

Azure Networking

Memilavi
www.memilavi.com

<https://t.me/learningnets>



Networking

- All aspects of networking in Azure
- Deals with resources' network connections, firewalls, etc.
- Might sound boring and not very important, but...

**Networking is the
foundation of cloud
security**

ReadIt!

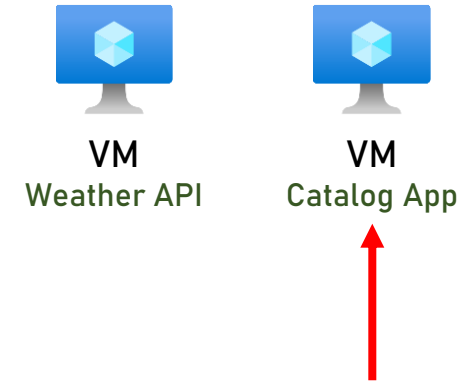
Cloud Architecture

A Word of Caution:

NEVER

leave a VM open to the internet this way

We will learn later on what should be done




- Directly accessible from the internet
- Can be RDPed from anywhere

ReadIt!

Cloud Architecture

Two main threats:

- Brute force attacks on port 3389 (RDP)
- No line of defense in front of the VM
web server



- Directly accessible from the internet
- Can be RDPed from anywhere

Networking knowledge is what makes a good cloud architect – an amazing cloud architect

Networking

- We'll talk about 4 networking-related cloud services:

VNets

SubNets

Load Balancer

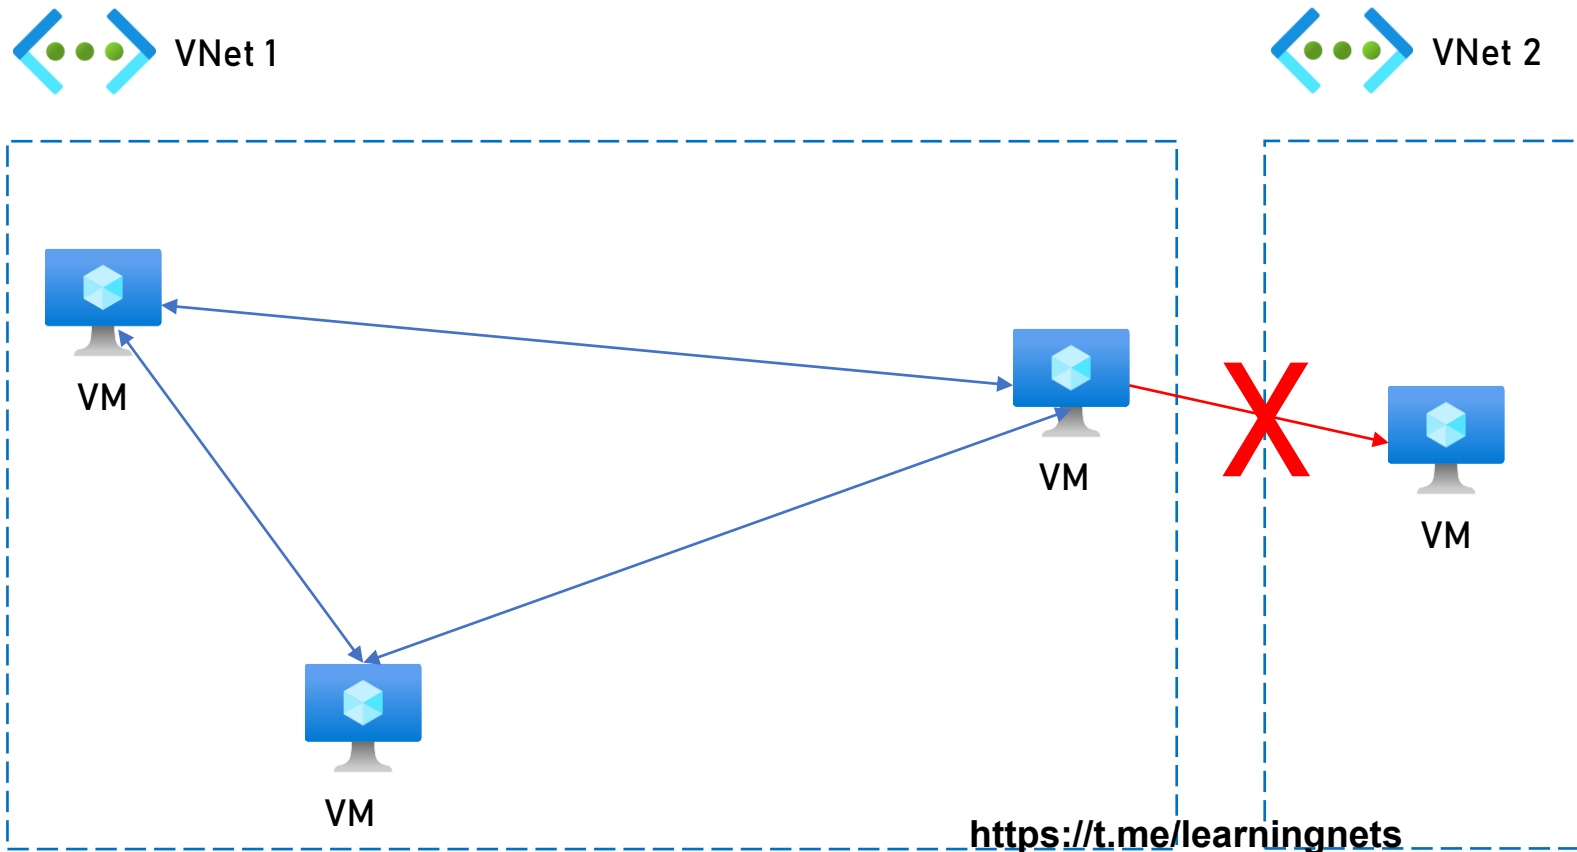
Application Gateway
<https://t.me/learningnets>

Virtual Networks

- A network in which you can deploy cloud resources
- Many cloud resources are deployed within Vnets
 - VMs, App Services, DBs, etc.
- “Virtual” as in “based on physical network and logically separated from other virtual networks”

Virtual Networks

- Resources in VNet can communicate with each other by default




... but not with
resources in other
VNets

Virtual Networks

- Think of it as your organization's private network
- In AWS it's called VPC – Virtual Private Cloud
- Other organizations' VNETs

cannot communicate with
your VNet



VNet Pricing

- VNets are free
- Limit of 50 VNets per subscription across all regions

Characteristics of VNets

- Scoped to a single Region
 - Cannot span multiple Regions
- Scoped to a single Subscription
- Can be connected via Peering
- Segmented using Subnets
- Protected using NSG (on the Subnets)

Security and VNets

- The most important thing to think about when designing network:

How to limit access to the resources in the VNet so that risk is minimized

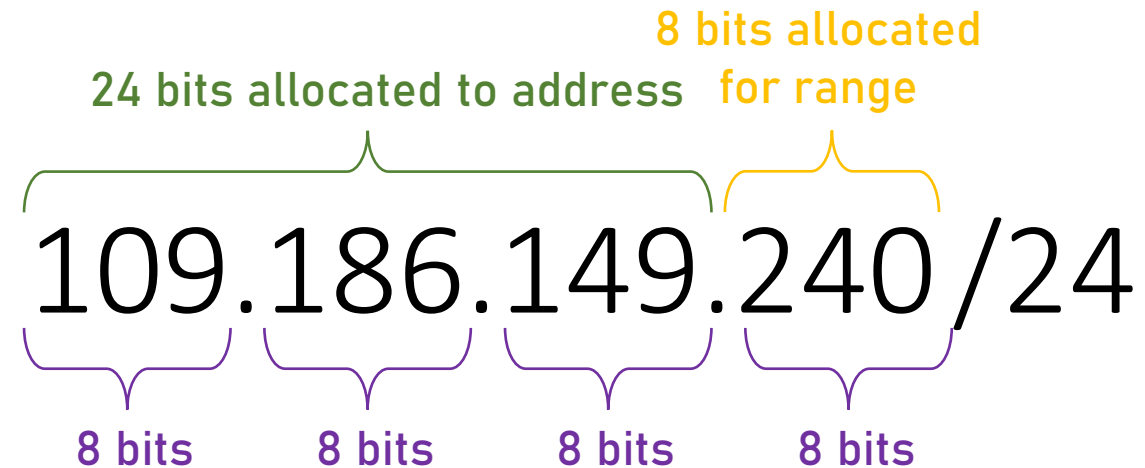
Addresses of VNets

- Each VNet has its own address range
 - Or IP Range
- By default – 65,536 addresses
- Can be customized
- All network devices must be in this address range
- Expressed using CIDR Notation

CIDR Notation

- Classless Inter-Domain Routing
- A method for representing an IP Range
- Composed of an address in the range and a number between 0 and 32
- The number indicates the number of bits that are allocated to the address. The smaller the number – the larger the range


CIDR Notation Example #1



109.186.149.000 – 109.186.149.255
256 Addresses

Bits refresher:
00000000 = 0
11111111 = 255

CIDR Notation Example #2



109.186.000.000 – 109.186.255.255

65,536 Addresses

Bits refresher:
00000000 = 0
11111111 = 255

Probably way too big...

<https://t.me/learningnets>

CIDR Notation Example #3

149 Dec = 1001 ~~0101~~ Bin

1001 0000 Bin = 144 Dec

109.186.144.000 – 109.186.159.255

4,096 Addresses

Bits refresher:
00000000 = 0
11111111 = 255

CIDR Notation

- The good news:

You don't have to remember!

- A lot of CIDR calculators
 - ie. <https://www.ipaddressguide.com/cidr>
<https://t.me/learningnets>

CIDR Notation

- More good news:
 - Azure usually shows the actual range

IPv4 address space	
192.168.0.0/16	192.168.0.0 - 192.168.255.255 (65536 addresses)

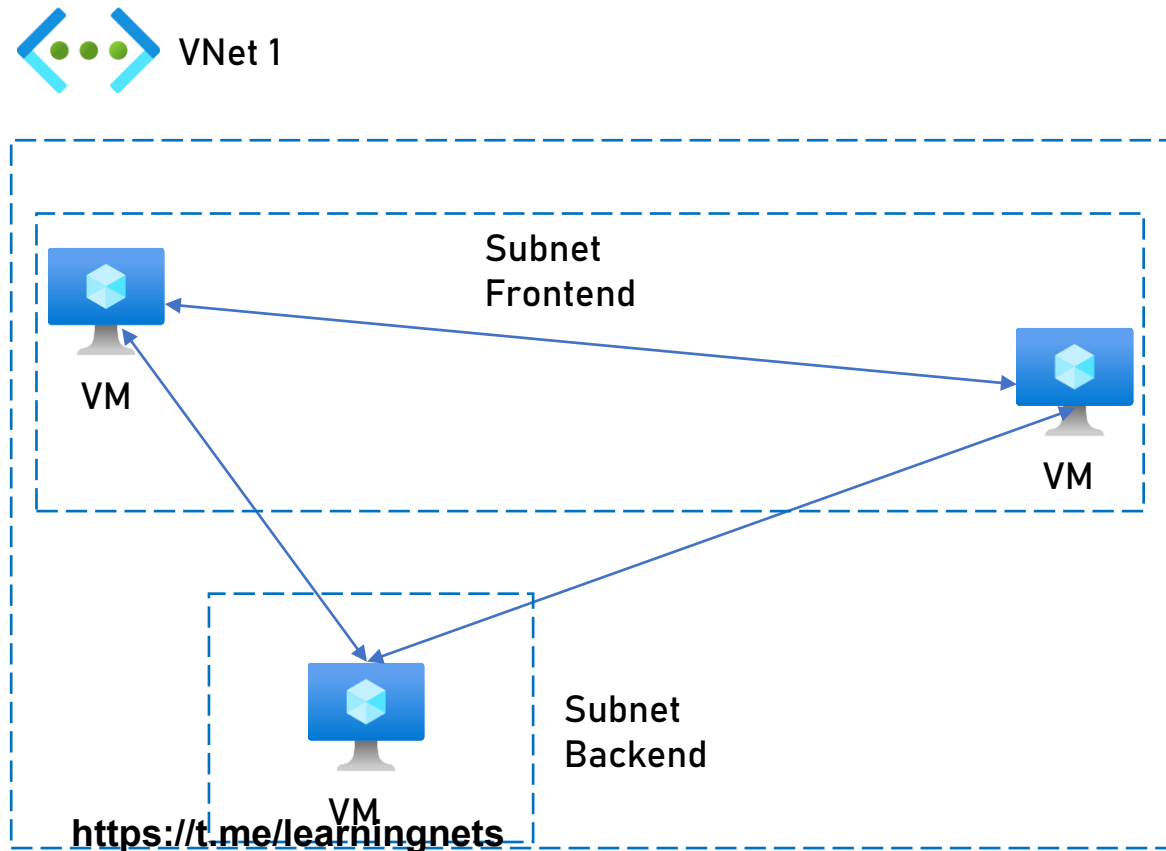
Subnet

- A logical segment in the VNet
- Shares a subset of the VNet's IP Range
- Used as a logical group of resources in the VNet
- Is a must. Resources must be placed in a Subnet, cannot be placed directly in the VNet

Subnet

- Resources in a subnet can talk to resources in other subnets in the same VNet*

*By default, can be customized



Addresses of Subnets

- Each Subnet gets a share of the parent VNet's IP Range
- NEVER use the full range of the VNet in a Subnet
- Extremely hard to modify the range later
- Makes it hard to add future Subnets

Subnet Pricing

- Subnets are free
- Limit of 3,000 Subnets per VNet

Network Security Group

- Usually called NSG
- A gatekeeper for Subnets
- Defines who can connect in and out of subnet
- Think of it as a mini-firewall
- Should be a standard part of Subnet creation
- Is free

How NSG Works?

- Looks at 5 tuples:
 - Source (=Where did the connection come from)
 - Source Port (=The port the source is using)
 - Destination (=Where does the connection request goes)
 - Destination Port (=To which port does it want to connect)
 - Protocol (=TCP, UDP, Both)

How NSG Works?

