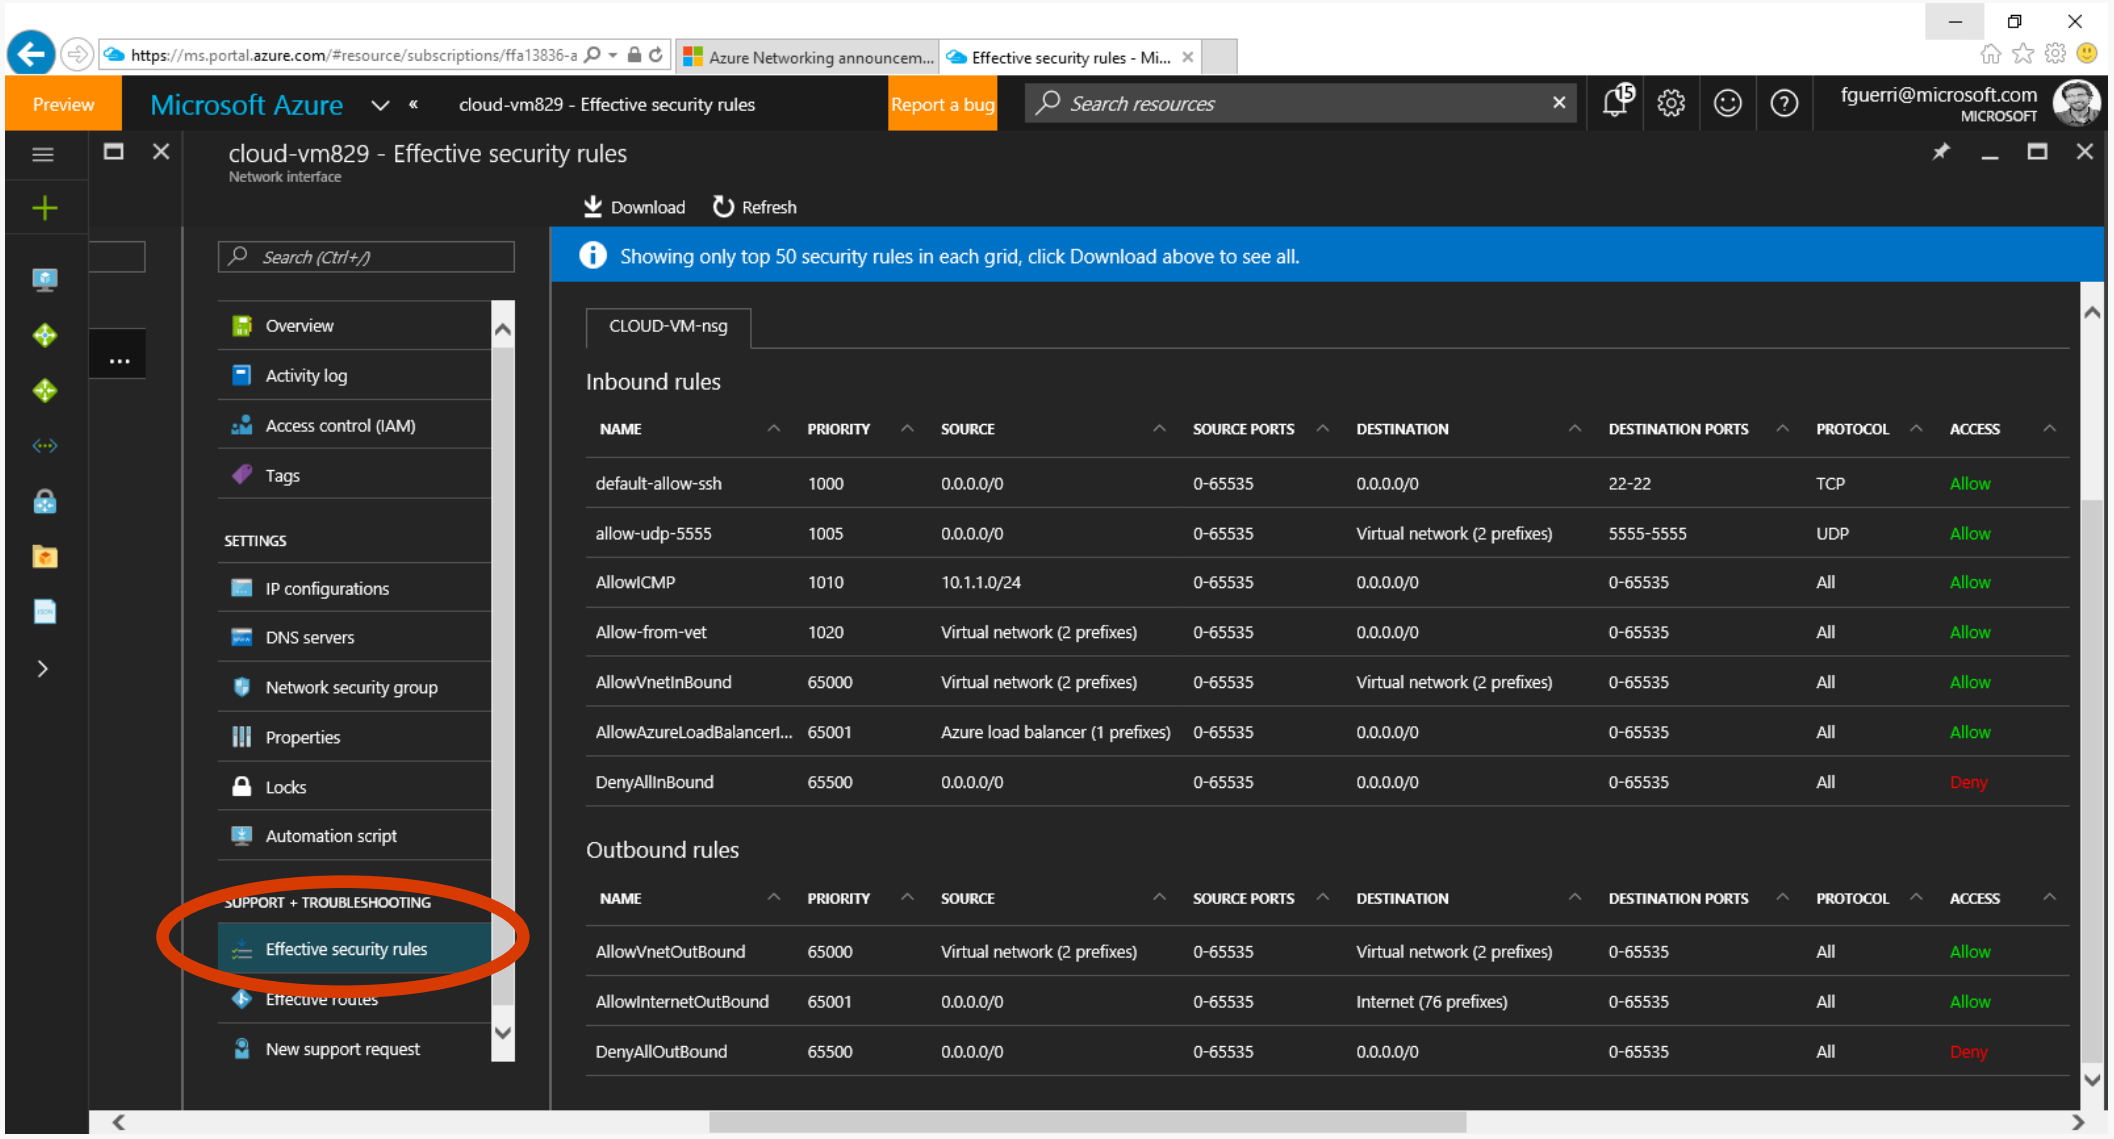

Network Virtual Appliance

A VM in your VNet that runs software

- Example: firewall, WAN optimization..etc
- You can create a route in Azure to route your VNet traffic through a virtual appliance to use its capabilities.
- NSGs provide security on your Vnet (layer 4 ACL on incoming/outgoing packets). NVA will offer a layer 7 security model.

