Azure Networking - Detailed Reference

Azure Complete Reference — Detailed Edition

Part 1: Networking (Topics 1-20)

Topic 1: Virtual Network (VNet)

A Virtual Network is a logically isolated network in Azure. Think of it as your own private network in the cloud — similar to a physical network in an on-prem data center, but virtual.

Key Properties

1. Region-bound - A VNet lives in one Azure region (e.g., East US, West Europe) - You cannot create a VNet that spans multiple regions - To connect VNets across regions, you use VNet peering

2. Subscription-bound - A VNet belongs to one subscription - But you can have multiple VNets in one subscription

3. Address Space (CIDR) - When you create a VNet, you define its IP address range using CIDR notation - Example: 10.0.0.0/16 gives you 65,536 addresses (10.0.0.0 – 10.0.255.255) - You can add multiple address spaces to one VNet (e.g., 10.0.0.0/16 + 192.168.0.0/16)

CIDR refresher:

CIDR	Total IPs	Usable IPs (Azure reserves 5)
/29	8	3
/28	16	11
/27	32	27
/26	64	59
/25	128	123
/24	256	251
/22	1,024	1,019
/20	4,096	4,091
/16	65,536	65,531
/8	16,777,216	16,777,211

Azure reserves 5 IPs in every subnet: - First address: network address (e.g., 10.0.1.0) - Second address: default gateway (e.g., 10.0.1.1) - Third and fourth: Azure maps to DNS and internal services - Last address: broadcast (e.g., 10.0.1.255)

4. Cannot overlap - If you peer two VNets or connect to on-prem, their address spaces must not overlap - This is why you plan IP addressing upfront — changing it later is painful

Subnets

A VNet is divided into subnets. Each subnet is a segment of the VNet's address space.

Why subnets? - Isolation — put different tiers in different subnets (web, app, database) - Security — apply NSGs at subnet level to control traffic between subnets - Organization — group related resources together

Subnet rules: - A subnet must be within the VNet's address space - Subnets within a VNet cannot overlap - Minimum subnet size: /29 (8 IPs, 3 usable — too small for anything useful) - Practical minimum: /27 (32 IPs, 27 usable) or /24 (256 IPs, 251 usable) - By default, all subnets in a VNet can communicate with each other automatically

Example: VNet: 10.0.0.0/16 - Frontend subnet: 10.0.1.0/24 (web servers) - Backend subnet: 10.0.2.0/24 (app servers) - Database subnet: 10.0.3.0/24 (database servers) - GatewaySubnet: 10.0.0.0/27 (for VPN/ExpressRoute gateway — must be named exactly GatewaySubnet)

Special Subnets

Subnet Name	Purpose	Special Rule
GatewaySubnet	VPN Gateway / ExpressRoute Gateway	Must be named exactly this. Min /27. No other resources allowed in it.
AzureBastionSubnet	Azure Bastion	Must be named exactly this. Min /26. No other resources allowed.
Delegated subnet	For specific services (e.g., App Service VNet integration)	You delegate the subnet to a specific service.

VNet Creation — What You Configure

Setting	What to Choose
Name	Follow naming convention (e.g., vnet-prod-eastus-001)
Region	Where your resources are
Address space	Plan for growth — do not paint yourself into a corner
Subnets	Define name + CIDR for each
DDoS Protection	Basic (free) or Standard (paid, ~$2,944/month)
Firewall	Skip at VNet creation, add later if needed

Default Behavior (No NSGs, No Custom Routes)

Traffic	Default Behavior
VM to VM in same subnet	Allowed
VM to VM in different subnet (same VNet)	Allowed
VM to internet	Allowed (outbound)
Internet to VM	Blocked (unless VM has public IP)
VM to on-prem	Not possible (no connectivity)
VM to other VNet	Not possible (no peering)

This is important — by default, everything inside the VNet can talk to everything else inside the VNet. If you want to restrict subnet-to-subnet traffic, you need NSGs (coming in Topic 2).

Common Mistakes

1. Using 10.0.0.0/16 everywhere — then when you need to peer VNets or connect on-prem, they overlap. Plan unique ranges.
2. Making subnets too small — /28 gives 16 IPs, only 11 usable. You will run out quickly.
3. Not reserving subnets for gateways and Bastion — you cannot create a VPN Gateway later if there is no GatewaySubnet.
4. Changing address space after resources are deployed — possible but disruptive. Plan upfront.