- Based on these 5 tuples the connection is either allowed or denied
- This is called Security Rule
- Each rule is assigned a number
- The lower the number – the higher the priority of the rule

NSG and VMs

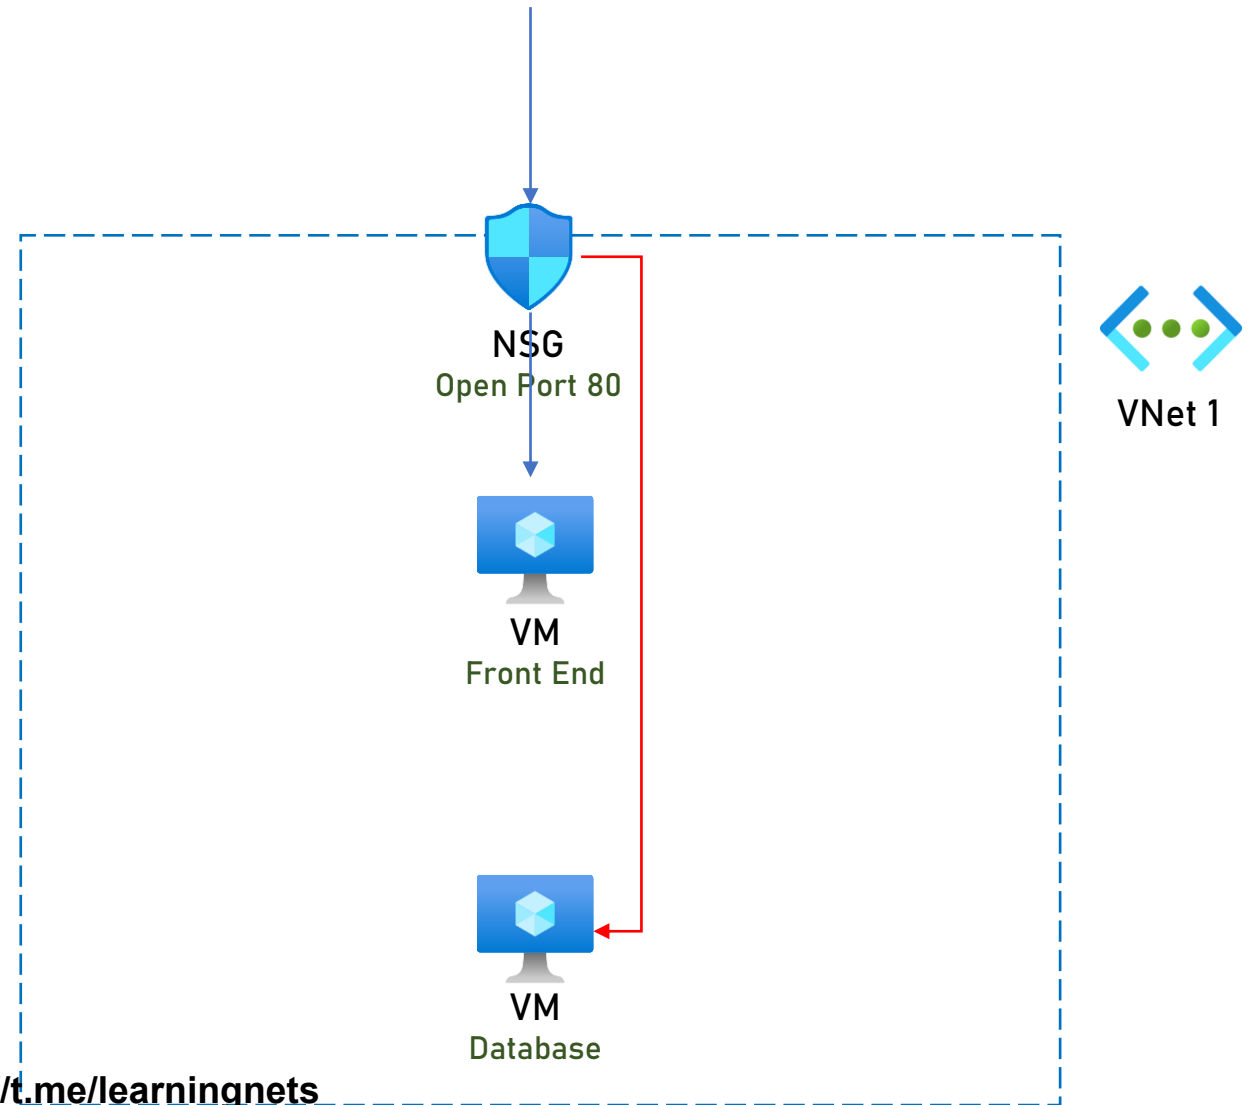
- An NSG is automatically created and attached to every newly-created VM's network interface
- By default – open RDP (on Windows) or SSH (on Linux) port to anyone
- **MUST** be handled first thing after creation

Network Peering

- Sometimes, to increase security, we want to place some resources in a completely different VNet
 - Not just Subnet!
- Examples:
 - Separate systems
 - System layers
 - Sensitive databases <https://t.me/learningnets>

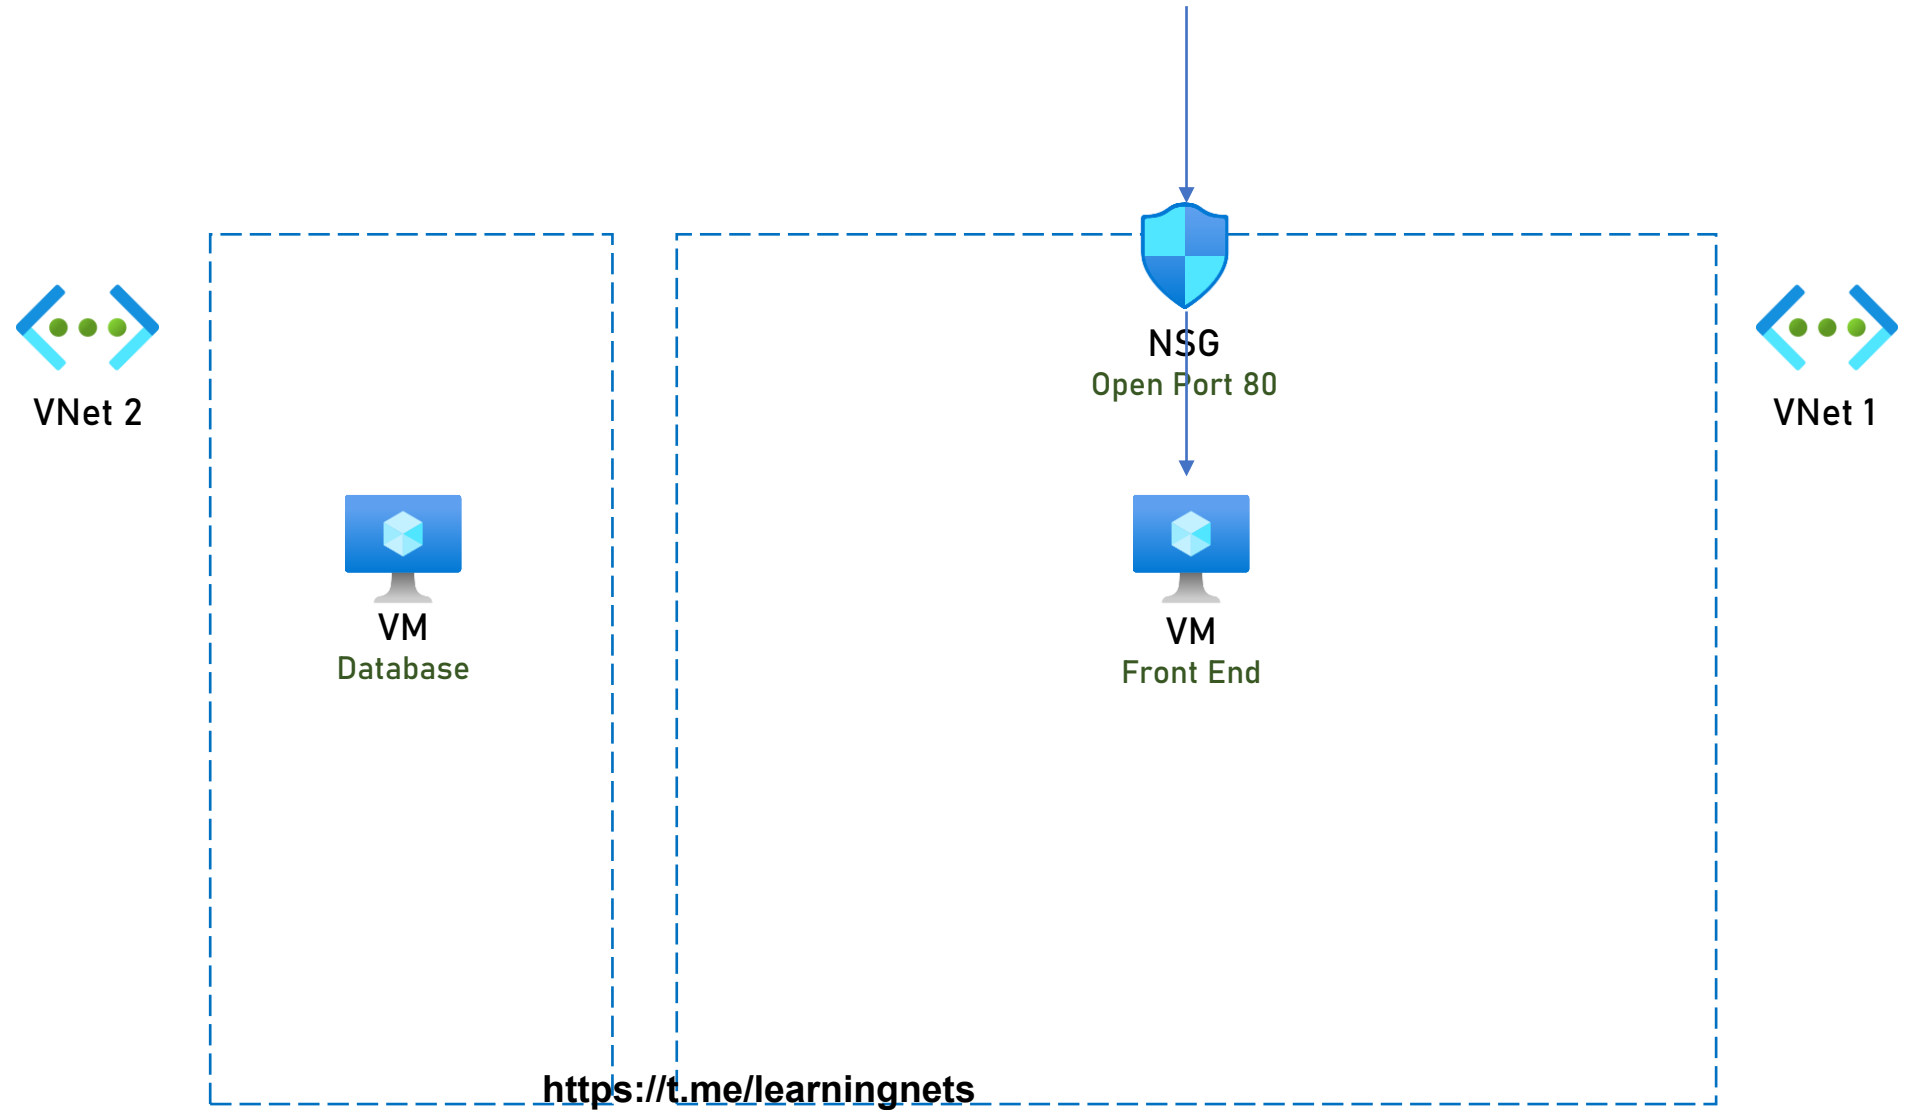
Network Peering

- Main reasoning:
- Not to place non-public resources in a VNet that has public access



Network Peering

- So...

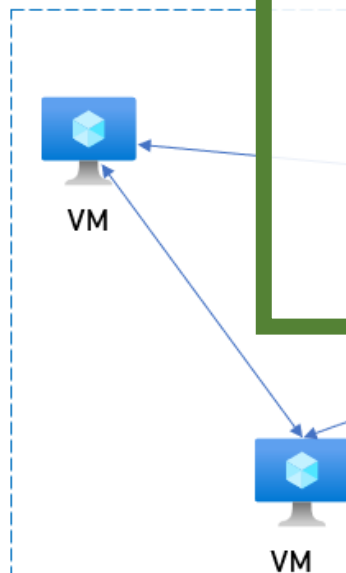


Network Peering

- But...


- Resources in VNet can communicate with each other by default

 VNet 1



Network Peering
to the rescue!

 VNet 2



... but not with
resources in other

VNets

Network Peering

- Allows two VNets to connect to each other
- From the user's point of view it's a single VNet
- Make sure address spaces are not overlapped!
- Use NSG for protection
- Can work across Regions