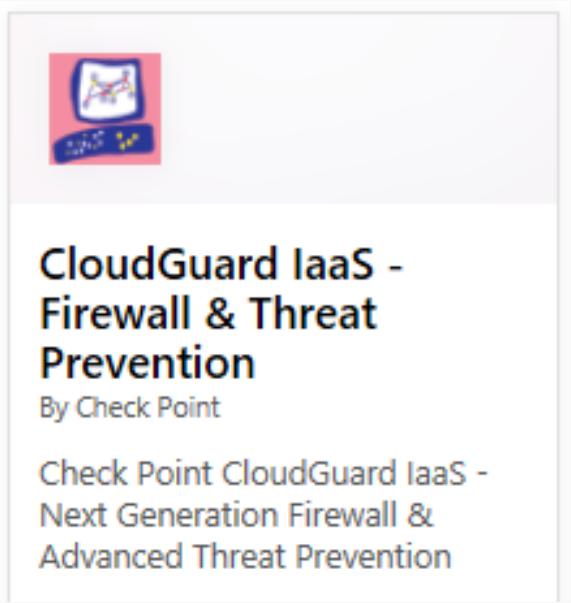

Name Resolution

By default, your VNet uses Azure-provided name resolution to resolve names inside the VNet, and on the public Internet.

If you connect your VNets to your on-premises data centers, you need to provide your own DNS server to resolve names between your networks.

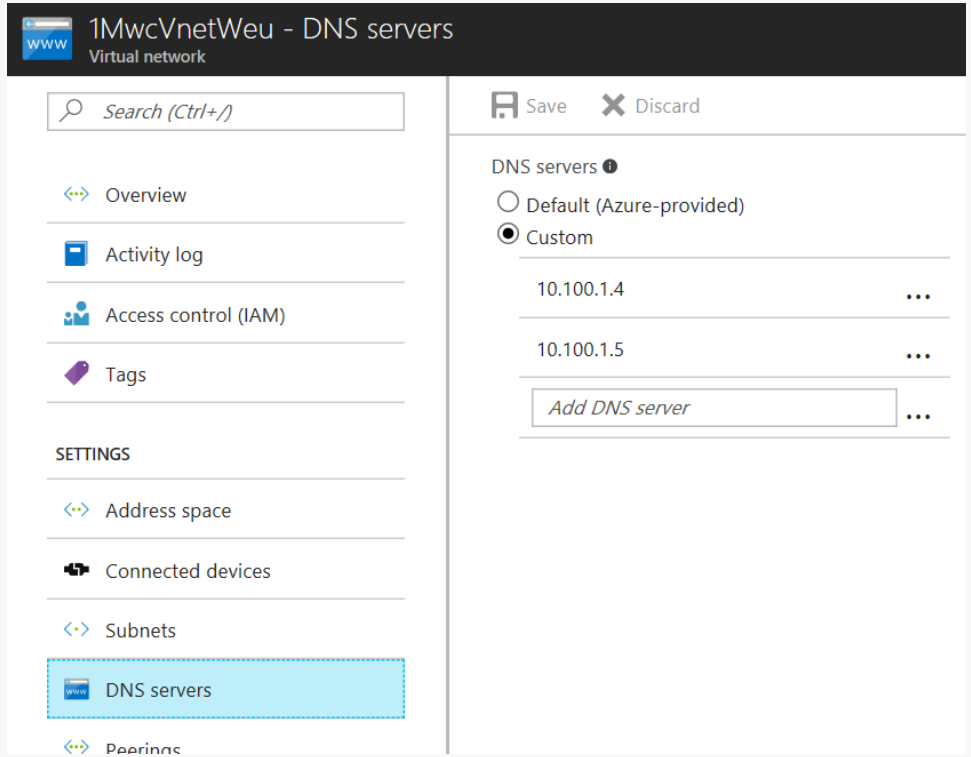

System Route Table

Default rules for routing/switching traffic in Azure VNETs

- Route table: set of rules that define where IP packets must be sent based on their destination IP address
- The default routing behavior for an Azure VNET is defined by the «System Route Table»

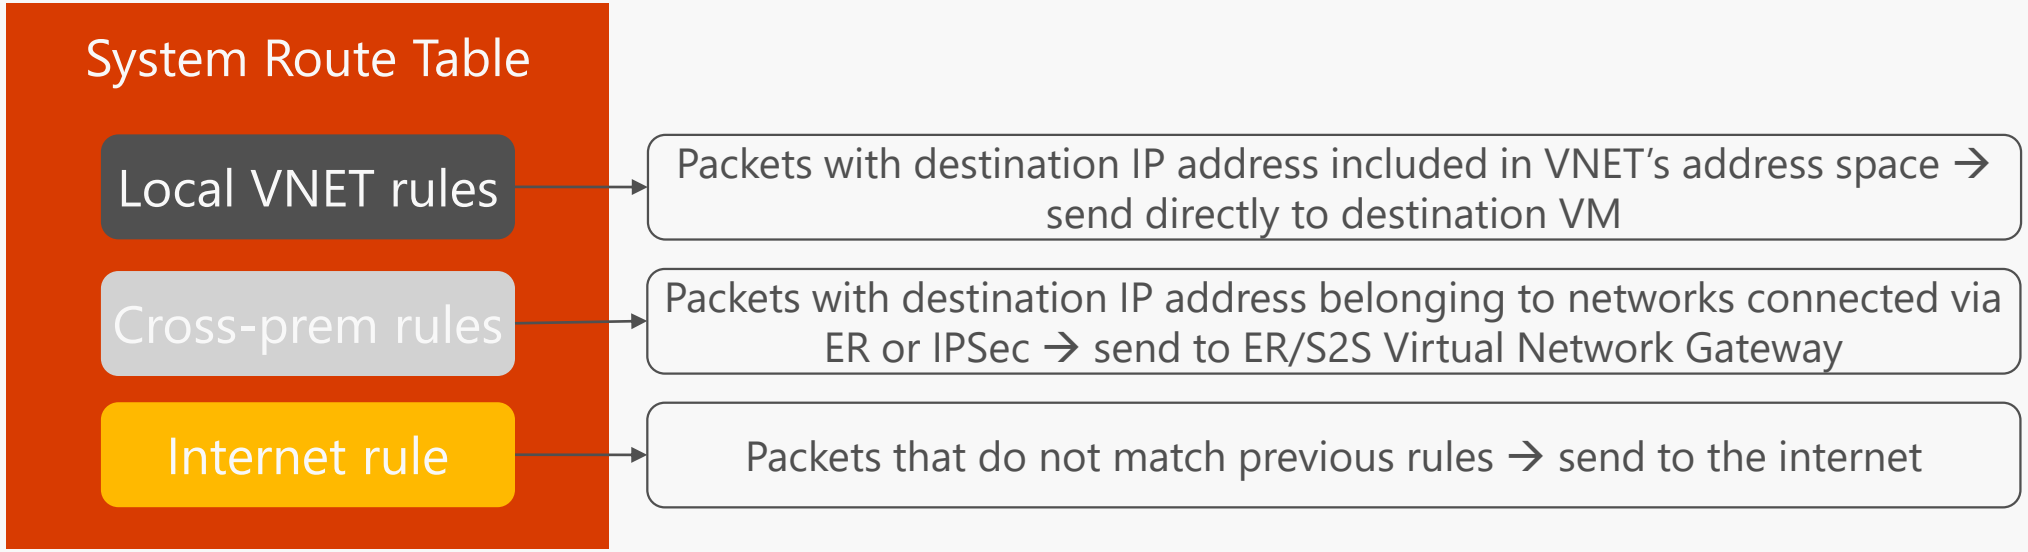

User Defined Routes (UDR's)