Topic 2: Network Security Groups (NSG)

An NSG is a firewall at the subnet or network interface level. It contains rules that allow or deny inbound and outbound traffic to/from Azure resources.

Think of it as a bouncer at the door — it checks every packet against its rules and decides whether to let it through or block it.

NSG Structure

An NSG contains two lists of rules: - Inbound rules — control traffic coming INTO the resource - Outbound rules — control traffic going OUT of the resource

Rule Properties

Every rule has:

Property	What it means	Example
Priority	Number 100-4096, lower = evaluated first	100 is checked before 200
Direction	Inbound or Outbound	Inbound
Action	Allow or Deny	Allow
Source	Where traffic is coming from	IP, CIDR, Service Tag, ASG
Source Port	Source port range	* (any), 80, 1024-65535

Property	What it means	Example
Destination	Where traffic is going to	IP, CIDR, Service Tag, ASG
Destination Port	Target port range	80, 443, 22
Protocol	TCP, UDP, ICMP, or Any	TCP
Name	Unique name for the rule	Allow-HTTP

How Rules Are Evaluated

Rules are processed lowest priority number first. When a match is found, processing stops — no further rules are checked.

Example: Priority 100: Allow TCP 443 from Any — MATCH — traffic allowed, stop Priority 200: Deny Any from Any — never checked

This is critical — rule order matters. A broad deny rule with low priority will block everything, even if you have specific allow rules with higher priority numbers.

Default Rules

Every NSG comes with 4 default rules that you cannot delete:

Default Inbound:

Priority	Action	Source	Port	Description
65000	Allow	VNet	Any	Allow traffic within the VNet
65001	Allow	AzureLoadBalancer	Any	Allow Azure health probes
65500	Deny	Any	Any	Deny all other inbound

Default Outbound:

Priority	Action	Destination	Port	Description
65000	Allow	VNet	Any	Allow outbound within VNet
65001	Allow	Internet	Any	Allow outbound to internet
65500	Deny	Any	Any	Deny all other outbound

What this means by default: - VMs in the same VNet can talk to each other - VMs can reach the internet - Internet cannot reach VMs (unless VM has public IP + you add allow rules) - Azure load balancer health probes work

Where NSGs Attach

An NSG can be attached to:

1. A Subnet (most common) - Rules apply to ALL resources in that subnet - Best practice for tier isolation (e.g., block database subnet from receiving traffic except from app subnet)

2. A Network Interface (NIC) - Rules apply to only that specific VM - Use for exceptions — e.g., one VM needs different rules than others in the same subnet

3. Both - Yes, you can attach NSGs to both subnet AND NIC - Effective rules = subnet NSG + NIC NSG combined - Evaluation: inbound = NIC NSG checked first, then subnet NSG. Outbound = subnet NSG first, then NIC NSG. - If either NSG denies, traffic is denied

Source and Destination Options

Instead of typing individual IPs, you can use:

Option	Example	When to Use
IP Address	203.45.67.89, 10.0.1.0/24	Specific IP or range
Service Tag	Internet, VirtualNetwork, AzureLoadBalancer, Storage, Sql	Predefined groups of IPs by service
Application Security Group (ASG)	asg-web, asg-db	Group VMs logically, reference in rules
Any	*	All sources/destinations

Service Tags — Important Ones

Tag	What it represents
VirtualNetwork	All IPs in your VNet address space (including peered VNets)
Internet	Everything outside Azure that is reachable publicly
AzureLoadBalancer	Azure infrastructure load balancer (health probes)
Storage	Azure Storage service public IPs (for private endpoint access control)
Sql	Azure SQL Database public IPs
AzureCloud	All Azure datacenter IPs (broad — use sparingly)

Application Security Groups (ASGs)

ASGs let you group VMs by role and reference the group in NSG rules instead of individual IPs.