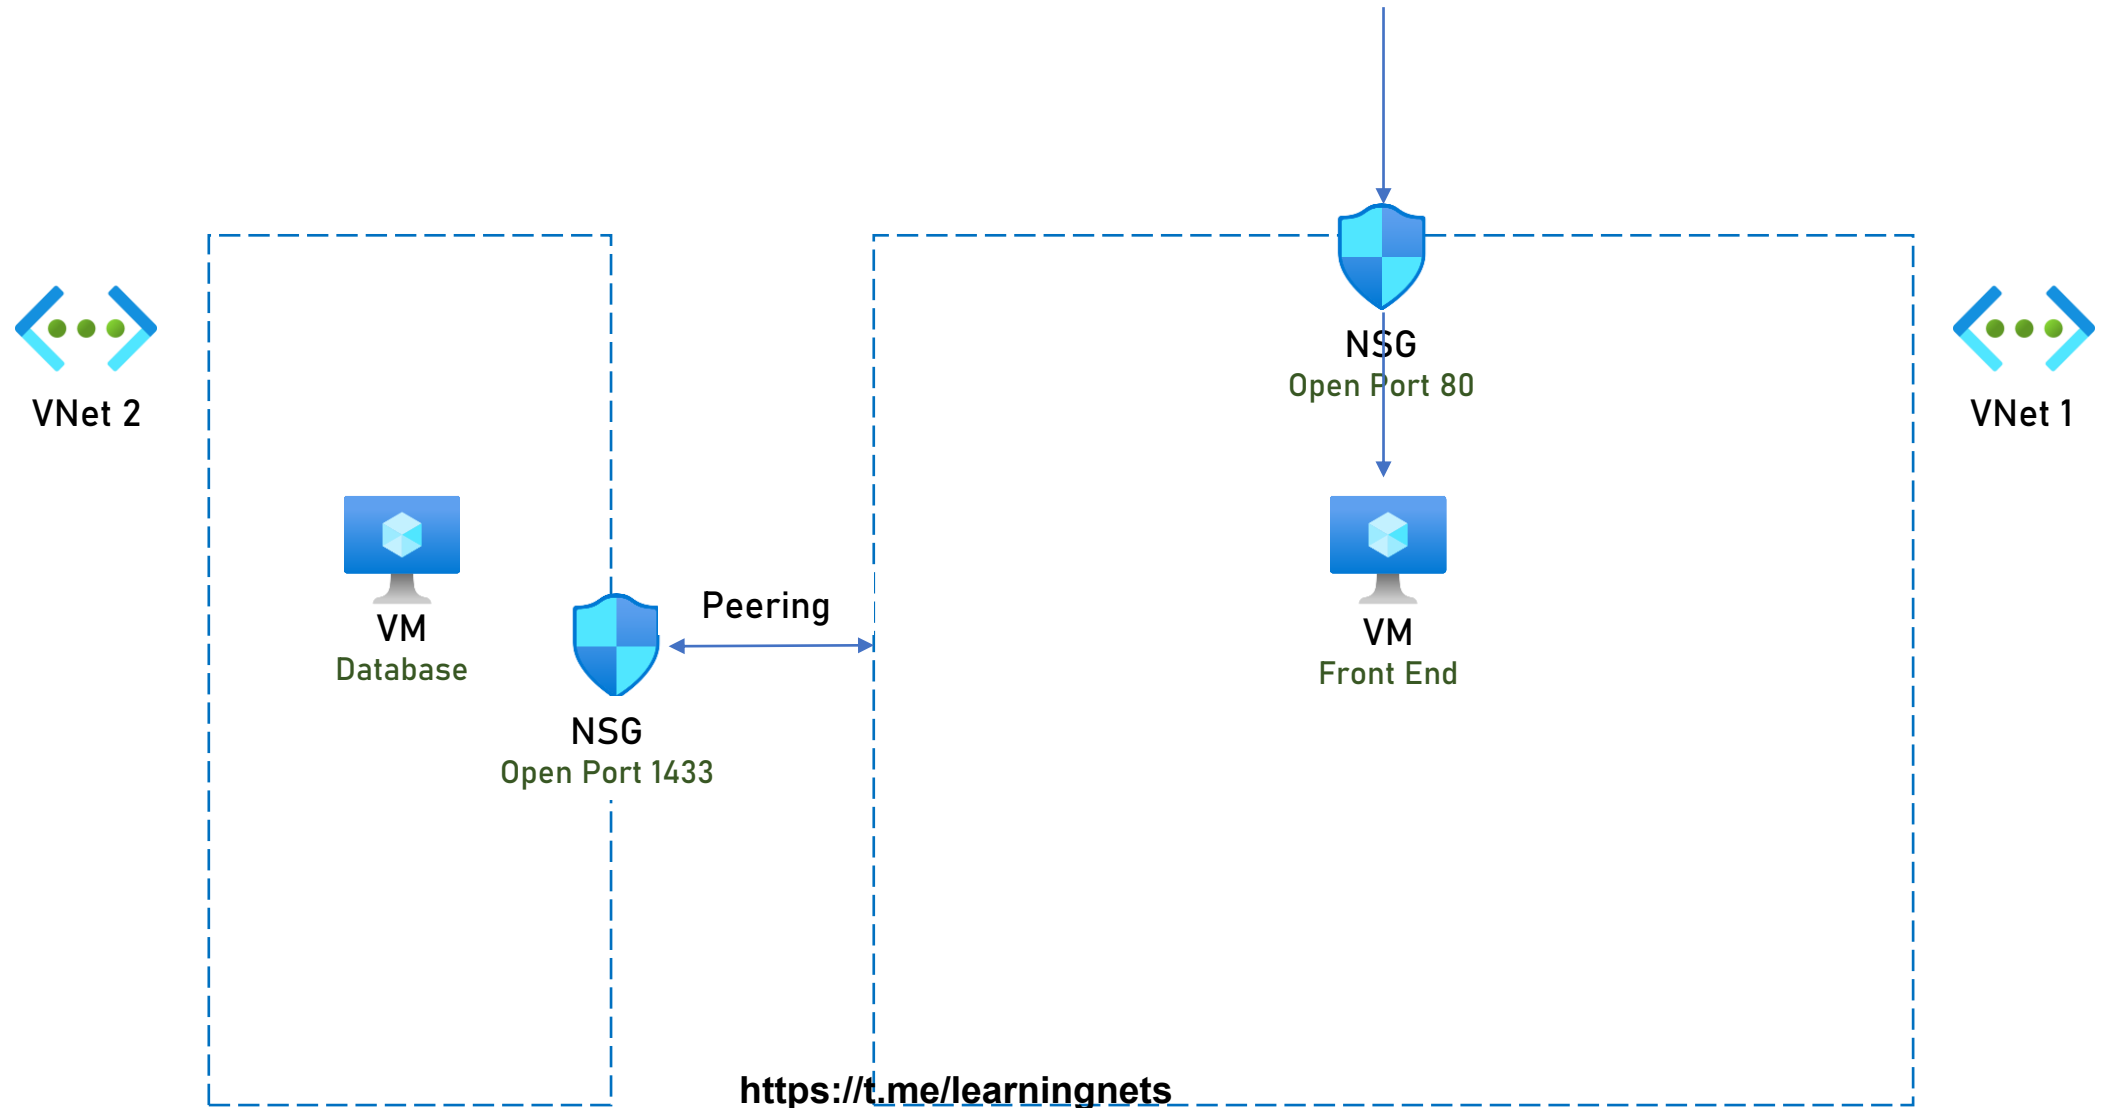
Network Peering

- Not free

Outbound data transfer		
100	×	\$0.0100
GB		Per GB

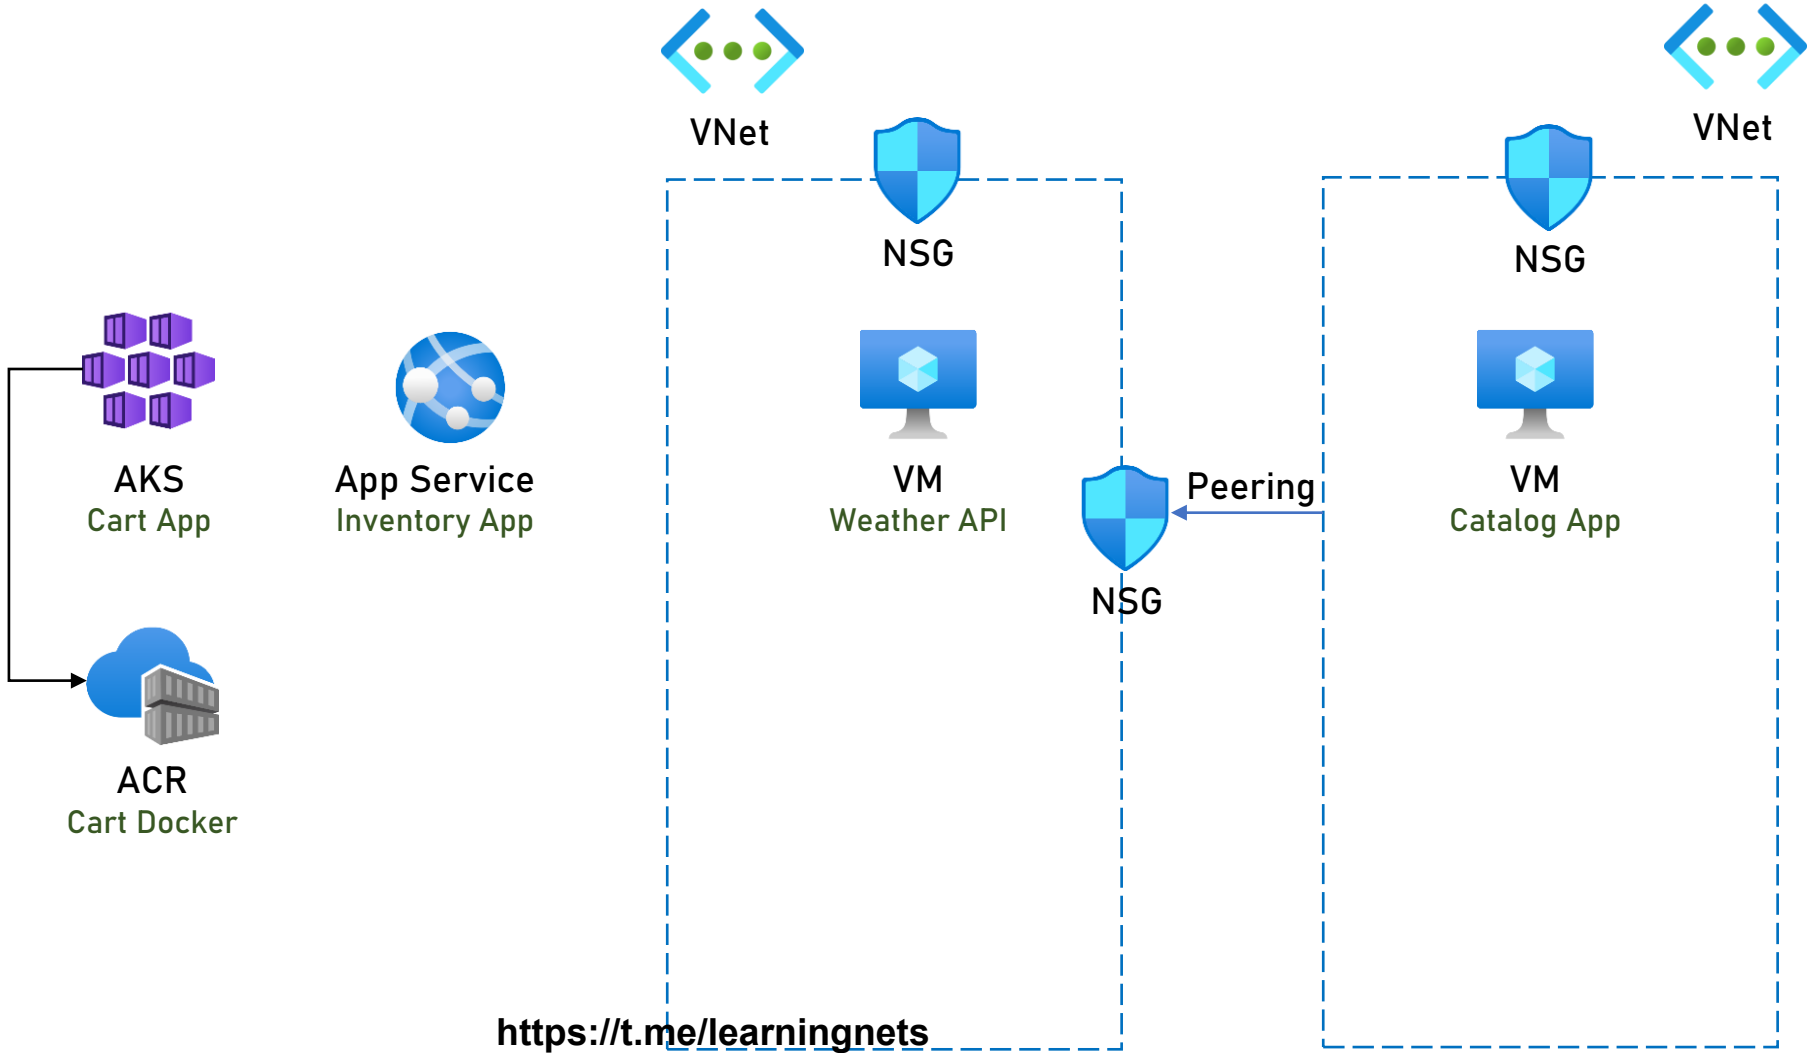
Inbound data transfer		
100	×	\$0.0100
GB		Per GB

Network Peering



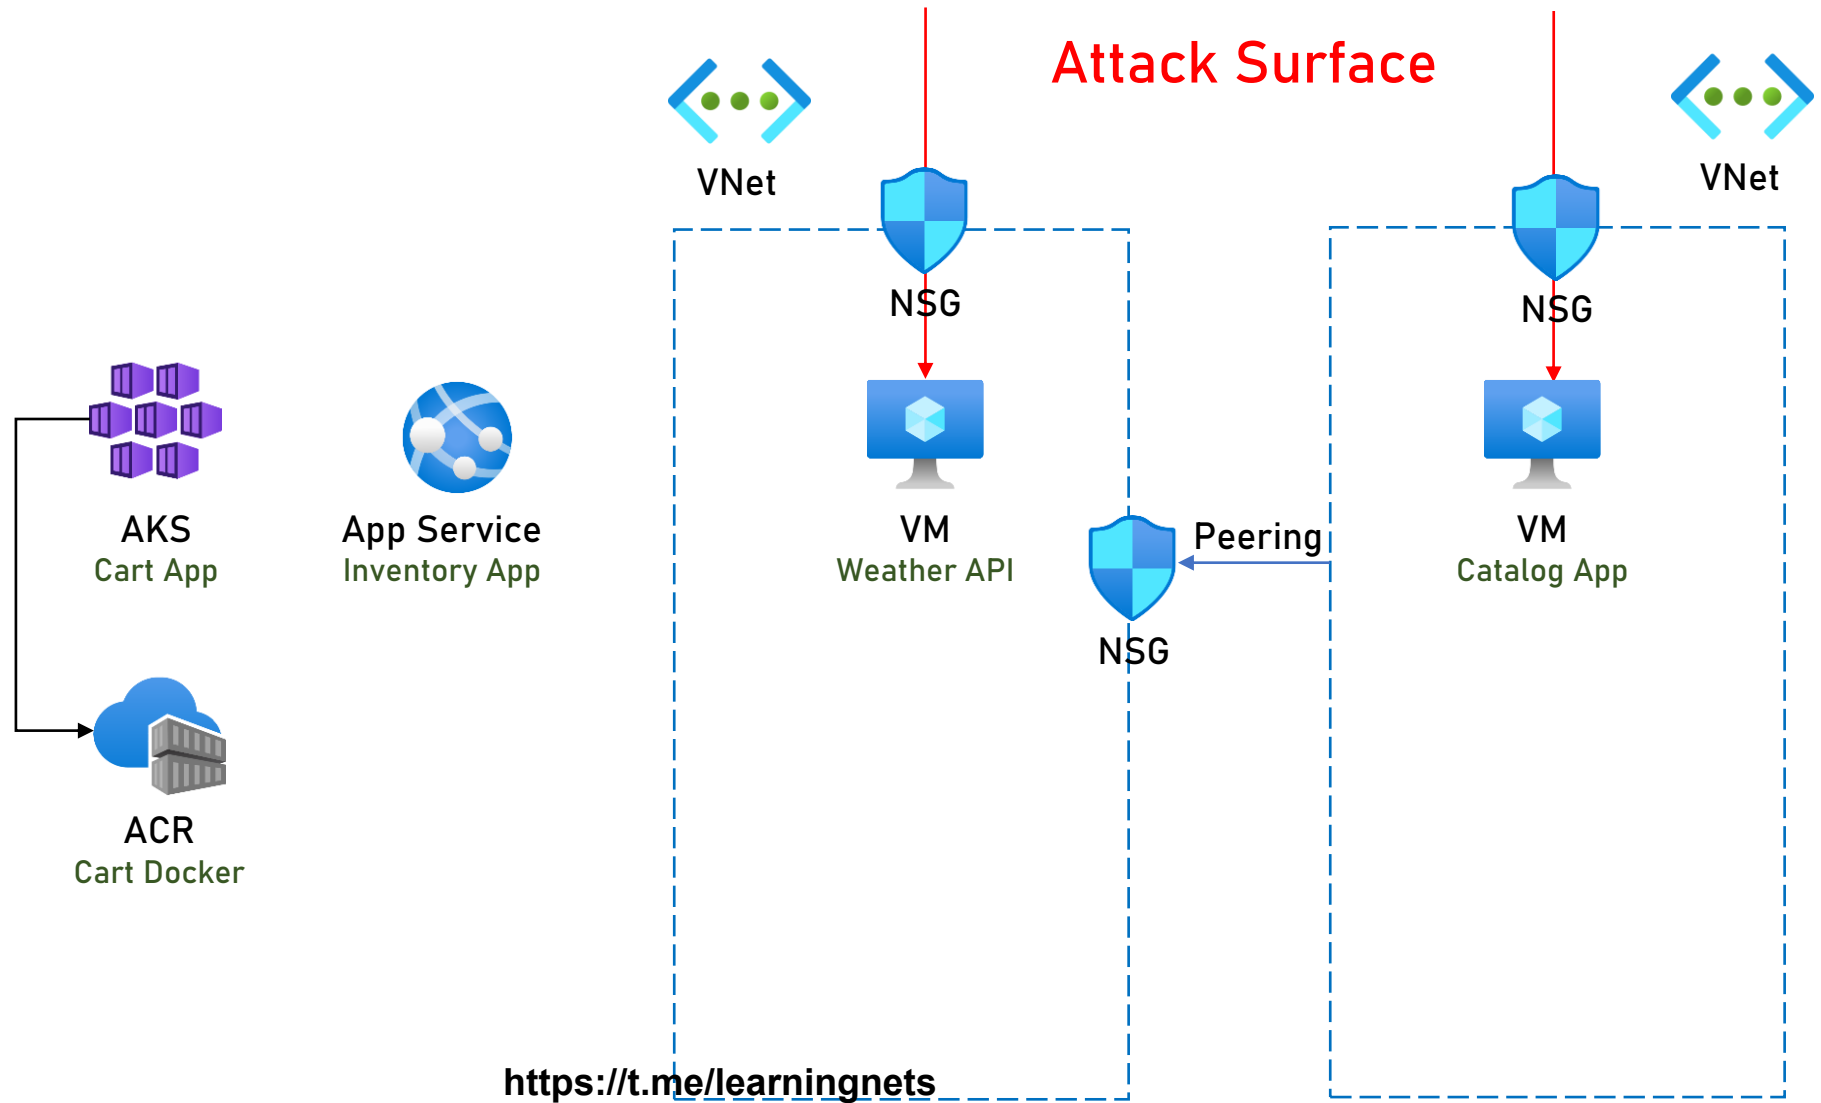
ReadIt!

Cloud Architecture



ReadIt!