Additional routes that modify a VNET's default routing policy

- A custom route table contains one or more UDR's AND the system routes
- UDR's are preferred over system routes with the same prefix length
- Each subnet in a VNET can be assigned a different custom route table
- A custom route table can be assigned to the Gateway Subnet

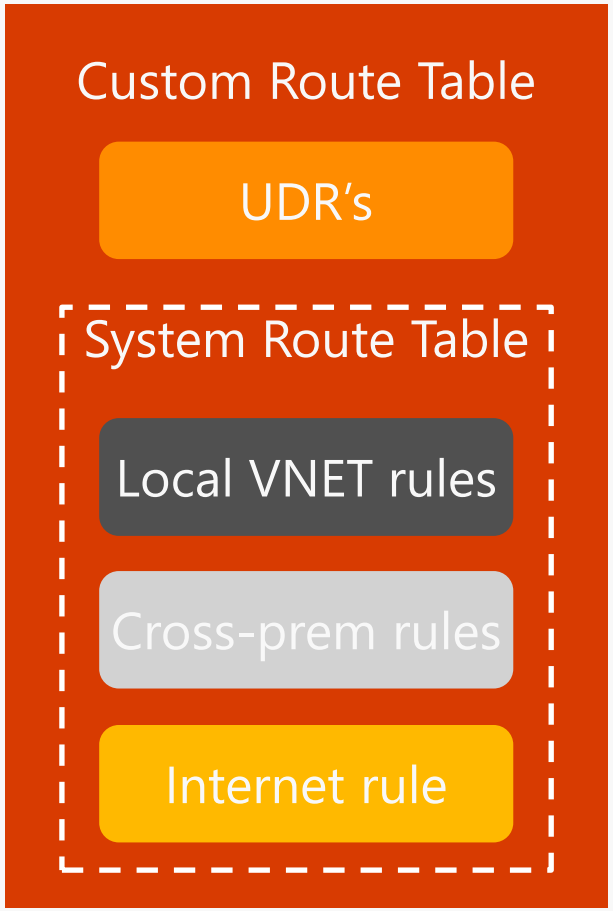

User Defined Routes (UDR's)

Use case 1: Virtual appliances

- According to the system route table, traffic will flow directly from VM1 to VM4

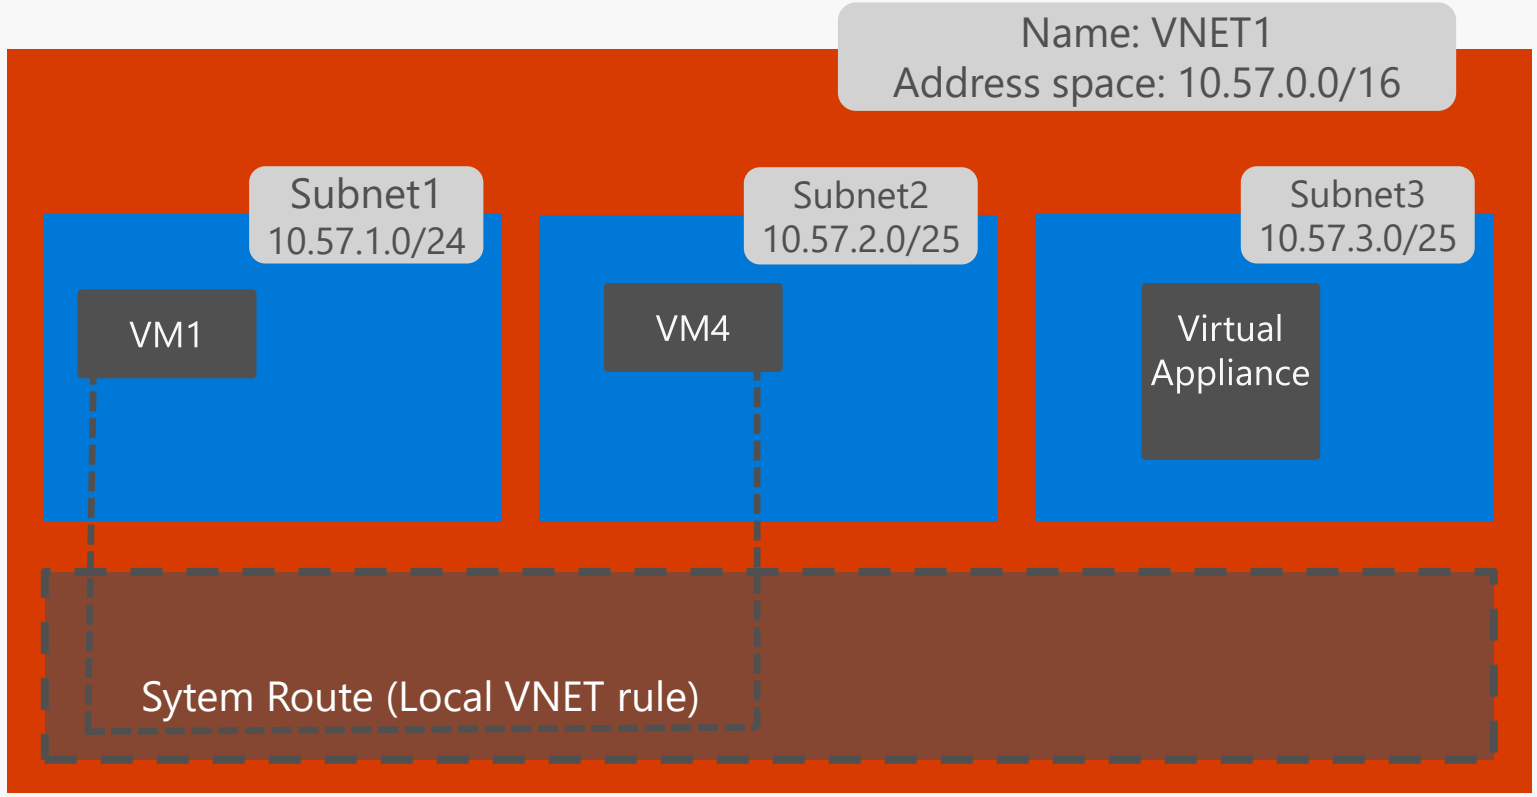

User Defined Routes (UDR's)

Use case 1: Virtual appliances

- According to the system route table, traffic will flow directly from VM1 to VM4
- A UDR can be used to override this behavior and send the traffic through an intermediate hop (e.g. a firewalling VA)
- UDR cannot be overridden by VM local route table