Example without ASGs: Allow port 443 from 10.0.1.4 (web-vm1) to 10.0.2.4 (app-vm1) Allow port 443 from 10.0.1.5 (web-vm2) to 10.0.2.4 (app-vm1) Allow port 443 from 10.0.1.4 (web-vm1) to 10.0.2.5 (app-vm2) Allow port 443 from 10.0.1.5 (web-vm2) to 10.0.2.5 (app-vm2) 4 rules and growing every time you add a VM.

With ASGs: Allow port 443 from ASG-Web to ASG-App 1 rule. Add 50 VMs to each ASG, the rule still works.

How to use ASGs: 1. Create ASG (e.g., asg-web, asg-app, asg-db) 2. Assign VM NICs to the appropriate ASG 3. Reference ASGs in NSG rules instead of IPs

Practical Example — 3-Tier App

VNet: 10.0.0.0/16

Subnet	CIDR	Resources
frontend	10.0.1.0/24	Web servers (ASG: asg-web)
backend	10.0.2.0/24	App servers (ASG: asg-app)

Subnet	CIDR	Resources
database	10.0.3.0/24	DB servers (ASG: asg-db)

NSG: nsg-frontend (attached to frontend subnet)

Priority	Direction	Source	Dest	Port	Action	Why
100	Inbound	Any	Any	80	Allow	HTTP from internet
110	Inbound	Any	Any	443	Allow	HTTPS from internet
200	Inbound	asg-app	Any	443	Allow	App tier responses (if needed)
300	Inbound	Your-IP/32	Any	22	Allow	SSH management
1000	Outbound	Any	asg-app	8080	Allow	Traffic to app tier
65000	Inbound	VNet	Any	Any	Allow	Default: VNet internal
65500	Inbound	Any	Any	Any	Deny	Default: deny all else

NSG: nsg-backend (attached to backend subnet)

Priority	Direction	Source	Dest	Port	Action	Why
100	Inbound	asg-web	Any	8080	Allow	Only from web tier
200	Inbound	Your-IP/32	Any	22	Allow	SSH management
65000	Inbound	VNet	Any	Any	Allow	Default: VNet internal
65500	Inbound	Any	Any	Any	Deny	Default: deny all else

NSG: nsg-database (attached to database subnet)

Priority	Direction	Source	Dest	Port	Action	Why
100	Inbound	asg-app	Any	3306	Allow	Only from app tier (MySQL)
200	Inbound	Your-IP/32	Any	22	Allow	SSH management
65500	Inbound	Any	Any	Any	Deny	Default: deny all else

Notice the pattern: Each tier only accepts traffic from the tier directly above it. Web accepts from internet. App accepts from web. DB accepts from app. Nothing else gets in.

NSG Limits

Limit	Value
Max rules per NSG	300 (including defaults)
Max NSGs per subnet	1
Max NSGs per NIC	1
Priority range	100-4096

Common Mistakes

1. Not restricting SSH/RDP — leaving port 22/3389 open to Any. Always restrict to your IP or use Bastion.
2. Forgetting outbound rules — by default outbound to internet is allowed. If you need to restrict outbound (e.g., prevent data exfiltration), add explicit deny rules.
3. Overlapping subnet and NIC NSGs — can cause confusion when troubleshooting. Pick one strategy: use subnet-level NSGs as the primary, NIC-level only for exceptions.
4. Not using ASGs — hardcoding IPs in NSG rules does not scale. ASGs make rules readable and maintainable.
5. Not leaving priority gaps — if you number rules 100, 101, 102, you cannot insert a rule between them later. Use gaps: 100, 110, 120, etc.

Topic 3: Azure Load Balancer

Distributes incoming network traffic across multiple backend VMs. Ensures no single VM is overwhelmed.

Types

Type	SKU	Layer	When to Use
Public Load Balancer	Basic/Standard	Layer 4 (TCP/UDP)	Distribute internet traffic to VMs
Internal Load Balancer	Basic/Standard	Layer 4 (TCP/UDP)	Distribute traffic internally between tiers

Basic vs Standard SKU

Feature	Basic	Standard
Backend pool size	100	1,000
Health probes	Yes	Yes (more options)
Availability Zones	No	Yes (zone-redundant)
Outbound SNAT	Yes	Yes (more control)
HTTPS probing	No	Yes
SLA	None	99.99%
Cost	Free	~$18/month
Security	Open by default	Closed by default (NSG required)
Diagnostic logs	Limited	Full

Always use Standard SKU for production. Basic is being deprecated.