Cloud Architecture



Secure VM Access

- The larger the attack surface – the greater the risk
- We want to minimize it as much as possible
- Leaving public IPs open is always a risk we want to avoid
- Not directly related to the app design but important nonetheless

Secure VM Access

- What can be done?

JIT Access

VPN

Jump Box

Bastion

<https://t.me/learningnets>

JIT Access

- **Just In Time Access**
- **Opens the port for access on demand, and automatically closes it**
- **Rest of the time – it's closed**
- **Can be configured from the VM's page in the portal**
- **Requires Security Center License Upgrade**

VPN

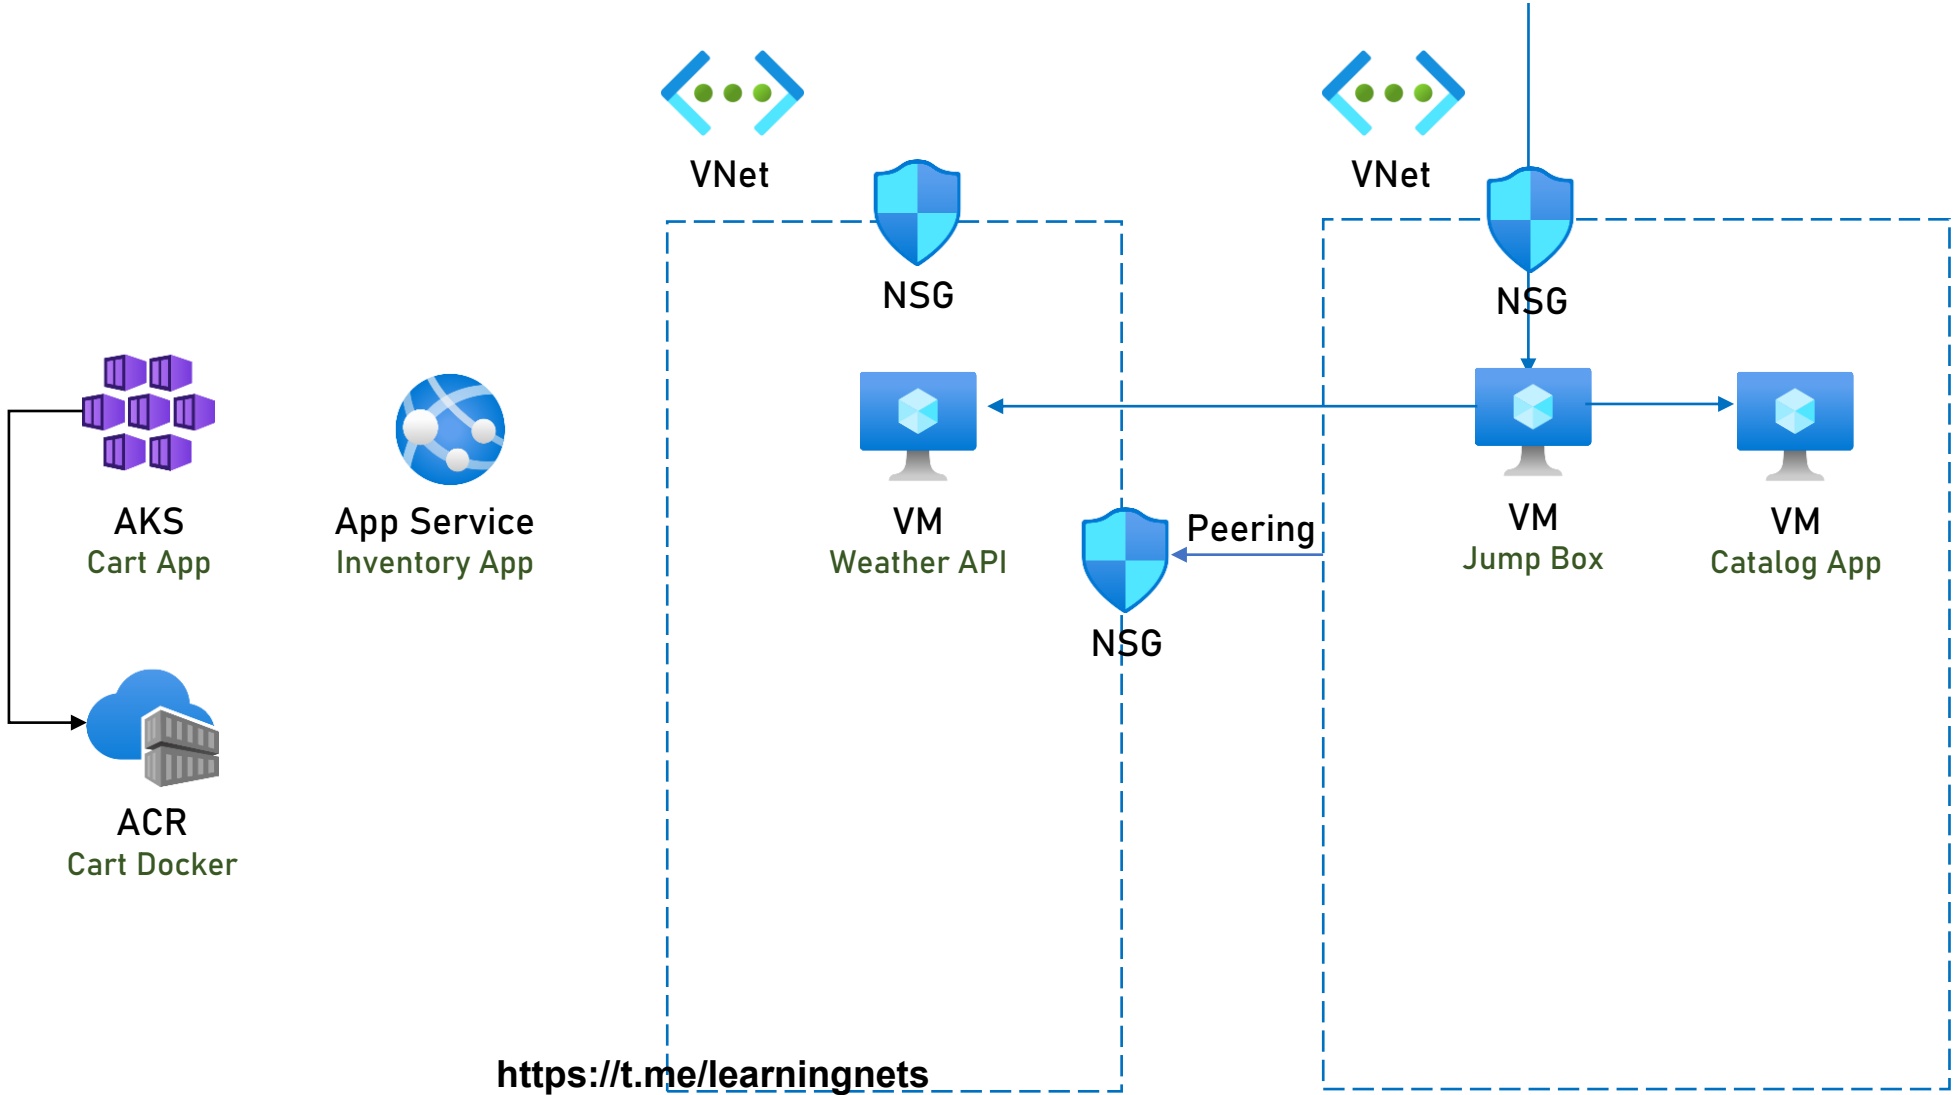
- A secure tunnel to the VNet
- Can be configured so that no one else can connect to the VNet
- Requires VPN software and license (not part of Azure)

Jump Box

- Place another VM in the VNet
- Allow access ONLY to this VNet
- When need to access one of the other VMs – connect to this one and connect from it to the relevant VM
- Only one port is open (still kind of a problem...)
- Cost: The additional VM (the Jump Box)

ReadIt!

Cloud Architecture



Bastion

- A web-based connection to the VM
- No open port is required
- Simple and secure
- Cost: ~140\$ / month

Bastion Downsides

- Cost
- Requires portal access
 - This is said to be handled by the Bastion team 😊

Service Endpoint

- A lot of managed services expose public IP
 - ie. Azure SQL Server, App Services, Storage and more
- Sometimes these resources are accessed only from resources in the cloud
 - ie. Database in the backend
- Might pose a security risk

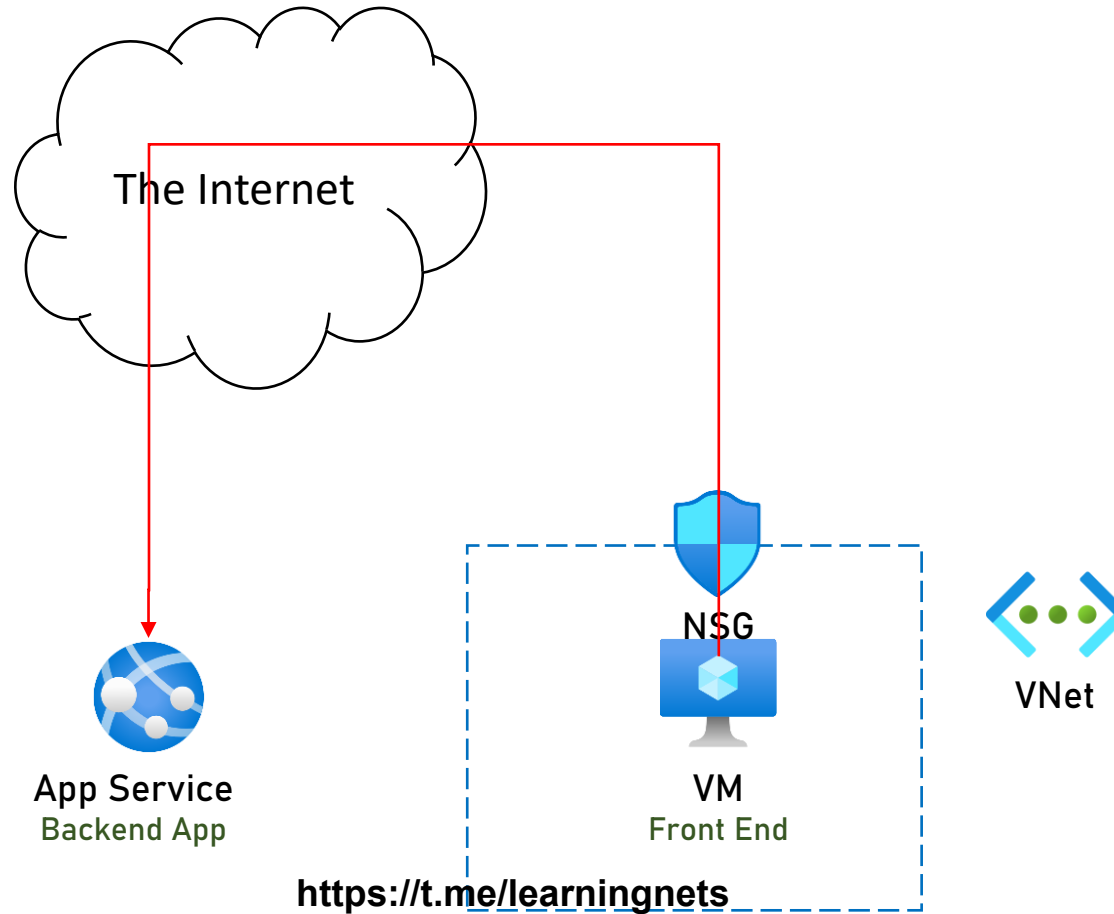
Service Endpoint

- Service Endpoint solves this security risk
- Creates a route from the VNet to the managed service
- The traffic never leaves Azure backbone
 - Although the resource still has a public IP
- Access from the internet can be blocked
- Is free!

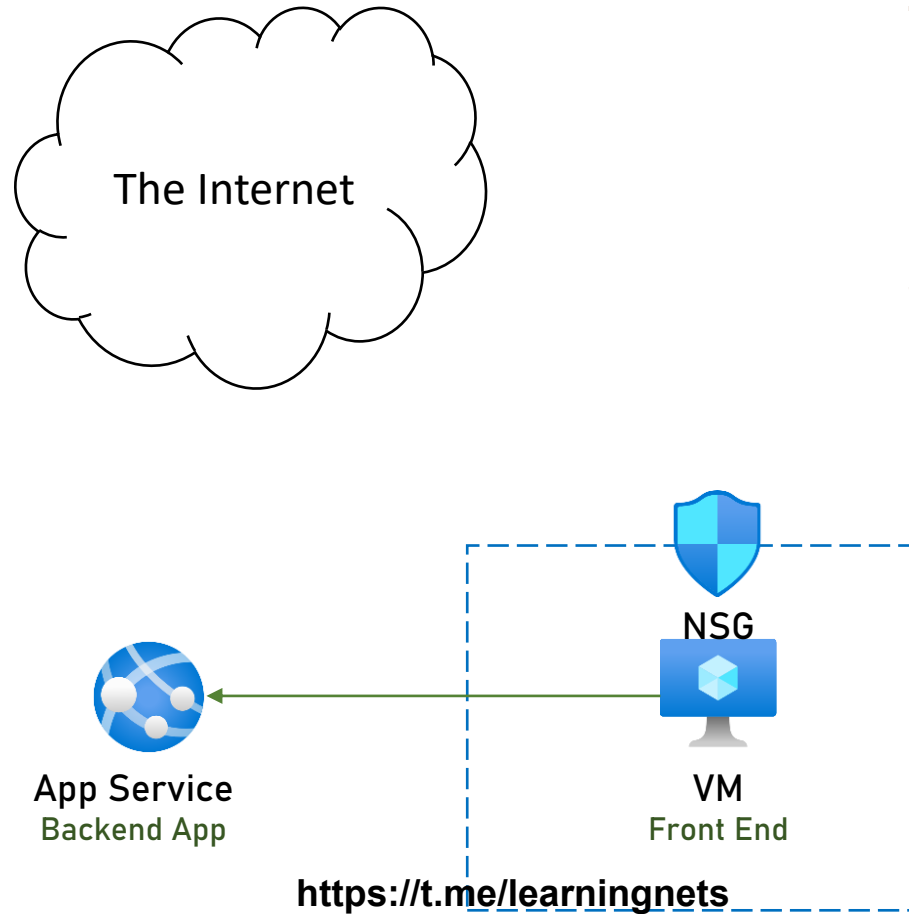
Service Endpoint

- How it's done:
 - Enable Service Endpoint on the Subnet from which you want to access the resource
 - On the resource, set the subnet as the source of traffic
 - Voila!

Without Service Endpoint



With Service Endpoint



Note:

1. Traffic leaves the VNet
2. There is a public IP on the PaaS service (App Service)
3. Can't be used from on-prem network
 - Almost...