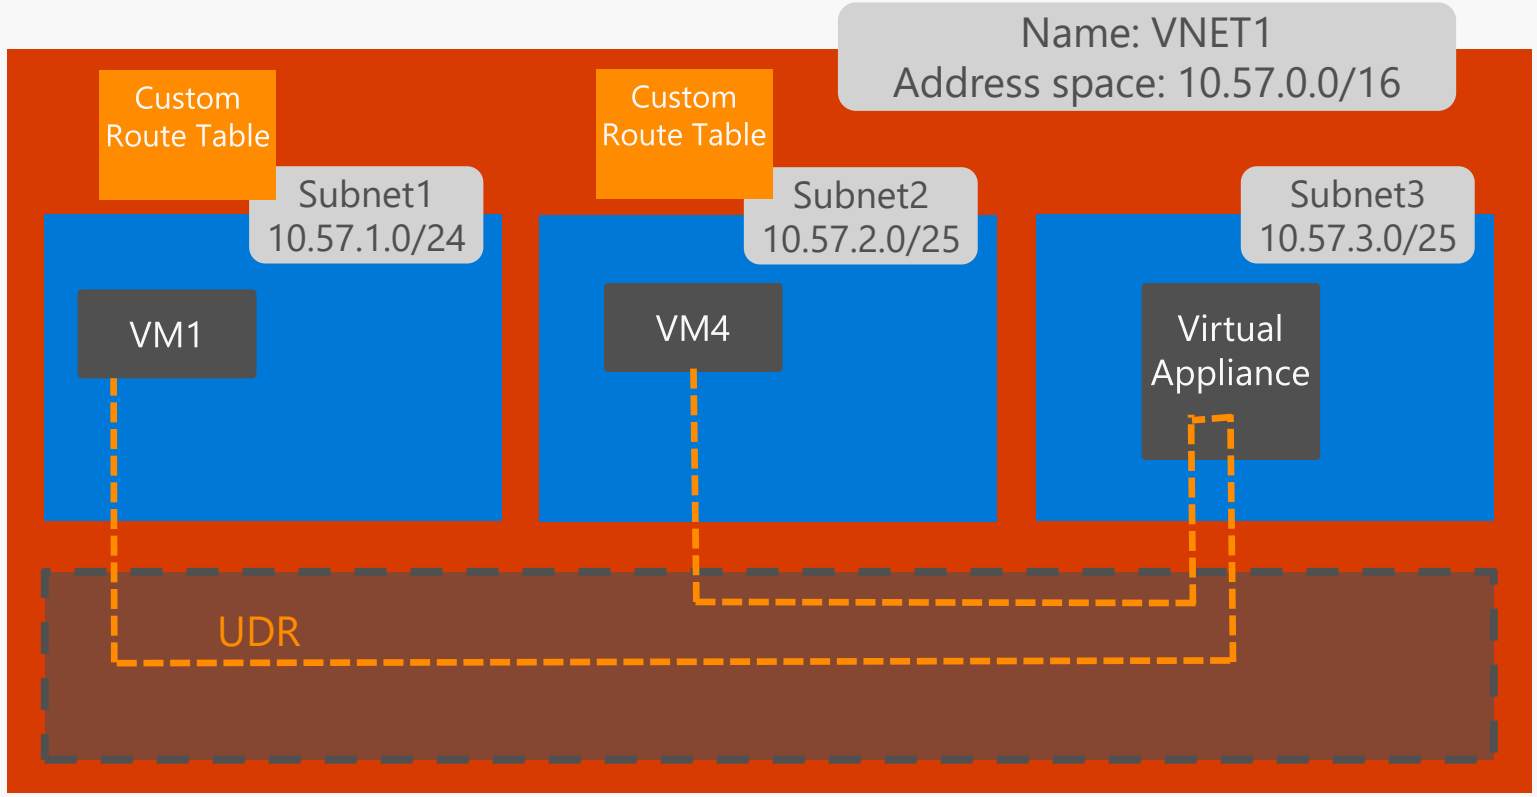

User Defined Routes (UDR's)