Key Components

Component	What it Does
Frontend IP	Public or private IP that receives traffic
Backend Pool	Set of VMs or NICs that receive distributed traffic
Health Probes	Check if backend VMs are healthy. Unhealthy VMs are removed from rotation
Load Balancing Rules	

Component	What it Does
	Define how traffic is distributed (port, protocol, algorithm)
Inbound NAT Rules	Forward specific port traffic to a specific VM (e.g., SSH to VM1 on port 50001)

Health Probes

Property	Description
Protocol	TCP, HTTP, or HTTPS
Port	Port to probe (e.g., 80)
Interval	How often to probe (default 5 sec)
Unhealthy threshold	Consecutive failures before marking unhealthy (default 2)
Request path	For HTTP/HTTPS probes (e.g., /health)

If a VM fails the health probe, Load Balancer stops sending traffic to it. When it passes again, traffic resumes automatically.

Custom probe path: Instead of just checking “/”, check “/health” endpoint that verifies the app is truly functional (database connected, dependencies available).

Load Distribution Methods

Method	How it Works
Round-robin (default)	Distribute evenly across all healthy VMs
Source IP affinity (session persistence)	Same client IP always goes to same backend VM

Session Persistence is useful for apps that store session state locally on the VM. Better approach: use external session store (Redis, database) and avoid session persistence.

Internal Load Balancer Use Case

Internet -> Public LB -> Web VMs -> Internal LB -> App VMs -> Database VMs

The internal LB distributes traffic between web tier and app tier. No public IP needed — purely internal.

HA Ports

- Load balance ALL ports (not just one specific port)
- Use case: Network Virtual Appliances (NVAs) — you do not know which ports traffic will use, so load balance everything
- Only available on Internal Standard Load Balancer

Azure Complete Reference — Detailed Edition

Part 2: Networking continued (Topics 4-10)

Topic 4: Azure Application Gateway

Application-level (Layer 7) load balancer with additional features like WAF, SSL termination, URL-based routing.

When to use Application Gateway instead of Load Balancer

Need	Use
Layer 4 (TCP/UDP) distribution only	Azure Load Balancer
Layer 7 (HTTP/HTTPS) routing	Application Gateway
SSL termination (offload HTTPS)	Application Gateway
URL-based routing (/api to one pool, /app to another)	Application Gateway
Cookie-based session affinity	Application Gateway
Web Application Firewall (WAF)	Application Gateway
Path-based routing	Application Gateway

SKUs

SKU	When to Use
Standard_v2	Auto-scale, zone-redundant, no WAF
WAF_v2	Same as Standard_v2 + WAF (OWASP protection)

Key Components

Component	What it Does
Frontend IP	Public and/or private IP for incoming traffic
Listeners	Listen on specific port/protocol/hostname. Can have multiple.
Routing Rules	Map listener to backend pool. Can include path-based rules.
Backend Pools	Target VMs, App Service, IP addresses
HTTP Settings	Port, protocol, cookie affinity, connection draining, health probe
Health Probes	Check backend health (HTTP/HTTPS, custom path)
SSL Certificate	For SSL termination at the gateway

SSL Termination

- Client sends HTTPS to App Gateway
- App Gateway decrypts (terminates SSL) using its certificate
- App Gateway sends HTTP or HTTPS to backend VMs
- Benefit: backend VMs do not need to handle SSL encryption/decryption = less CPU load
- You can also re-encrypt: App Gateway decrypts then re-encrypts then sends HTTPS to backend (end-to-end SSL)

URL-Based Routing

Listener on port 443: /app/* goes to Backend Pool: web-vm-pool /api/* goes to Backend Pool: api-vm-pool /images/* goes to Backend Pool: storage-pool

Multi-Site Hosting

- One App Gateway can serve multiple websites
- Listener 1: www.shop.com to shop-pool
- Listener 2: www.blog.com to blog-pool
- Saves cost vs deploying separate gateways

WAF Modes

Mode	Behavior
Detection	Logs violations, lets traffic through
Prevention	Blocks violations (returns 403), logs them