Service Endpoint

- Resources support Service Endpoint:
 - Storage
 - SQL Database
 - Synapse Analytics
 - PostgreSQL
 - MySQL
 - Cosmos DB
 - KeyVault
 - Service Bus
 - Event Hub
 - App Service
 - Cognitive Services

Private Link

- A lot of managed services expose public IP
 - ie. Azure SQL Server, App Services, Storage
- Sometimes these resources are accessed only, *I've heard that before...* resources in the cloud
- ie. Database in the backend
- Might pose a security risk



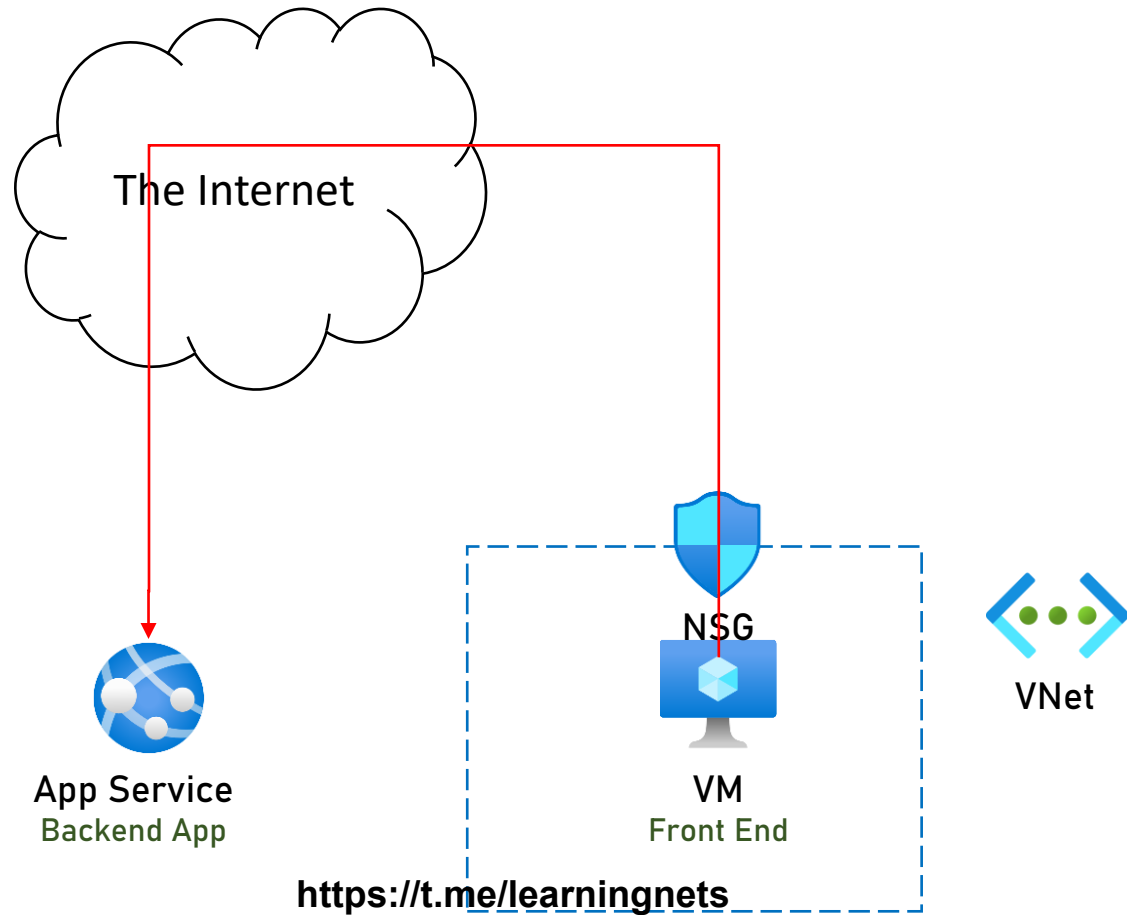
Private Link

- A newer solution to this problem
- Extends the managed service into the VNet
- The traffic never leaves the VNet
- Access from the internet can be blocked
- Can be used from on-prem networks
- Isn't free


Private Link

- How it's done:
 - Configure the resource to connect to the VNet
 - Configure private DNS
 - Might cause a problem if you have your own DNS

Without Private Link

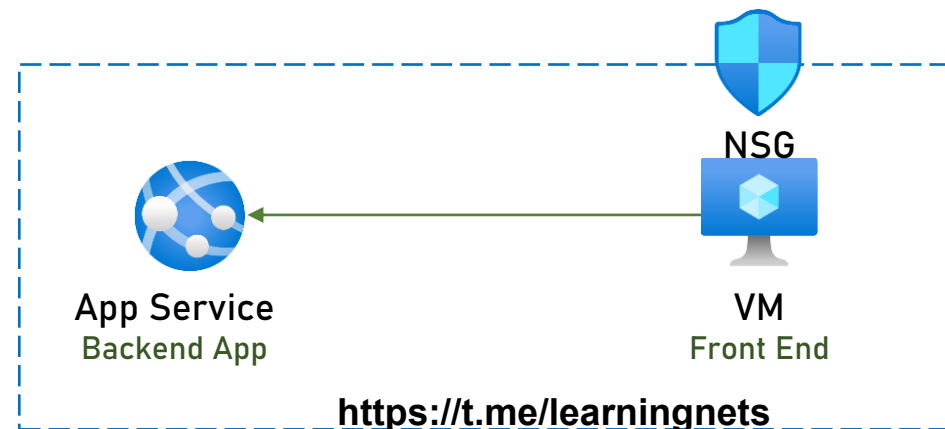


With Private Link



Note:

1. Traffic never leaves the VNet
2. The VM talks to the App Service via private IP
3. Can be used from on-prem network



Private Link

- Resources support Private Link:
 - Storage
 - SQL Database
 - Synapse Analytics
 - PostgreSQL
 - MySQL
 - Cosmos DB
 - KeyVault
 - Redis
 - AKS
 - Search
 - ACR
 - App Configuration
 - Backup
 - Service Bus
 - Event Hub
 - Monitor
 - Relay
 - Event Grid
 - App Service
 - Machine Learning
 - Automation
 - IOT Hub
 - SignalR
 - Batch

Service Endpoint vs Private Link

	Service Endpoint	Private Link
Security	Connects via Public IP	Connects via Private IP
Simplicity	Very simple	More complex
Price	Free	Not free
Supported services	Limited list	Large list, probably will get larger
On-Prem connectivity	Quite complex	Supported

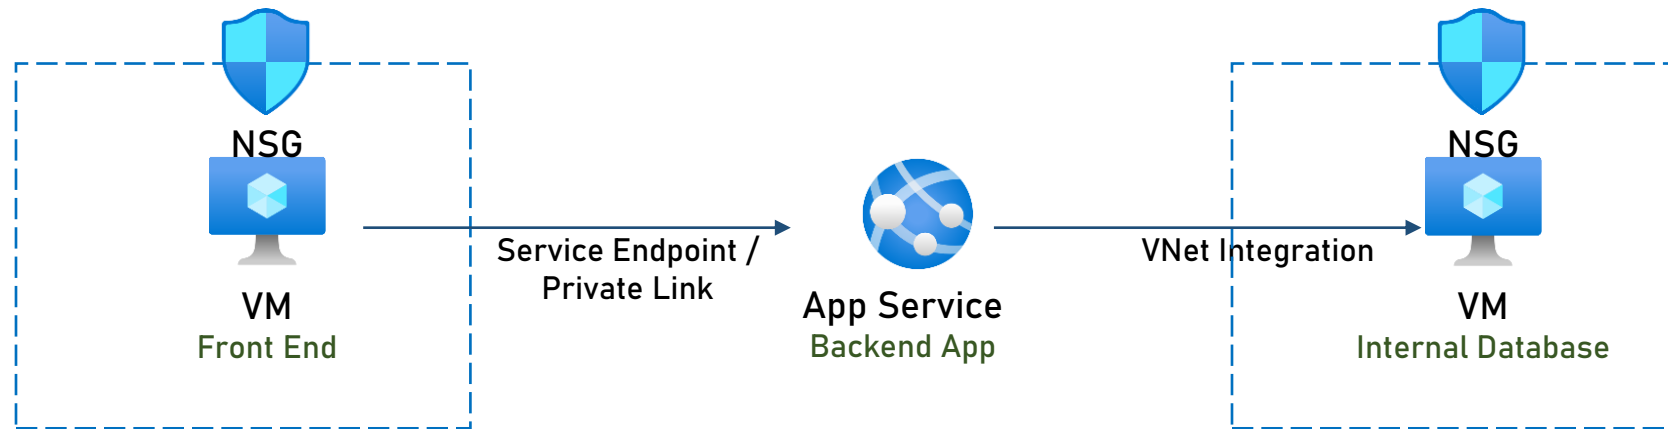
Service Endpoint vs Private Link

- We'll demonstrate Service Endpoint and Private Link later
 - Service Endpoint – on Application Gateway and Azure SQL Server
 - Private Link – on KeyVault

App Service VNet Integration

- Allows access from App Service to resources within VNet
 - So that these resources should not be exposed on the internet
- Extremely useful when App Service needs access to a VM with some internal resources
- Supports same-region VNets. For VNets in other regions – a gateway is required

SE / PL vs VNet Integration



App Service Access Restrictions

- Similar to NSG – but for App Services
- Restricts traffic to the App Service
- By default – all inbound traffic is allowed (in relevant ports)
- Using access restrictions inbound traffic is restricted to the allowed IPs / VNets / Service Tag

App Service Access Restrictions

- Main use cases:
 - Backend App Service that should be accessed from front end
App Service / VM only
 - App Service that sits behind Application Gateway / Load Balancer and shouldn't be accessible directly
 - Open App Service to a specific customer only

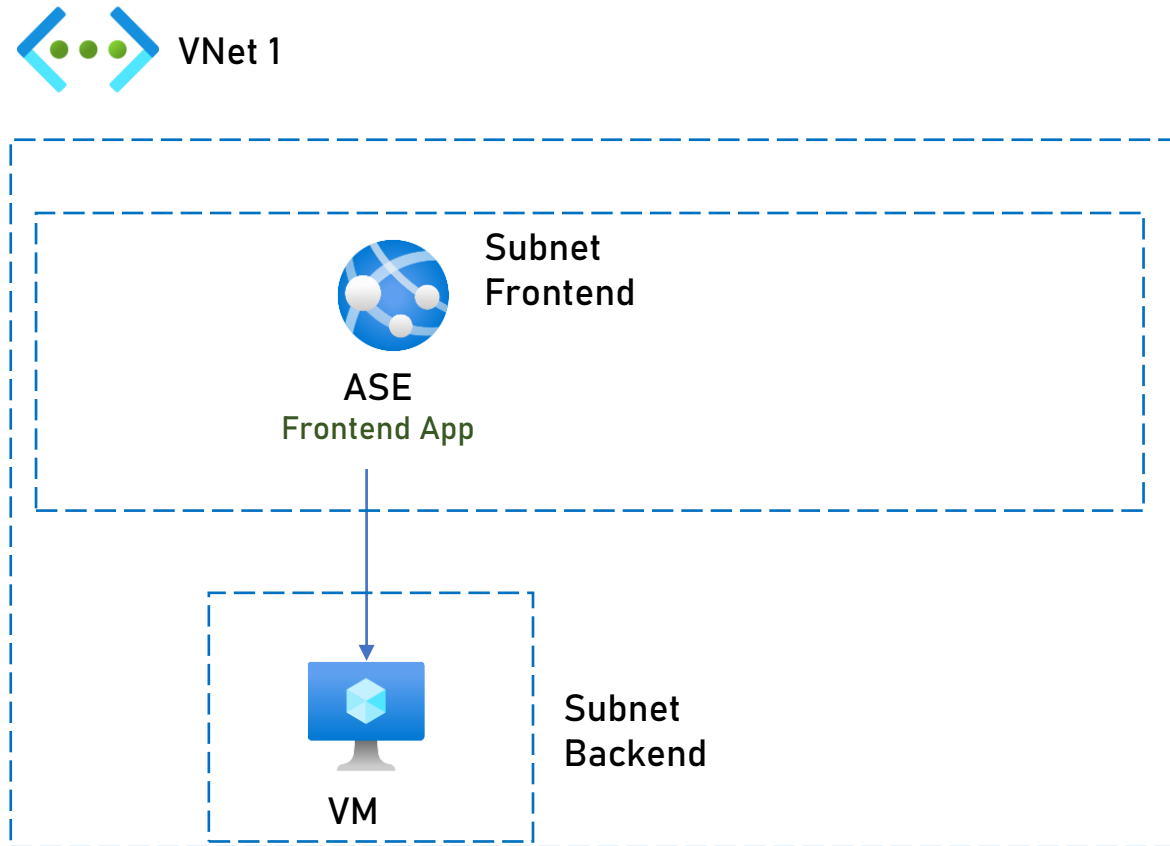
ASE

- App Service Environment
- Special type of app service deployed directly to a dedicated VNet
- VNet can be configured like any other VNet – Subnets, NSGs, etc
- Created on dedicated hardware
- Quite expensive...

ASE

- Major use cases:
 - Elevated security – complete isolation
 - Very high scale requirements