Use case 2: Inbound ER or S2S traffic

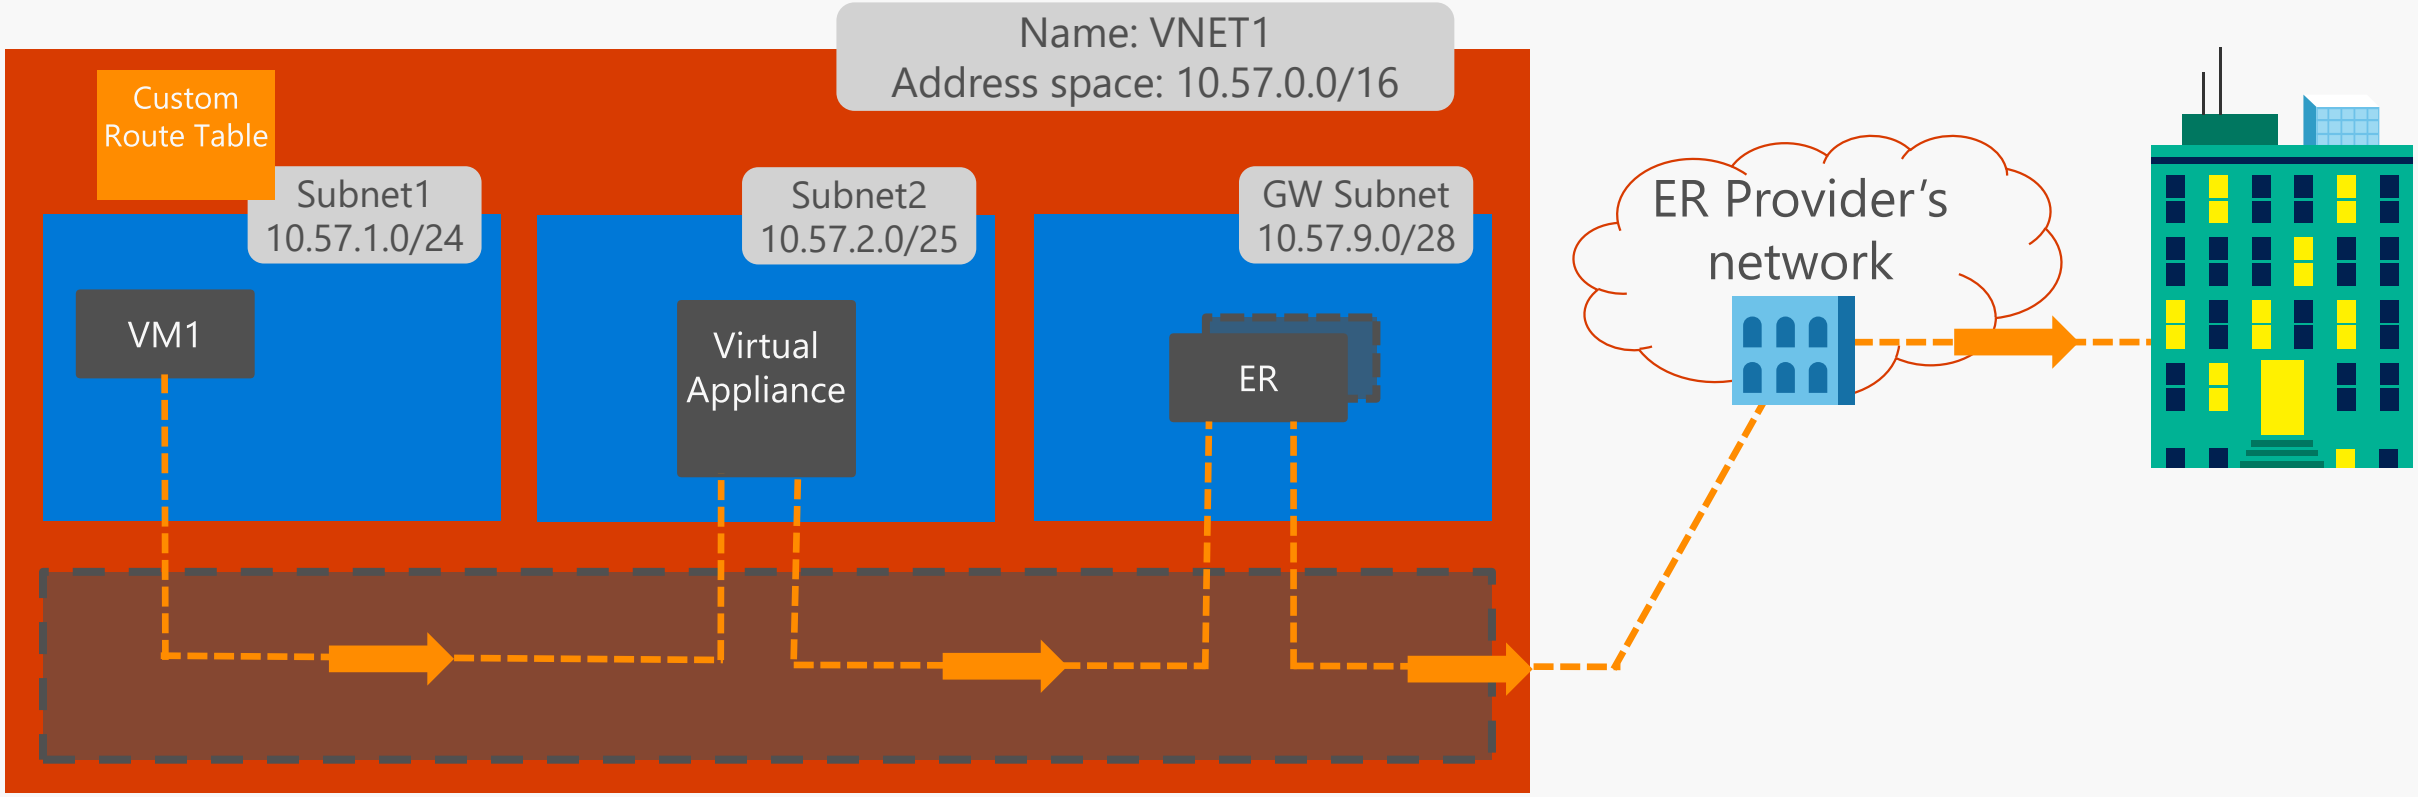

User Defined Routes (UDR's)