WAF Rule Sets

Rule Set	Covers
OWASP 3.2	Top 10 web vulnerabilities (SQL injection, XSS, etc.)
Bot Protection	Block bad bots, allow good bots
Custom Rules	IP-based blocking, rate limiting, geo-filtering

Typical WAF Implementation Steps

1. Deploy WAF in Detection mode first
2. Monitor logs for 1-2 weeks — see what gets flagged
3. Tune rules — exclude false positives (some apps trigger OWASP rules with legitimate traffic)
4. Switch to Prevention mode
5. Set up alerts for blocked requests
6. Review WAF logs weekly

App Gateway vs Load Balancer Decision

Scenario	Choose
Simple TCP/UDP distribution	Load Balancer
HTTP/HTTPS web traffic	Application Gateway
Need WAF	Application Gateway
Need SSL termination	Application Gateway
Multiple sites on one IP	Application Gateway
Non-HTTP traffic (database, custom protocol)	Load Balancer
Internal tier-to-tier (non-HTTP)	Internal Load Balancer

Topic 5: Azure Firewall

A fully managed, cloud-native firewall service. Inspects all outbound and inbound traffic at the network level.

When to use Azure Firewall

- You need centralized network-level firewall for all VNets
- You need to inspect outbound traffic (what VMs are connecting to on the internet)
- You need DNAT (inbound NAT from internet to VMs)
- You need FQDN-based filtering (allow VMs to reach *.microsoft.com but nothing else)

Azure Firewall vs NSG

Feature	NSG	Azure Firewall
Layer	Layer 3/4	Layer 3/4/7
Filtering	IP, port, protocol	IP, port, protocol, FQDN, application
Stateful	No (stateless)	Yes (stateful inspection)
FQDN filtering	No	Yes (allow *.windowsupdate.com)
Threat intelligence	No	Yes (block known malicious IPs/domains)
DNAT	No	Yes (port forwarding from internet)
Cost	Free	~$500/month + per-GB processing
Where	Subnet or NIC	Dedicated subnet (AzureFirewallSubnet)
Centralization	Per subnet	Centralized for all VNets

Both are used together. NSG = first line per subnet. Azure Firewall = centralized deep inspection.

SKUs

SKU	Features	Cost
Standard	FQDN filtering, DNAT, threat intel	~$500/month
Premium	Standard + TLS inspection, IDPS, URL filtering, web categories	~$900/month
Basic	Simplified, for small businesses	~$150/month

Premium Features

- TLS Inspection: decrypt traffic, inspect payload, re-encrypt. Detects threats inside encrypted traffic.
- IDPS: Intrusion Detection and Prevention System. Signature-based + anomaly-based.
- URL Filtering: block specific URLs, not just FQDNs (e.g., allow facebook.com but block facebook.com/games)
- Web Categories: block entire categories (gambling, social media, malware sites)

Deployment

- Requires a dedicated subnet named AzureFirewallSubnet (min /26)
- Typically deployed in the hub VNet of a hub-spoke topology
- All spoke VNets route traffic through the firewall using UDRs (User Defined Routes)

Hub-Spoke with Firewall

Spoke-1 -> VNet Peering -> Hub (Azure Firewall) -> Internet Spoke-2 -> VNet Peering -> Hub (Azure Firewall) -> Internet Spoke-3 -> VNet Peering -> Hub (Azure Firewall) -> On-prem via ExpressRoute

Every spoke outbound internet traffic goes through the firewall. Firewall inspects and allows/denies.

Firewall Rules

Rule Type	Collection	What it Filters
NAT rules	NAT rule collection	DNAT — translate inbound traffic to internal IPs
Network rules	Network rule collection	IP/port/protocol — allow/deny any network traffic
Application rules	Application rule collection	FQDN/URL — allow/deny outbound to specific domains

Rule evaluation order: NAT rules -> Network rules -> Application rules. First match wins within each collection.