ASE

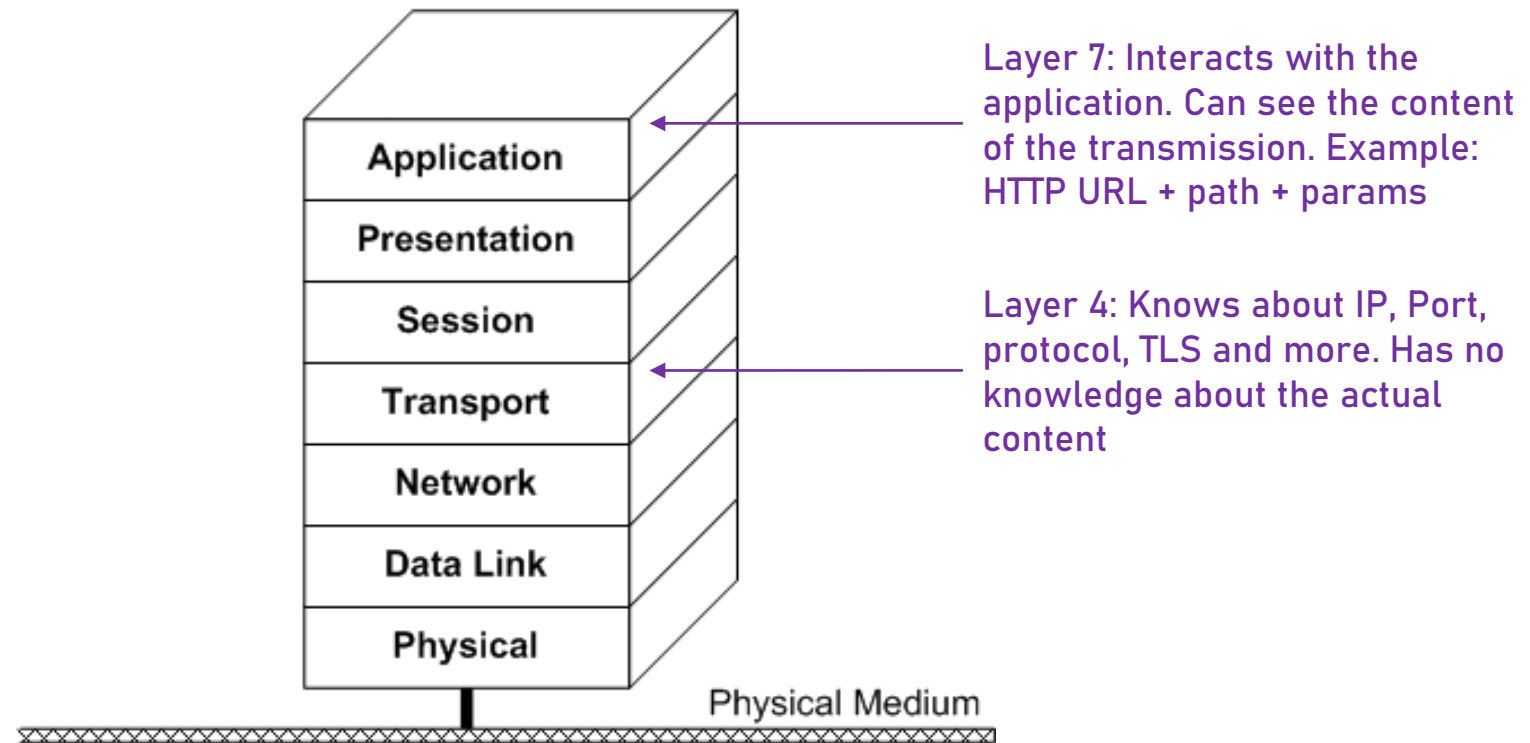


Load Balancer

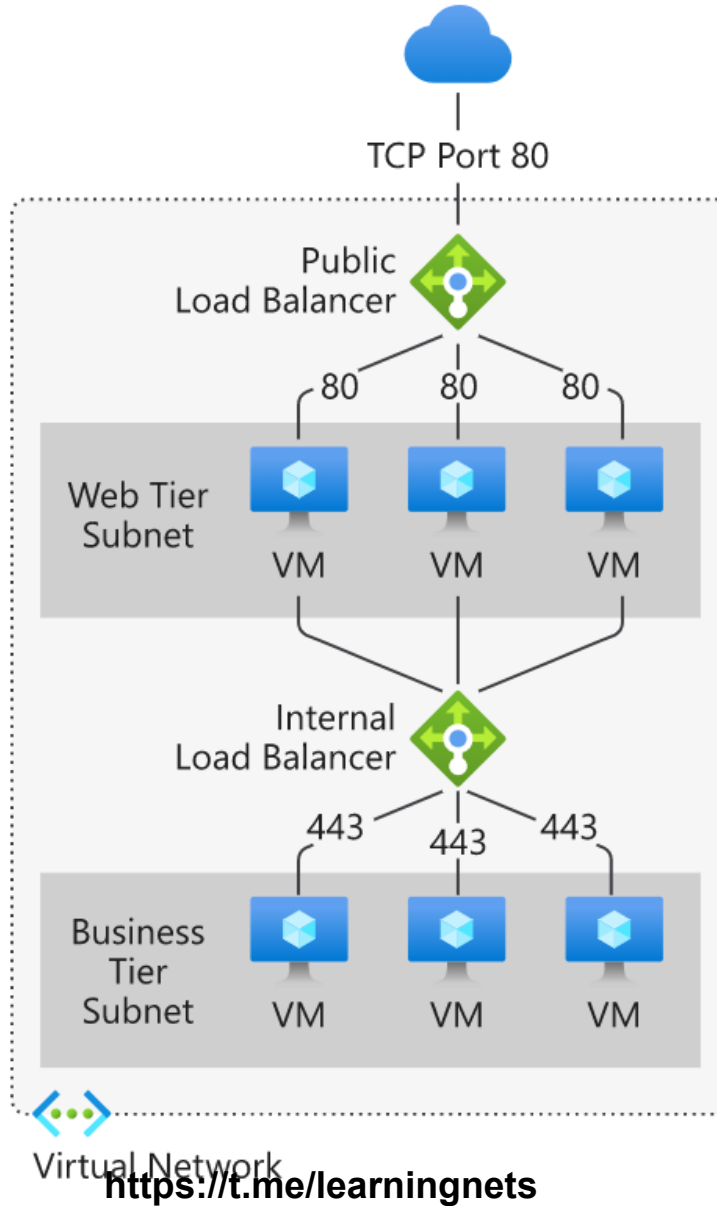
- Azure service that distributes load and checks health of VMs
- When a VM is not healthy – no traffic is directed to it
- Can work with VMs or Scale Set
- Can be public or private
- Operates at layer 4 of the OSI model

7 Layers Model

The OSI Reference Model



Load Balancer

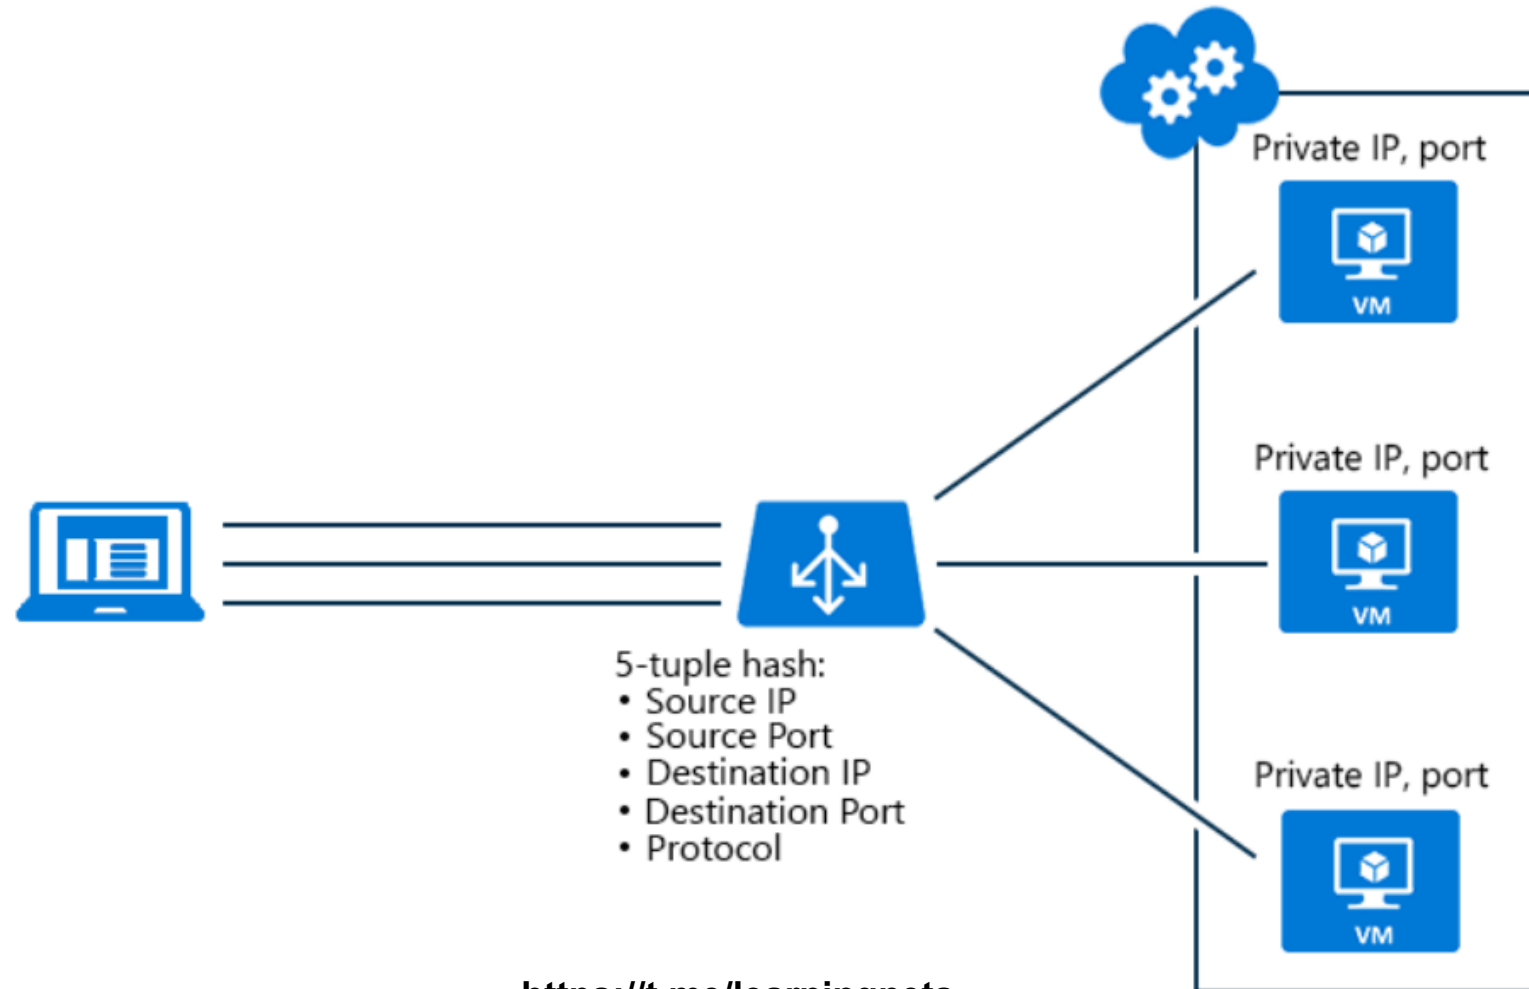


Load Balancer Distribution Algorithm

- Based on 5 tuple hash:
 - Source IP
 - Source port
 - Destination IP
 - Destination port
 - Protocol type
- Same tuples used by NSG



Load Balancer Distribution Algorithm



Load Balancer Types

Basic	Standard

Load Balancer Types

Basic	Standard
No redundancy	Redundant

Load Balancer Types

Basic	Standard
No redundancy	Redundant
Open by default	Secure by default

Load Balancer Types

Basic	Standard
No redundancy	Redundant
Open by default	Secure by default
Up to 300 instances	Up to 1000 instances

Load Balancer Types

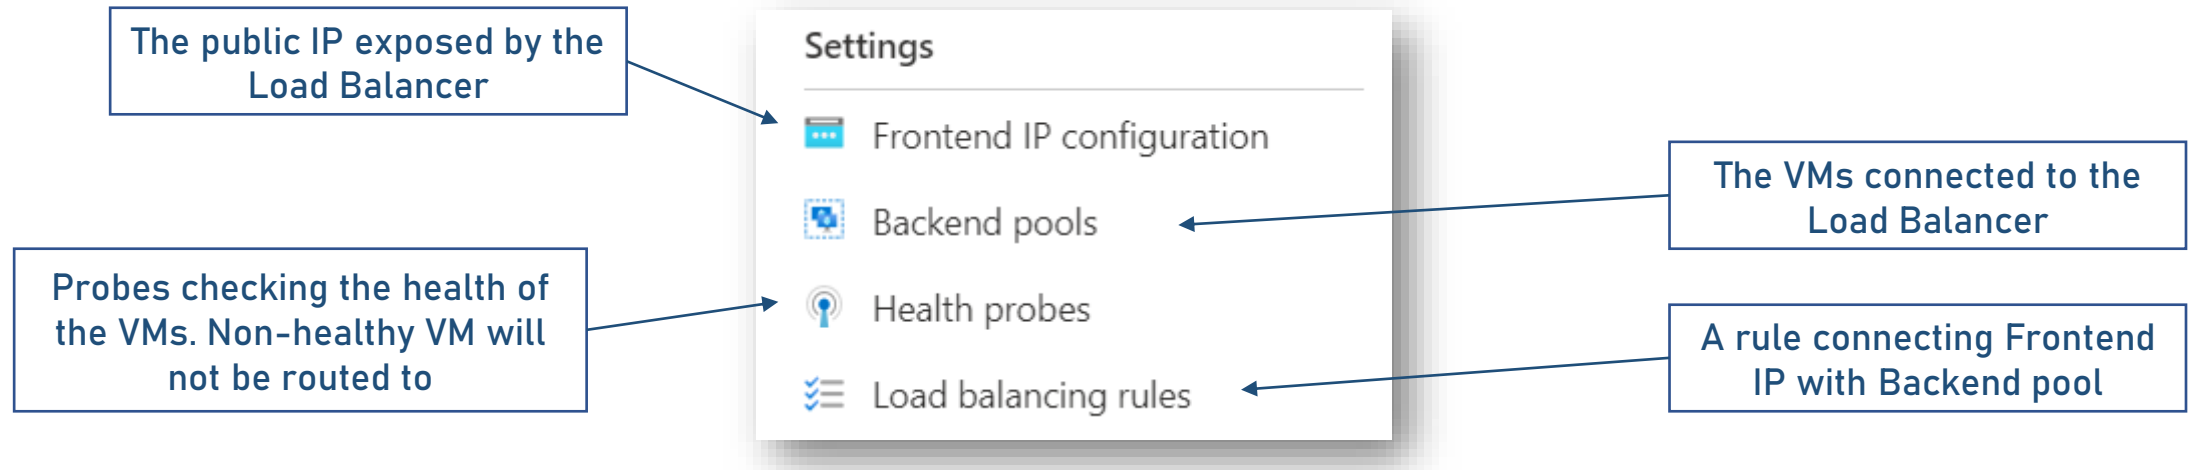
Basic	Standard
No redundancy	Redundant
Open by default	Secure by default
Up to 300 instances	Up to 1000 instances
No SLA	99.99% SLA

Load Balancer Types

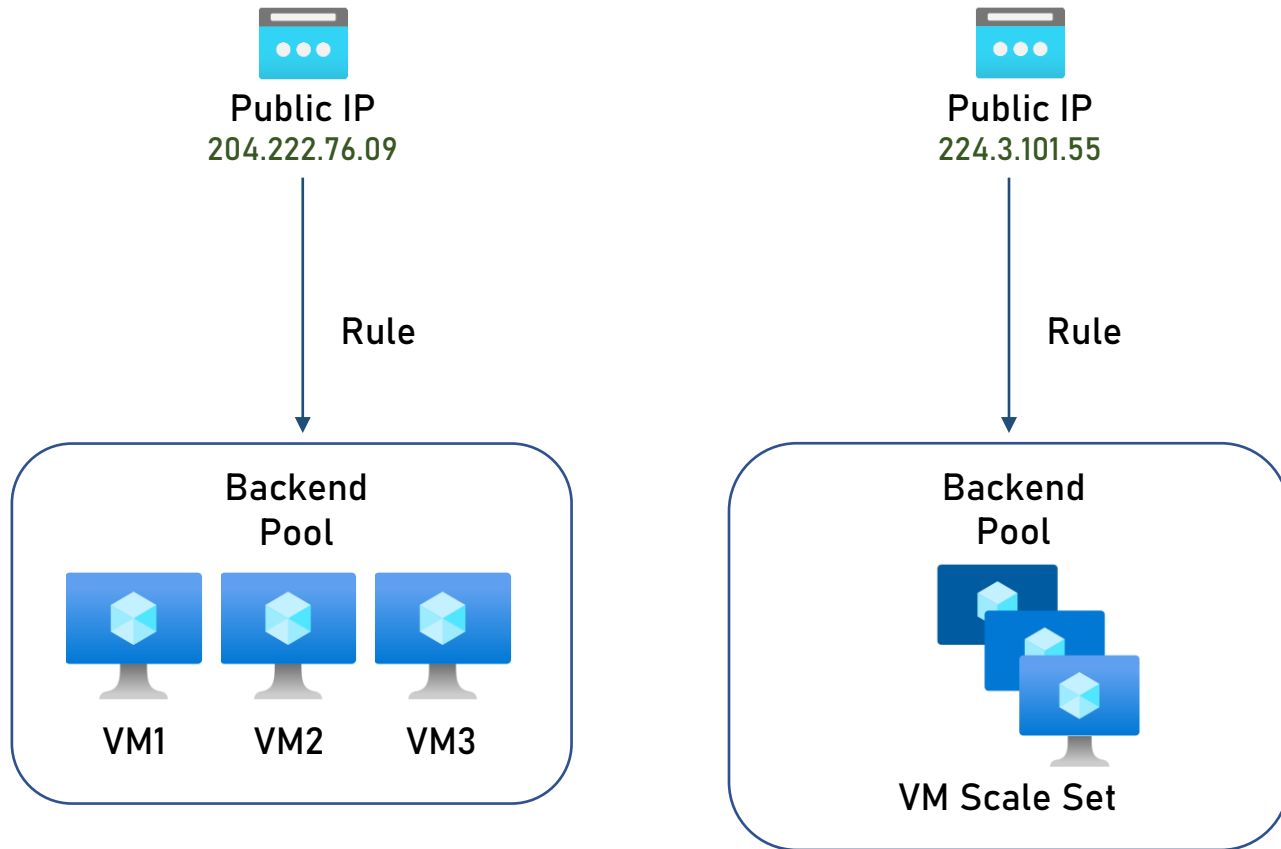
Basic	Standard
No redundancy	Redundant
Open by default	Secure by default
Up to 300 instances	Up to 1000 instances
No SLA	99.99% SLA
Free	Not Free

Configuring Load Balancer

- 4 main configurations:



Example



Settings

- Frontend IP configuration
- Backend pools
- Health probes
- Load balancing rules

Health Probes

- Check the health of the VM
- A non-healthy VM will be marked as Down and will not be routed to
- Run in intervals (usually a few seconds)
- Can run on TCP, HTTP, HTTPS (Standard only)
- Configurable unhealthy threshold – how many times a check should fail for the VM to be marked as Down (default is 2)

Health Probes

Add health probe

mylb

Name *

Protocol ⓘ

Port * ⓘ

Interval * ⓘ

seconds

Unhealthy threshold * ⓘ

consecutive failures

Health Probes

- Run on the VM's host
- No network traffic outside the host
- Originate from the same IP: 168.63.129.16
- Allowed by default in NSG

Inbound port rules Outbound port rules Application security groups Load balancing

Network security group **az-vnet-test2-nsg** (attached to network interface: **az-vnet-test2665**)
Impacts 0 subnets, 1 network interfaces Add inbound port rule

Priority	Name	Port	Protocol	Source	Destination	Action	
300	⚠ RDP	3389	TCP	Any	Any	✔ Allow	...
65000	AllowVnetInBound	Any	Any	VirtualNetwork	VirtualNetwork	✔ Allow	...
65001	AllowAzureLoadBalancerInBound	Any	Any	AzureLoadBalancer	Any	✔ Allow	...
65500	DenyAllInBound	Any	https://t.me/learningnets	Any	Any	✘ Deny	...