Use case 2: Inbound ER or S2S traffic

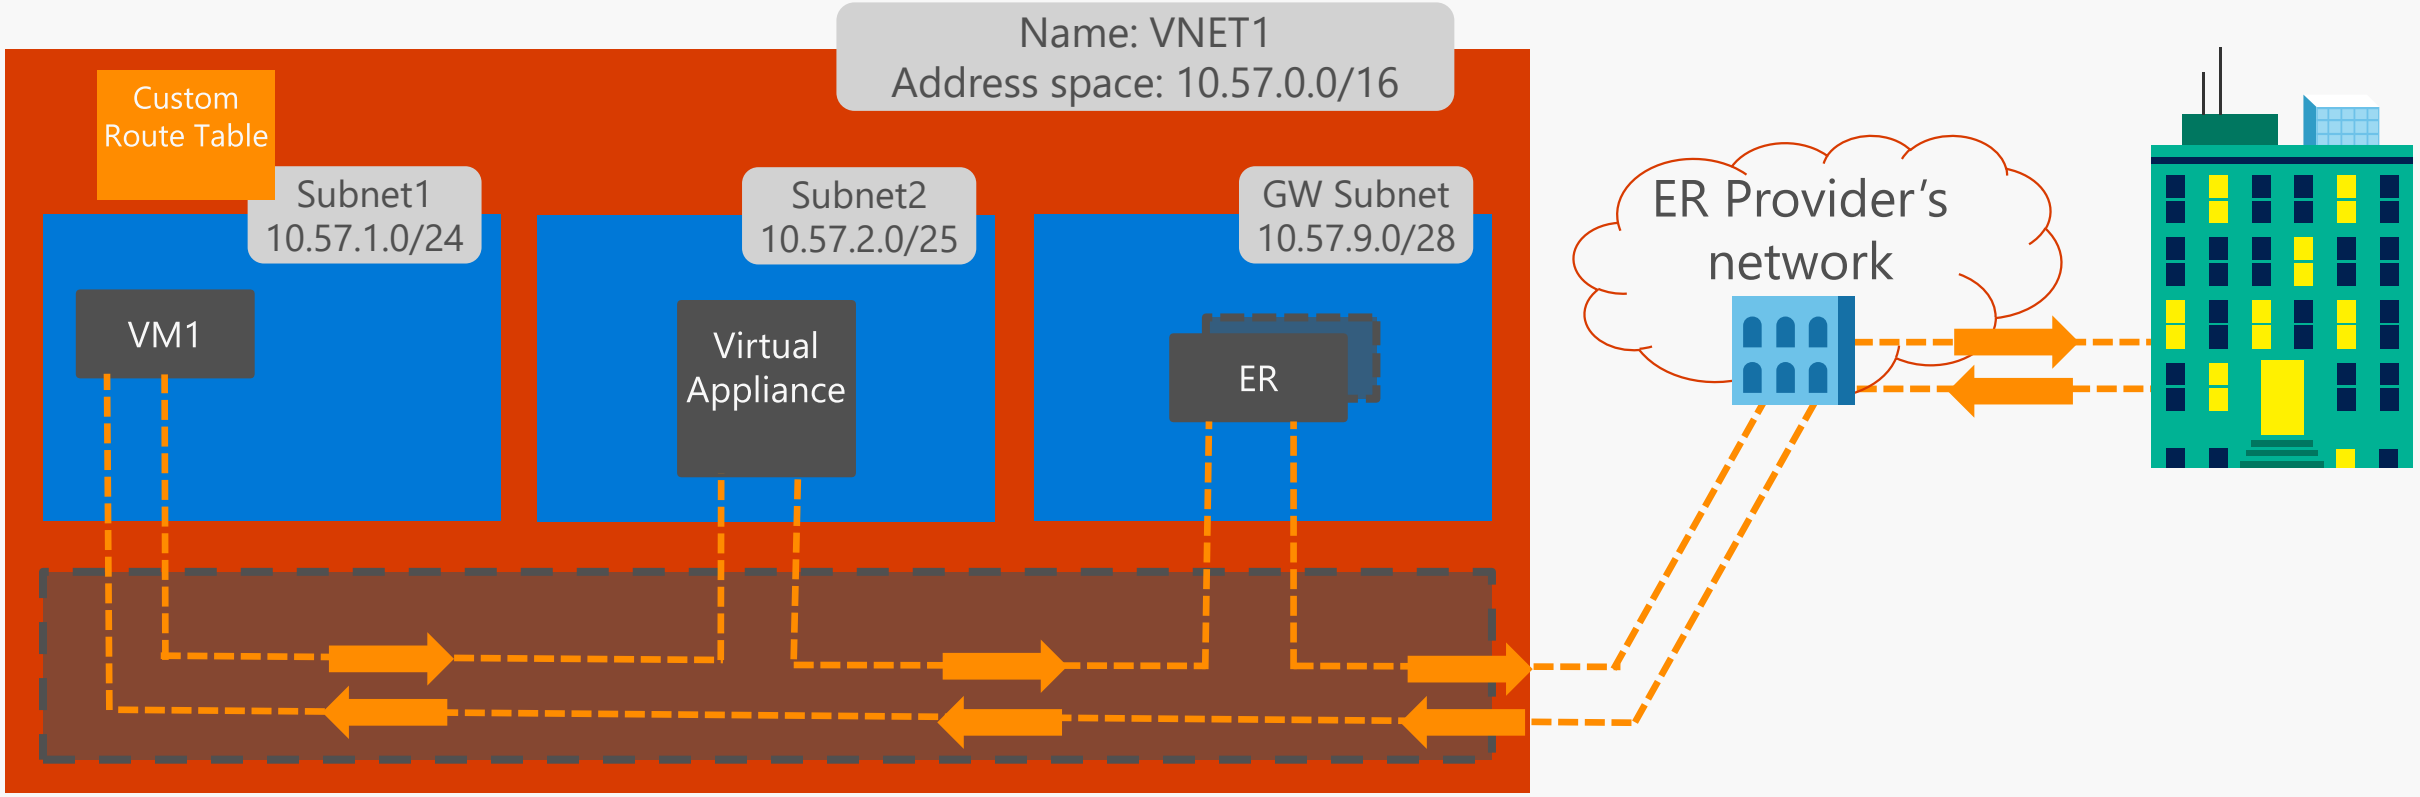

User Defined Routes (UDR's)