Example Rules

Type	Source	Destination	Port	Action
NAT	Any	Frontend IP:80	80	DNAT to 10.0.1.4:80
Network	10.0.2.0/24	External-DB-IP	3306	Allow
Application	10.0.1.0/24	*.microsoft.com	443	Allow
Application	10.0.1.0/24	*.windowsupdate.com	443	Allow
Application	Any	Any	Any	Deny

Threat Intelligence

- Built-in feed of known malicious IPs and domains
- Can operate in Alert only or Alert and Deny mode
- Automatically updates — no maintenance required

Topic 6: VNet Peering

Connects two VNets so they can communicate privately over the Azure backbone (not the internet).

Key Properties

- Non-transitive: if VNet A peers with VNet B, and VNet B peers with VNet C, A cannot reach C through B. You must peer A-C directly or use a hub.
- Cross-region: VNets in different regions can be peered (Global VNet Peering)
- Cross-subscription: VNets in different subscriptions can be peered

- No downtime: peering is created without affecting resources in either VNet

Peering Configuration Options

Setting	What it Does	When to Enable
Allow virtual network access	Let traffic flow between peered VNets	Always yes (otherwise why peer?)
Allow forwarded traffic	Allow traffic from a third VNet that has been routed through this VNet	Yes if this VNet has a firewall/NVA that routes traffic
Allow gateway transit	Let the peered VNet use this VNet's VPN/ExpressRoute gateway	Yes if hub has gateway, spokes need on-prem access
Use remote gateways	Use the other VNet's gateway instead of deploying your own	Yes for spoke VNets in hub-spoke

Cost

Traffic Type	Cost per GB
Inbound (into VNet)	~$0.01/GB
Outbound (out of VNet)	~$0.01/GB
Same region peering	Free (inbound + outbound)
Cross-region peering	Charged both directions

Hub-Spoke Pattern

Hub peers with each spoke. Spokes do NOT peer with each other (no spoke-to-spoke directly). Spoke-to-spoke traffic goes through hub firewall. All spokes use hub's gateway for on-prem connectivity (gateway transit).

Common Mistakes

1. Forgetting that peering is non-transitive: assuming A-B and B-C means A-C
2. Not enabling Allow gateway transit on hub: spokes cannot reach on-prem
3. Overlapping IP address spaces: peering creation will fail
4. Creating too many peering relationships: use hub-spoke, not mesh (unless necessary)

Topic 7: ExpressRoute

A private, dedicated connection from your on-premises network to Azure. Does NOT go over the public internet.

Why ExpressRoute instead of VPN Gateway?

Feature	VPN Gateway	ExpressRoute
Connection type	Internet (IPsec encrypted)	Private (through provider)
Bandwidth	Up to 10 Gbps	Up to 100 Gbps
Latency	Variable (internet)	Low, consistent (<10ms typical)
Reliability	Internet-dependent	SLA-backed (99.95%)
Cost	~$150-500/month	~$1,000-5,000/month (circuit)

Feature	VPN Gateway	ExpressRoute
Setup time	Minutes	Weeks (provider provision)
Use case	Low bandwidth, non-critical	High bandwidth, latency-sensitive

ExpressRoute Components

Component	What it Is
ExpressRoute Circuit	The logical connection between your on-prem and Azure. Created in Azure.
ExpressRoute Provider	A telecom/partner that provides the physical connectivity (e.g., AT&T, Equinix, Airtel)
Peering	How traffic is routed. Three types: Private, Microsoft, Public (legacy)
ExpressRoute Gateway	Azure VNet gateway that connects the circuit to your VNet. Required.

ExpressRoute Peering Types

Peering	What it Connects To	Address Space
Private Peering	Azure VNets (IaaS VMs, internal services)	Your private IP ranges
Microsoft Peering	Azure PaaS services (M365, Azure SQL, Storage)	Public IP ranges owned by Microsoft

ExpressRoute Gateway SKUs

SKU	Bandwidth	Connections	Cost/month
ErGw1AZ	1 Gbps	4	~$300
ErGw2AZ	2 Gbps	8	~$600
ErGw3AZ	10 Gbps	16	~$1,500

Gateway must be in the GatewaySubnet (min /27).