When to Use Load Balancer

- Great for internal resources
- Do not use for external resources
 - Especially on Web Apps / Web API / etc.
 - Can't handle HTTP
 - Doesn't route based on path
 - No protection
- For this we have the Application Gateway
 - And demo too...

Application Gateway

- Web traffic load balancer
- Can function as the external endpoint of the web app
- Works with:
 - VMs
 - VM Scale Sets
 - App Services
 - Kubernetes (requires some hacking...)

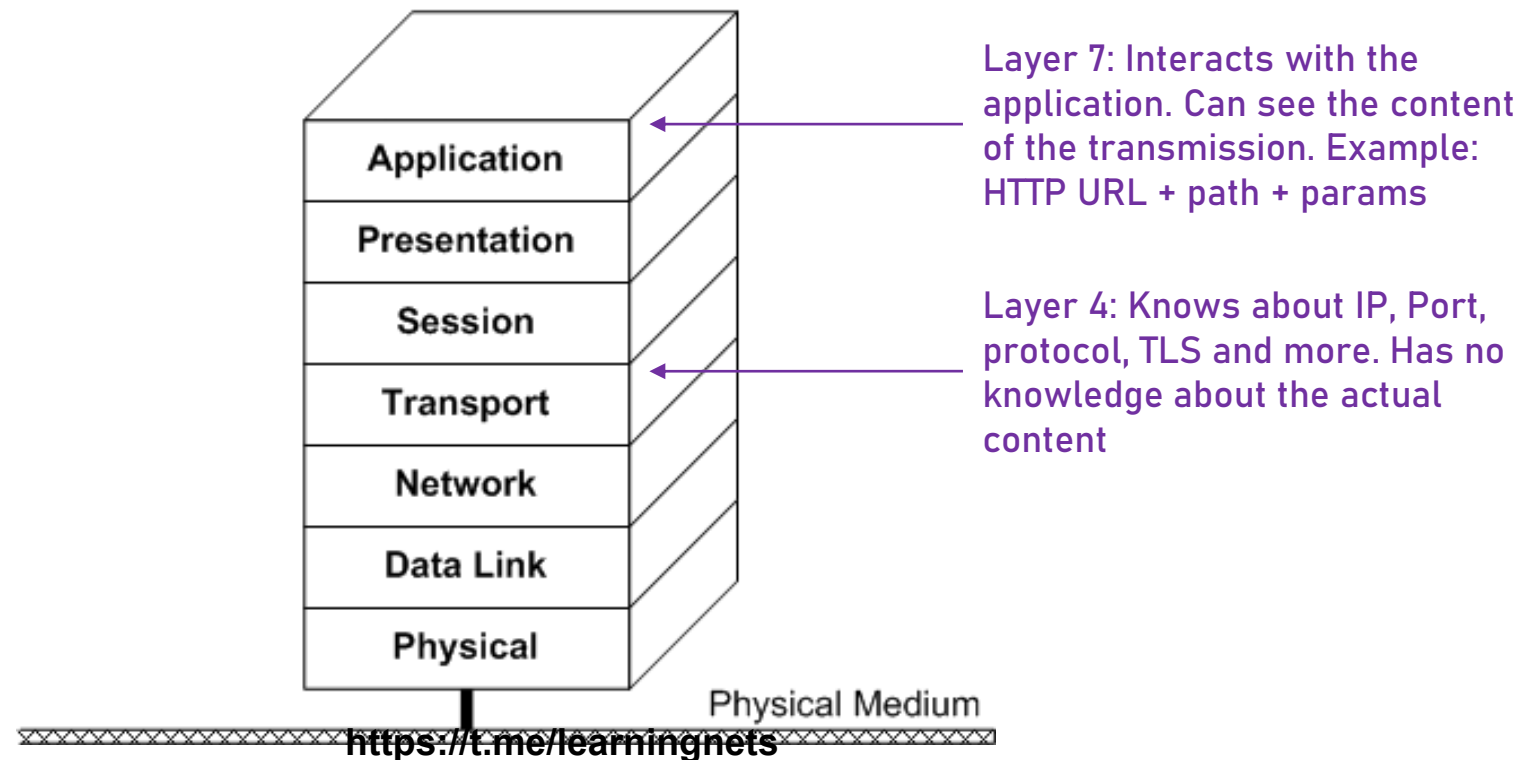
Application Gateway

- Similar to the Load Balancer...
- With additional features:
 - SSL Termination
 - Autoscaling
 - Zone redundancy
 - Session affinity
 - URL based routing
 - WebSocket and HTTP/2 support
 - Custom error pages
 - Header & URL rewrite
 - WAF
 - And more...

Application Gateway

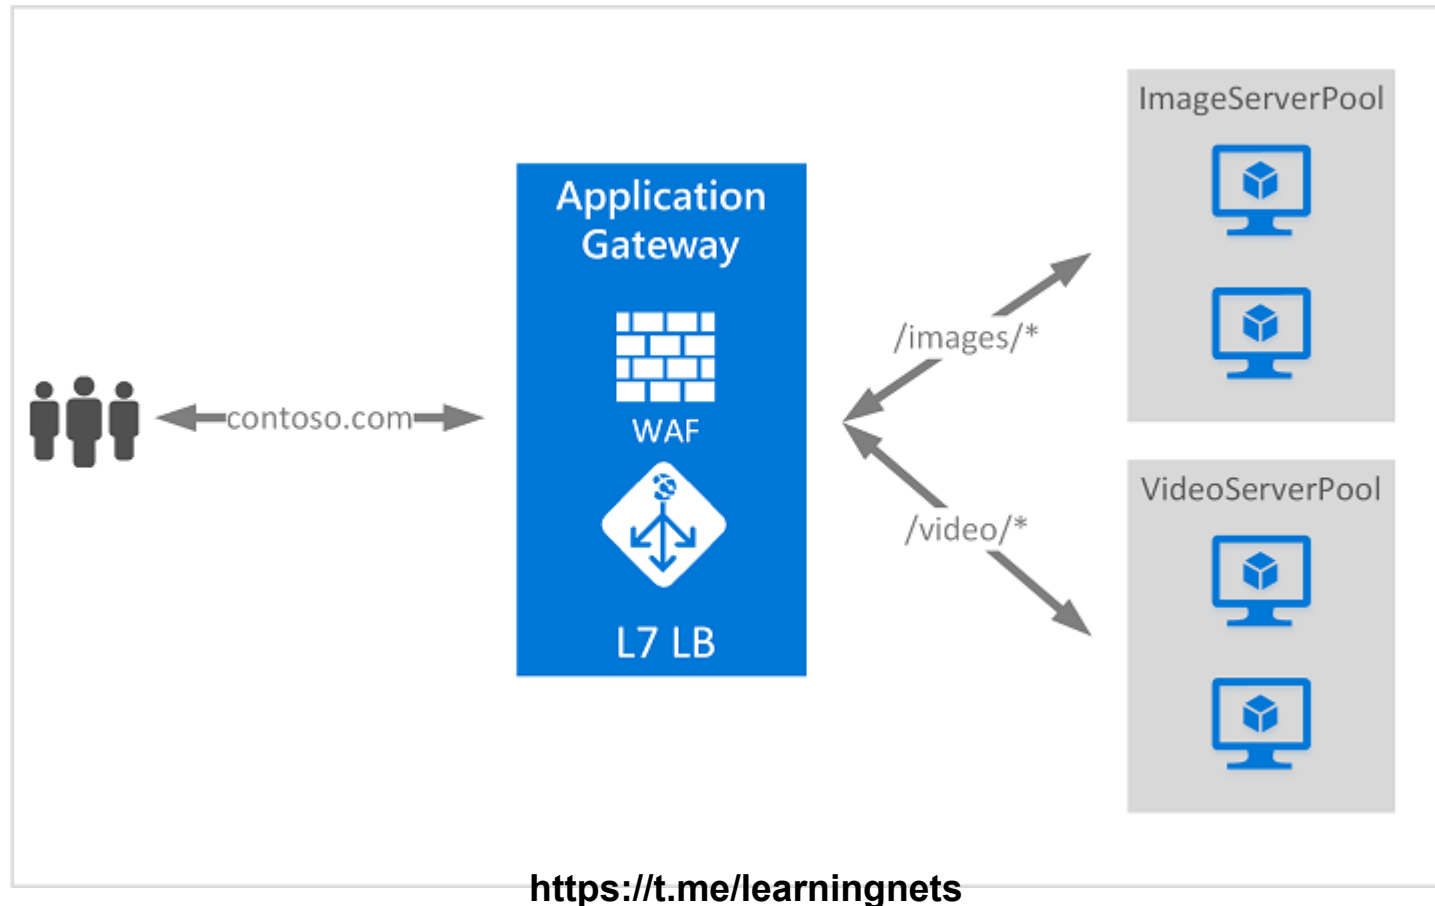
- Operates at layer 7 of the OSI model

The OSI Reference Model



Application Gateway

- Operates at layer 7 of the OSI model



WAF

- Web Application Firewall
- Protects web apps against common attacks
 - ie. Cross-site scripting, SQL injection, etc.
- Protection rules based on OWASP Core Rule Set
- Updates continuously
- Works in Detection or Prevention mode

WAF

- Many organizations have their own WAF deployment
- Usually based on 3rd party products (Palo Alto, Fortinet, Imperva etc.)
- In these cases – there's no need for the WAF in the Application Gateway

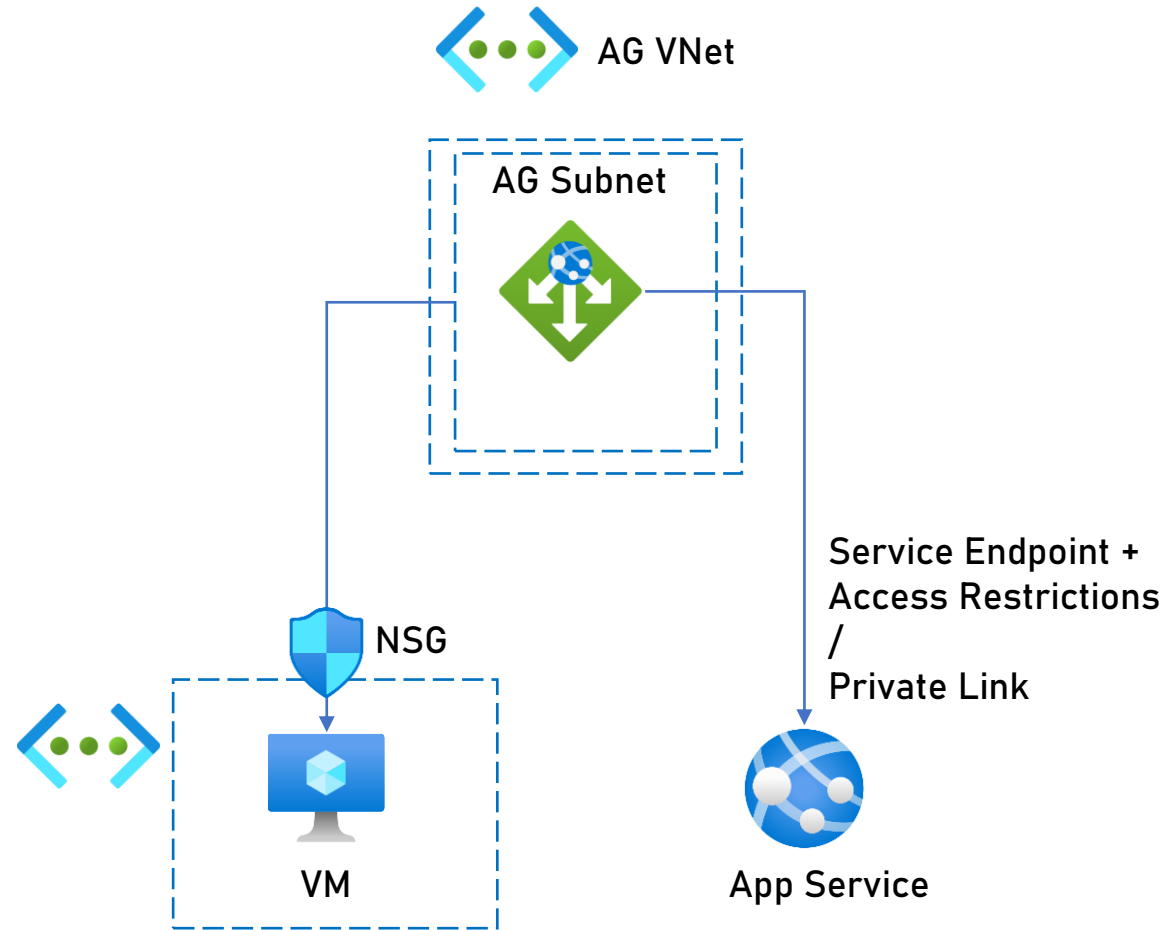
Application Gateway SKUs

- **Standard_V2** – includes all the features mentioned, excluding WAF
- **WAF_V2** – Includes everything (almost double the price...)