Use case 2: Inbound ER or S2S traffic

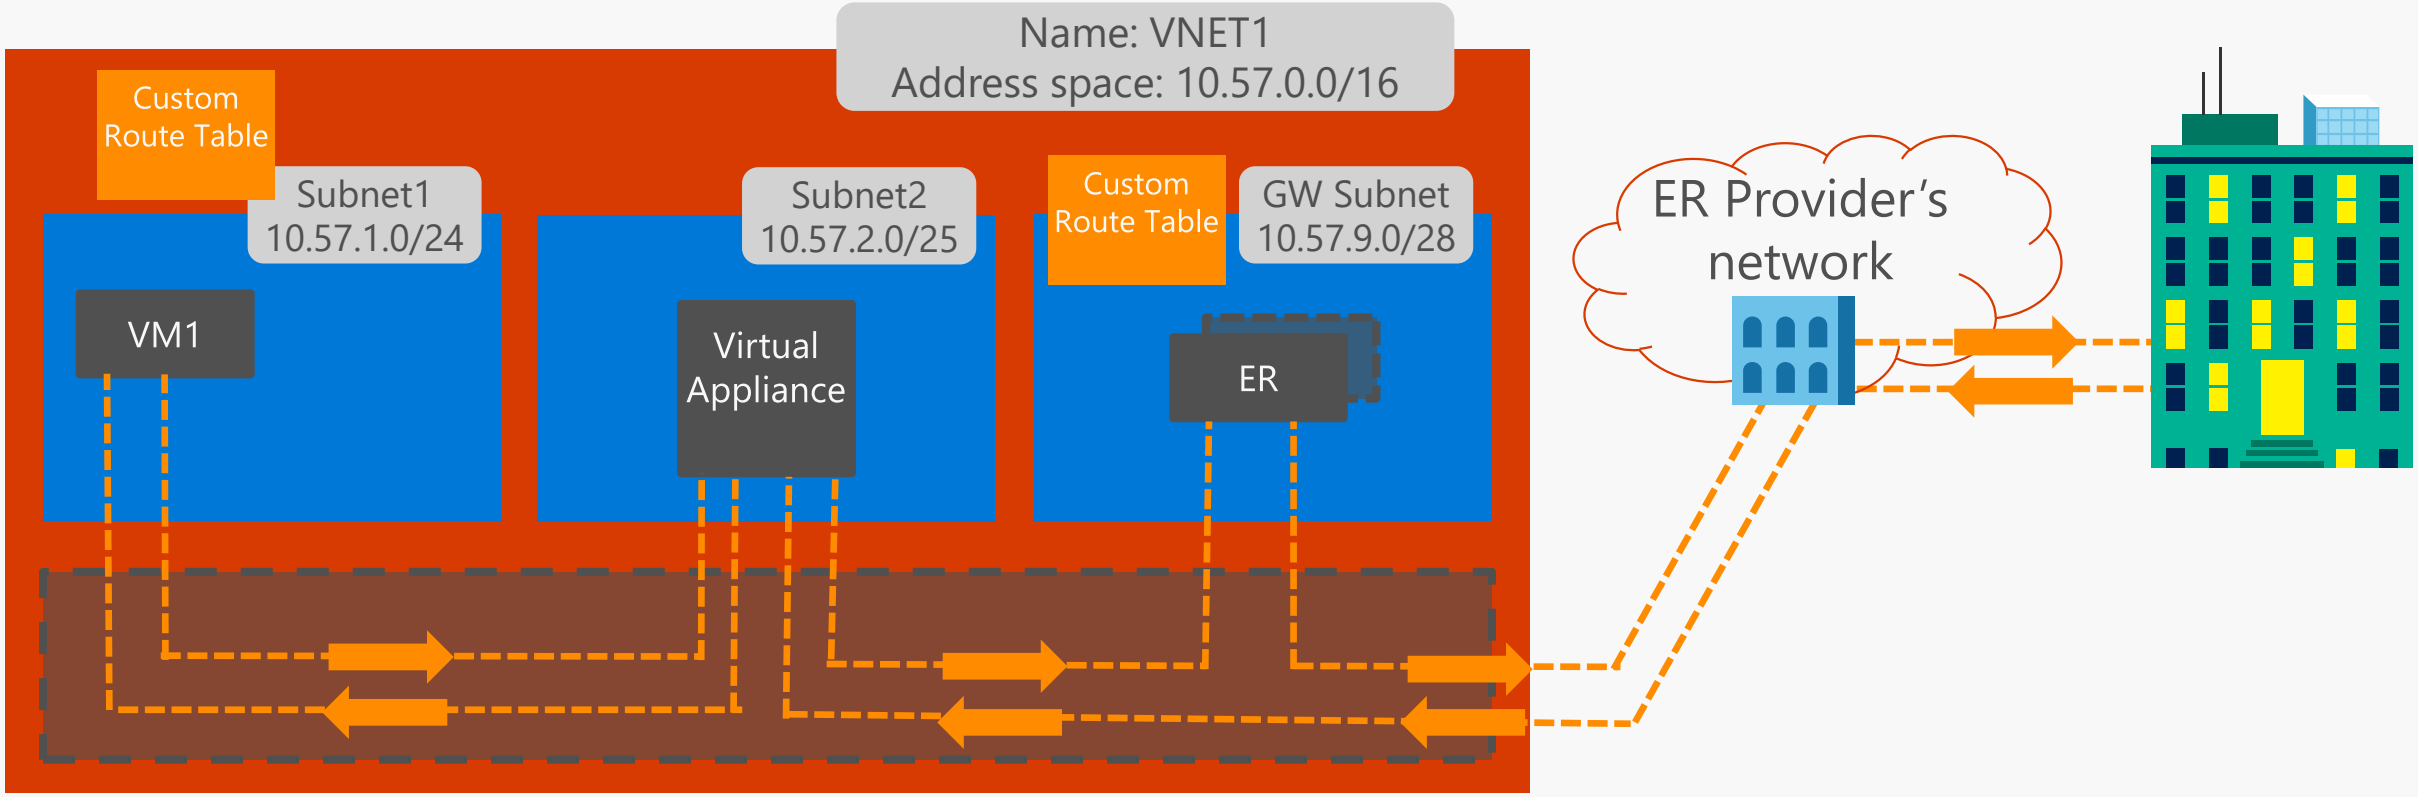

HA NVAs

Ensuring flow symmetry with Source-NAT

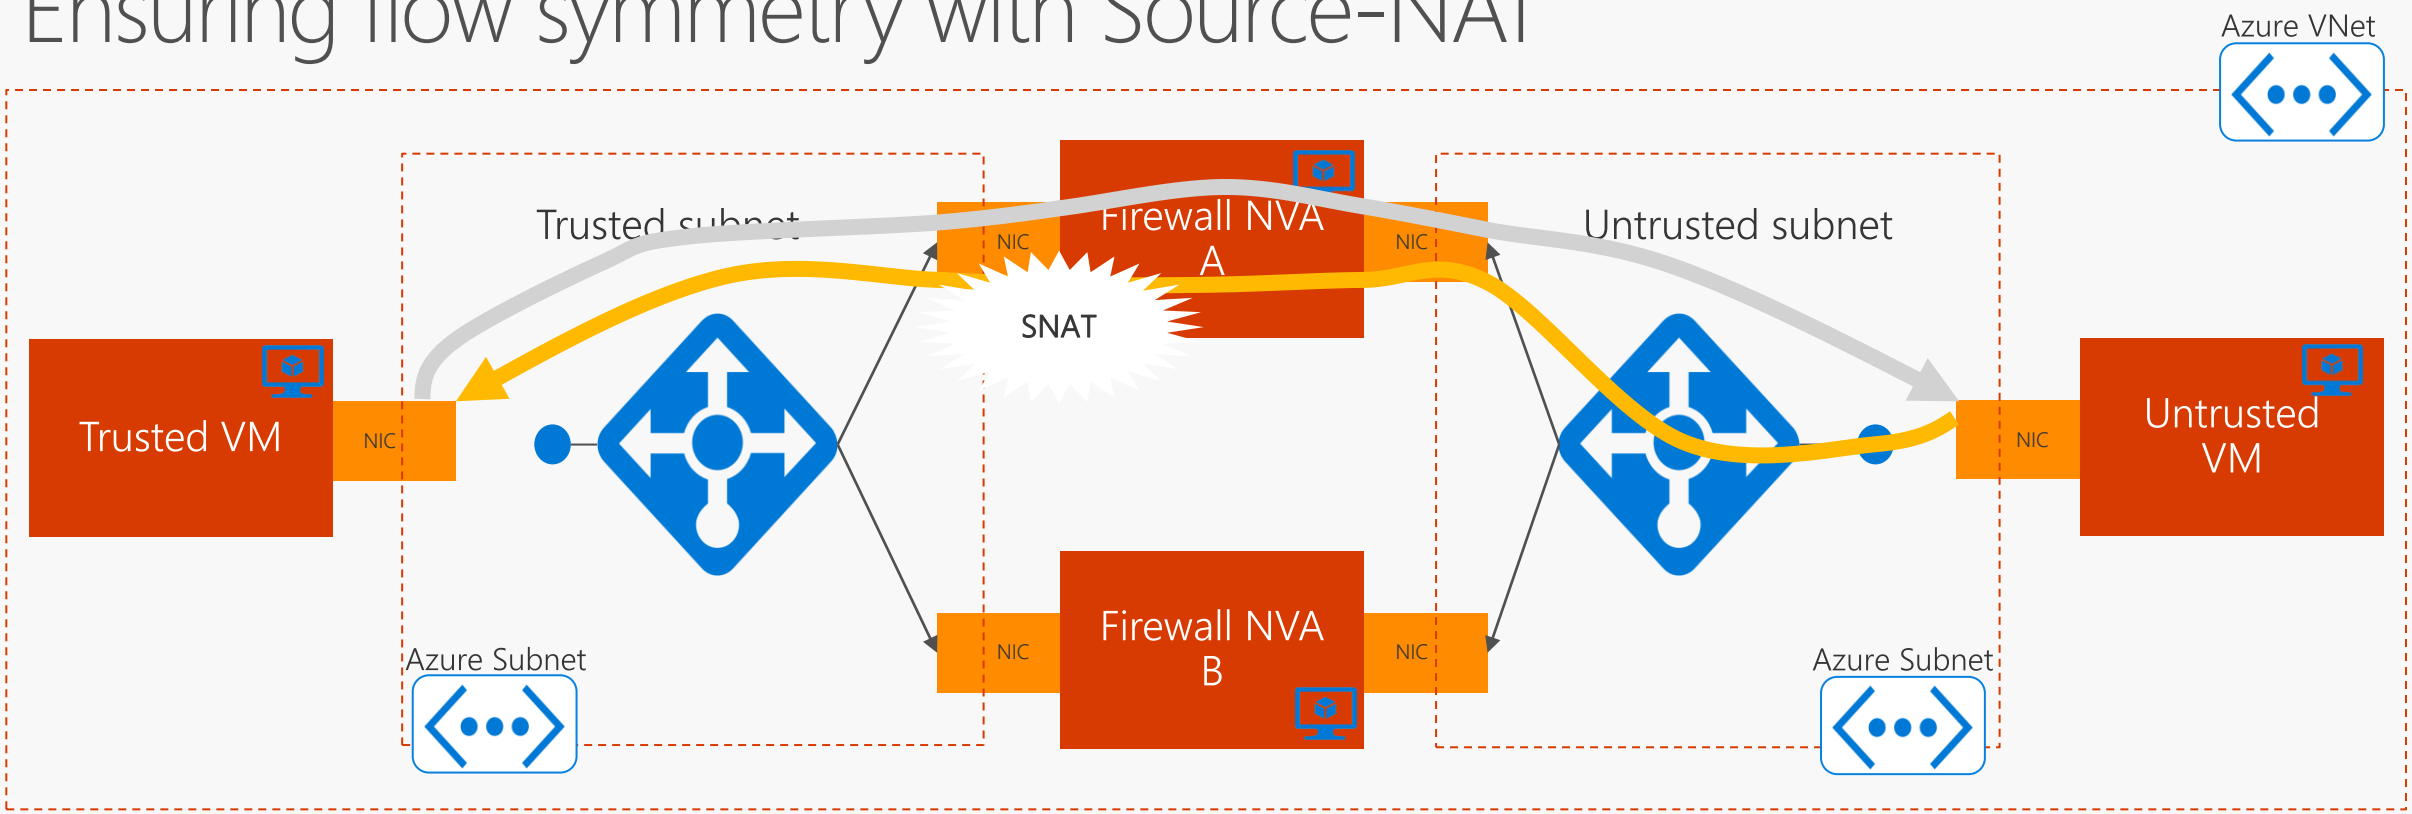

- Load balancer in “Untrusted” subnet assigns outbound flow to NVA “A”
- NVA “A” source-NATs traffic behind its “Trusted” subnet interface’s IP
- Return flow goes to NVA “A” without hitting the load balancer

Flow symmetry with single NIC configuration

Why is it not an issue?