ExpressRoute + Hub-Spoke

- ExpressRoute gateway in hub VNet
- All spokes use gateway transit to reach on-prem through hub
- No need for separate gateways in each spoke

ExpressRoute Global Reach

- Connects two ExpressRoute circuits together
- Example: on-prem Mumbai to Azure East US, AND on-prem London to Azure West Europe
- Global Reach connects the two circuits: Mumbai on-prem can reach London on-prem via Azure backbone

ExpressRoute FastPath

- Bypasses the ExpressRoute gateway for data plane traffic
- Improves performance (lower latency, higher packets per second)
- Supported on ErGw1AZ/2AZ/3AZ SKUs
- Some features not compatible (VNet peering, UDRs on gateway subnet): check before enabling

---

Topic 8: VPN Gateway

Creates an encrypted IPsec/IKE tunnel between your on-prem network and Azure over the public internet.

Types

Type	What it Does
Site-to-Site (S2S)	Connect on-prem network to Azure VNet. Requires VPN device on-prem.
Point-to-Site (P2S)	Connect individual clients (laptops) to Azure VNet. No VPN device needed.
VNet-to-VNet	Connect two Azure VNets via IPsec tunnel (alternative to VNet peering).

VPN Gateway SKUs

SKU	S2S Tunnels	P2S Connections	Bandwidth	Cost/month
VpnGw1	10	128	650 Mbps	~$140
VpnGw2	10	500	1 Gbps	~$370
VpnGw3	30	1,000	1.25 Gbps	~$900
VpnGw4	30	5,000	2.5 Gbps	~$1,400
VpnGw5	30	10,000	5 Gbps	~$2,800

Active-Active vs Active-Standby

- Active-Standby (default): One gateway instance. If it fails, standby takes over (2-4 seconds failover).
- Active-Active: Two gateway instances. Both active simultaneously. Better failover and more bandwidth.

Site-to-Site Setup

On-Prem Requirements: - VPN device (hardware or software) that supports IPsec/IKE - Public-facing IP address - Shared key (pre-shared key for authentication)

Azure Side: - GatewaySubnet (min /27) - VPN Gateway (created in GatewaySubnet) - Local Network Gateway (represents on-prem: contains on-prem public IP and address ranges) - Connection object (links VPN Gateway to Local Network Gateway)

BGP with VPN Gateway

- BGP enables dynamic route exchange between Azure and on-prem
- Without BGP: you manually configure address prefixes in Local Network Gateway
- With BGP: routes are learned automatically. If on-prem adds a new subnet, Azure learns it automatically.
- Requires ASN (Autonomous System Number): Azure assigns one by default, or you can use your own

Point-to-Site (P2S) Authentication Methods

Method	How it Works
Certificate	Generate root cert, upload to Azure, issue client certs. Clients connect with client cert.
RADIUS	Authenticate against on-prem RADIUS server (integrates with AD, MFA)
Entra ID	Authenticate with Azure AD credentials (OpenID Connect). Supports MFA natively.

When to use VPN Gateway vs ExpressRoute

Scenario	Choose
Budget is tight	VPN Gateway
Need connection in minutes/days	VPN Gateway
Need low, consistent latency	ExpressRoute
Need high bandwidth (>1 Gbps sustained)	ExpressRoute
Compliance requires no internet traversal	ExpressRoute
Hybrid connectivity for dev/test	VPN Gateway
Production hybrid for enterprise	ExpressRoute (with VPN as backup)

Best practice for enterprise: ExpressRoute primary + VPN Gateway as backup. If ExpressRoute goes down, VPN takes over.

Topic 9: Azure Bastion

A fully managed PaaS service that provides secure RDP and SSH access to VMs over TLS (port 443) directly from the Azure Portal.

Why Bastion instead of public RDP/SSH?

- No public IP needed on VMs
- No open port 22/3389 on NSG
- No VPN or ExpressRoute required
- Access from Azure Portal via browser
- All traffic stays inside Azure backbone
- Audited, logged

How it works

1. You deploy Bastion in AzureBastionSubnet (min /26)
2. In Azure Portal, click Bastion on any VM
3. Browser opens a TLS session to Bastion
4. Bastion connects to the VM’s private IP via RDP/SSH
5. Session appears in your browser

SKUs

SKU	Features	Cost
Basic	RDP/SSH, 2 concurrent sessions	~$140/month
Standard	Basic + more sessions, shareable links, clipboard upload/download	~$140/month+
Premium		Higher