Application Gateway Networking

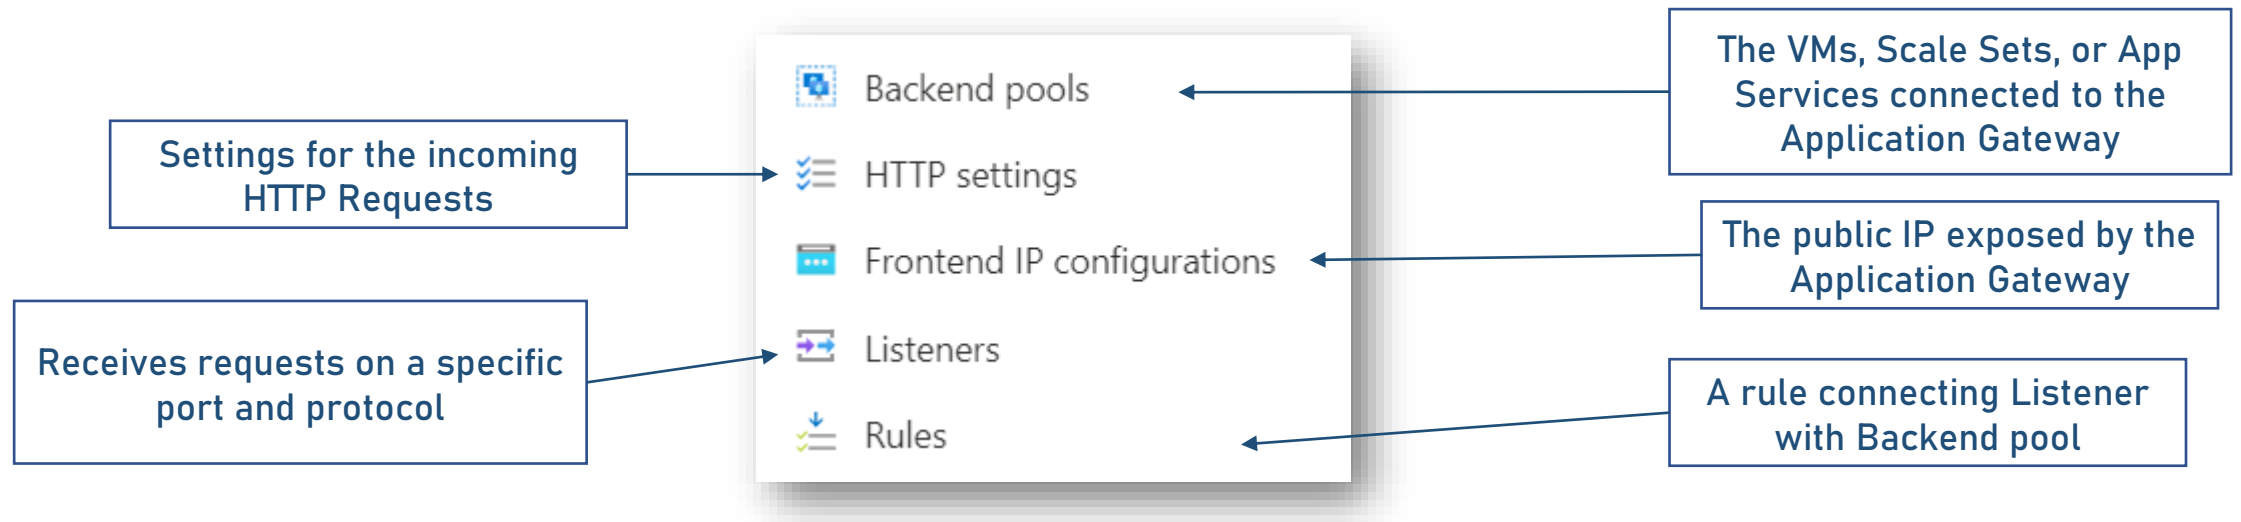
- Application Gateway is placed in its own Subnet
- Often in its own VNet
- Must make sure backend resources are:
 - Accessible from the AG Subnet
 - Not accessible from anywhere else...

Application Gateway Networking



Configuring Application Gateway

- 5 main configurations:



Application Gateway and AKS

- No built-in integration with AKS
- AKS has kind-of gateway (=services)
- There's Application Gateway Ingress Controller (AGIC) that does this
 - In preview mode, quite buggy, not recommended
- Better use 3rd party products

Application Gateway and Functions

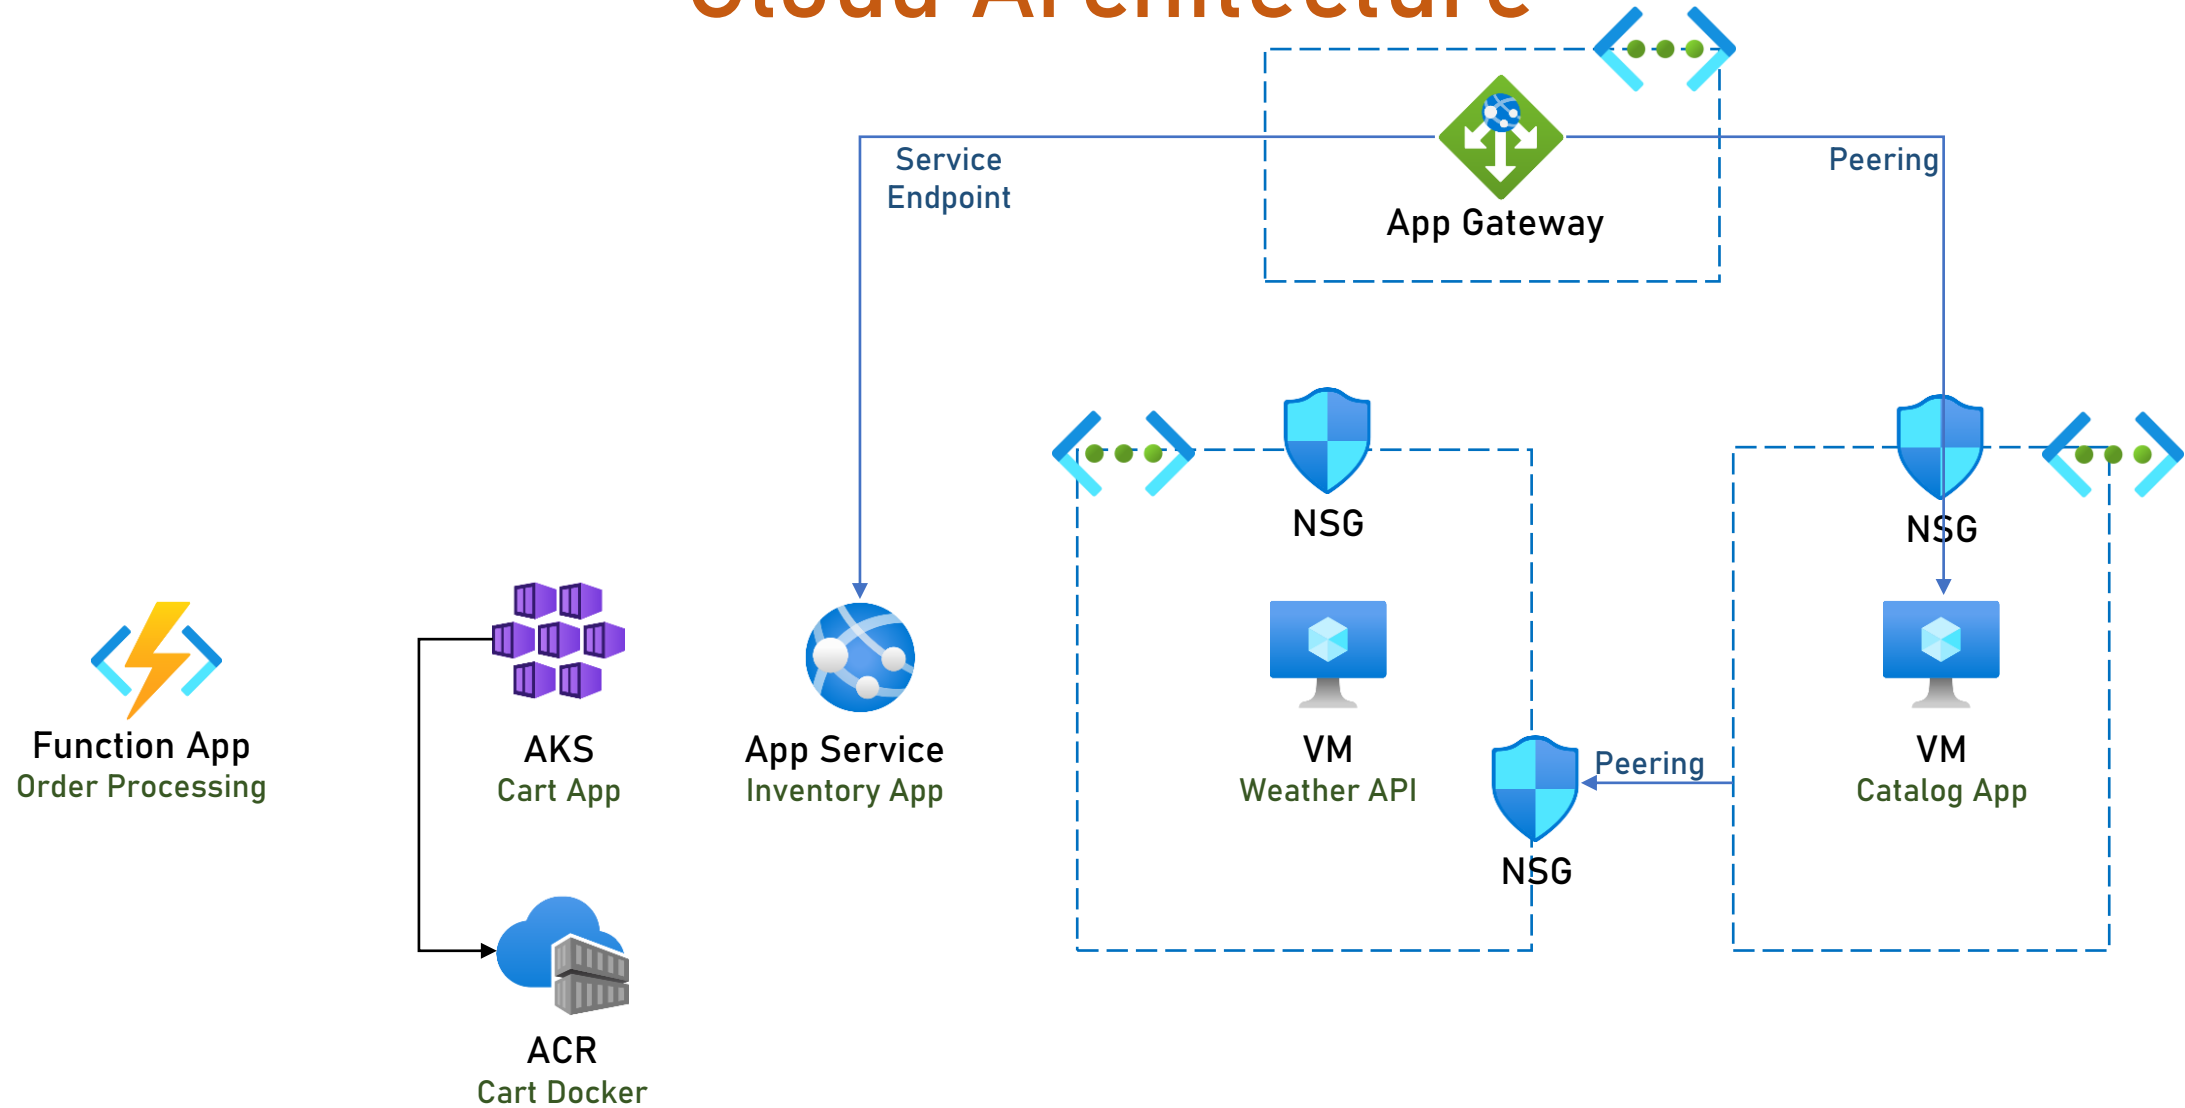
- Functions Apps are basically App Services
- They can be protected by Application Gateway the same way

App Services are

- Configure in Backend pool
- Configure Access Restrictions
- Our function won't be accessible from the web so no demo...

ReadIt!

Cloud Architecture



Affinity

Add HTTP setting

HTTP settings name

inventory-settings

Backend protocol

HTTP HTTPS

Backend port *

80

Additional settings

Cookie-based affinity ⓘ

Enable Disable

Connection draining ⓘ

Enable Disable

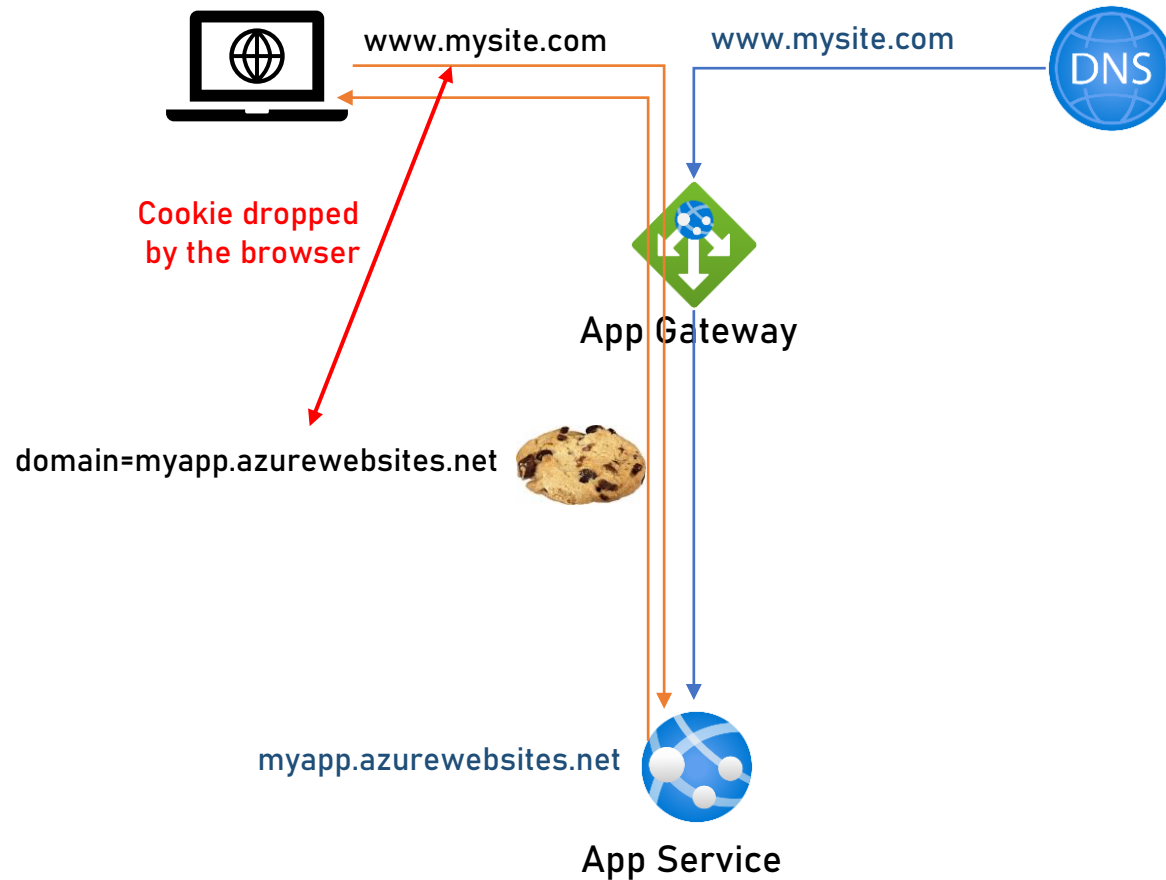
Affinity

- Makes sure user will always be directed to the same instance (VM / App Service) it began with
- Should be avoided when possible
- Usually required in Stateful apps
- Usually a sign of bad design
- Always try to design Stateless app

Application Gateway and Cookies

Super Advanced

Application Gateway and Cookies



Application Gateway and Cookies

- The solution:
 - Set custom domain for the App Service to be the same one of the Application Gateway

Application Gateway and Cookies

App Service

Search (Ctrl+ /)

- Deployment slots
- Deployment Center
- Deployment Center (Preview)

Settings

- Configuration
- Authentication / Authorization
- Application Insights
- Identity
- Backups
- Custom domains**
- TLS/SSL settings

Custom domains

Refresh Troubleshoot FAQs

Custom Domains

Configure and manage custom domains assigned to your app [Learn more](#)

IP address: ⓘ
23.97.208.18

Custom Domain Verification ID: ⓘ
99DC541397E5FF77F04C286775933BD298BF14ECB845028CC6F834E02ED62914

HTTPS Only: ⓘ
 Off

+ Add custom domain

Status Filter
All (1) Not Secure (0) Secure (1)

<https://t.me/learningnets>

Secure Network Design