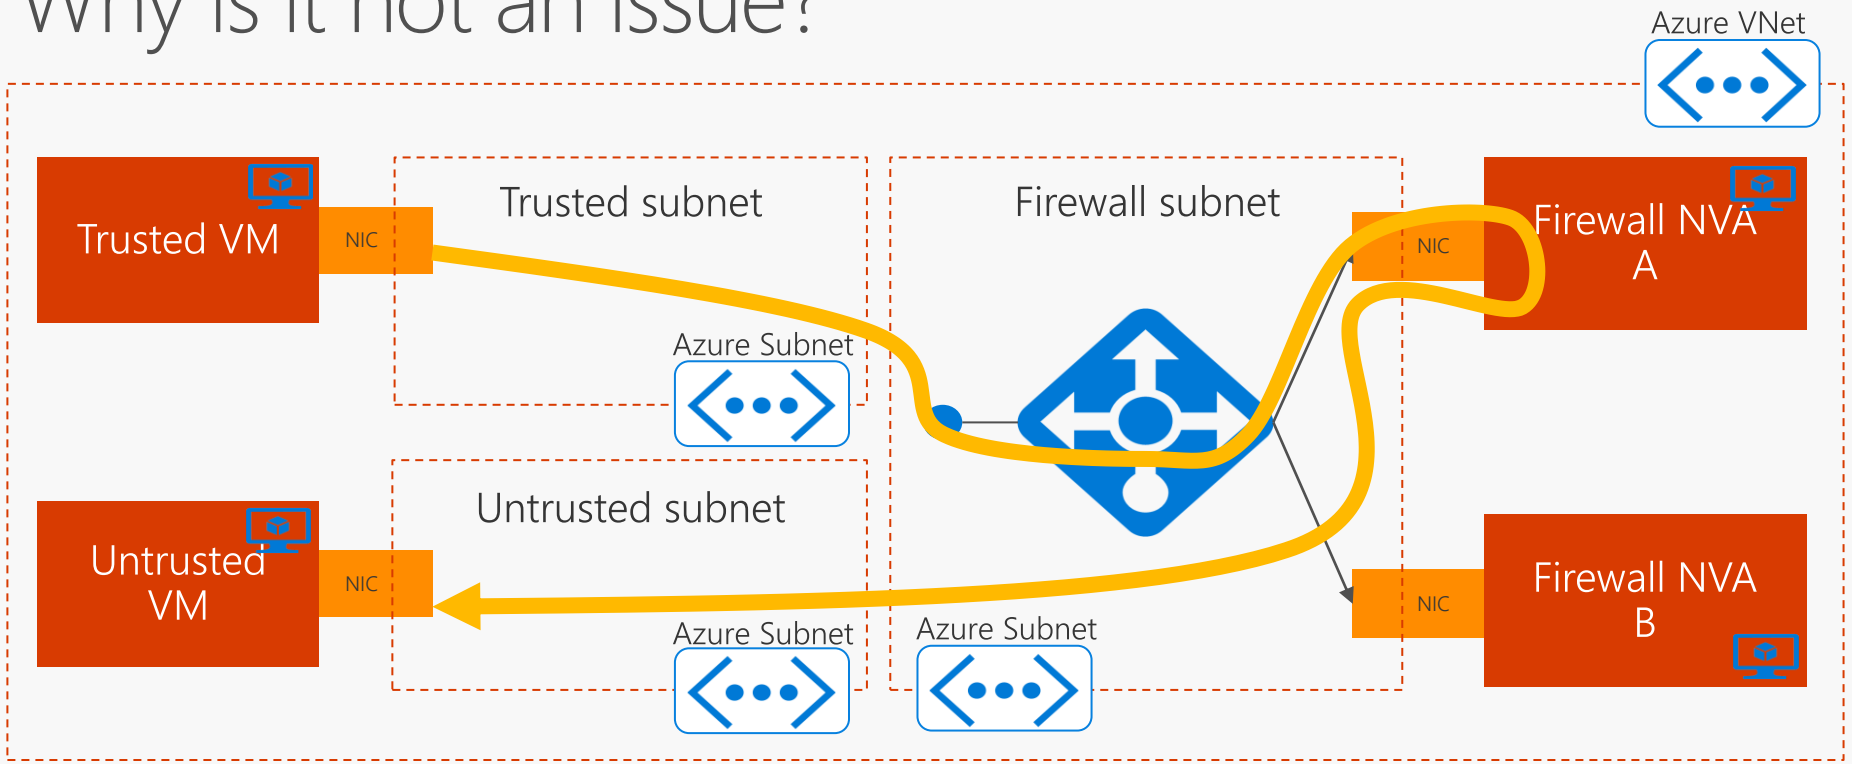

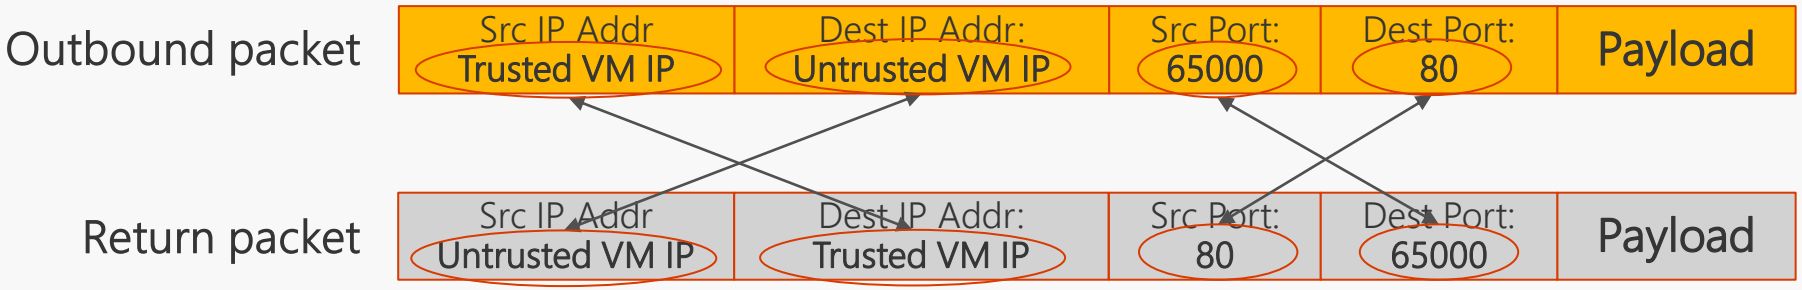

- Both packets have the same src/dest IP addresses and ports, in reverse order
- The load balancer's hashing algorithm assigns both packets to the same backend instance

VNet Peering

Peering connects 2 VNets together seamlessly

Works globally (across regions)!

No additional hop

Non-transitive (except gateway)

Can only ever have a single Gateway in a VNet (local or remote)

Important Addresses

KMS

kms.core.windows.net
https://docs.microsoft.com/en-us/azure/virtual-machines/windows/troubleshoot-activation-problems

DNS & monitoring

Including Load Balancer probes
168.63.129.16/32
https://blogs.msdn.microsoft.com/mast/2015/05/18/what-is-the-ip-address-168-63-129-16/

Q&A