SKU	Features	Cost
	Standard + private-only connect, Kerberos authentication	

Bastion vs Jumpbox (Bastion Host VM)

Feature	Azure Bastion	Jumpbox VM
Management	Fully managed (PaaS)	You manage OS, patching, hardening
Public IP on VM	No	Yes (or VPN required)
NSG changes	Not needed on target VM	Need to allow RDP/SSH
Access	Browser-based	RDP/SSH client
Cost	~$140/month fixed	VM cost (variable)
Audit	Built-in	Manual (enable logging)
Risk of compromise	Very low	Higher (exposed VM)

Always prefer Azure Bastion for production. Jumpbox only if Bastion is not available in your region or budget is extremely tight.

Topic 10: User Defined Routes (UDR) / Route Tables

Custom routing rules that override Azure's default routing. Control where traffic goes next.

Why UDRs?

- Force traffic through Azure Firewall or NVA (network virtual appliance)
- Override default routing (e.g., prevent direct internet access)
- Route specific subnets through specific gateways

How it works

- Create a Route Table
- Add routes to it (destination + next hop)
- Associate the Route Table with a subnet
- All VMs in that subnet follow these routes

Route Properties

Property	What it Means
Address prefix	Destination CIDR (e.g., 0.0.0.0/0 = all traffic)
Next hop type	VirtualAppliance, VirtualNetworkGateway, Internet, None, VNetLocal
Next hop IP	IP of the firewall/NVA (when next hop = VirtualAppliance)

Common UDR Patterns

1. Force all internet traffic through Azure Firewall:

Destination	Next Hop Type	Next Hop IP
0.0.0.0/0	VirtualAppliance	10.0.0.4 (firewall IP)

2. Force on-prem traffic through ExpressRoute gateway:

Destination	Next Hop Type
192.168.0.0/16	VirtualNetworkGateway

3. Blackhole (drop) traffic to specific range:

Destination	Next Hop Type
10.99.0.0/16	None

Azure Default Routes (system routes, always present)

Destination	Next Hop	What it Does
VNet address space	VNetLocal	Route within VNet
0.0.0.0/0	Internet	Default internet outbound
10.0.0.0/8	None	Blackhole private ranges
192.168.0.0/16	None	Blackhole private ranges
172.16.0.0/12	None	Blackhole private ranges

UDRs override system routes. More specific routes (longer prefix) always win over less specific.

Hub-spoke UDR setup

- Each spoke subnet has a route table with: 0.0.0.0/0 to Azure Firewall IP (in hub)
- This forces all spoke internet-bound traffic through the firewall for inspection
- Without this UDR, spoke VMs would go directly to internet

Topic 11: Azure DNS

DNS hosting service in Azure. Resolves domain names to IP addresses.

Two distinct services

Service	What it Does
Azure DNS (Public)	Host public DNS zones (e.g., mycompany.com). Resolves for anyone on the internet.
Azure Private DNS	Host private DNS zones (e.g., internal.corp). Resolves only within VNets.

Azure DNS (Public)

- Host your domain DNS records in Azure
- Same cost whether 1 record or 1000
- ~$0.50/month per zone + ~$0.20 per million queries
- Record types: A, AAAA, CNAME, MX, NS, PTR, SOA, SRV, TXT
- You delegate your domain NS to Azure at your domain registrar

Azure Private DNS

- Resolves internal names (e.g., db.internal.corp to 10.0.3.4)
- Only accessible from VNets it is linked to
- Auto-registration: VMs automatically get DNS records when created
- Link to VNet: enable auto-registration on the VNet link

Why Private DNS?

- Without it: you access VMs by IP (10.0.3.4) — hard to remember, breaks if IP changes
- With it: you access VMs by name (db.internal.corp) — readable, auto-updates if IP changes
- Critical for Private Endpoints (Topic 12)

Custom DNS for VNets

Option	When to Use
Default (Azure-provided)	Simple setups, no on-prem DNS needed
Custom DNS (on-prem DNS servers)	Hybrid environments, need to resolve on-prem names
Azure Private DNS	Cloud-only, need internal name resolution
Custom DNS + Azure Private DNS	Hybrid — forward on-prem queries to on-prem DNS, resolve Azure in Private DNS

Common DNS Record Types

Type	What it Does	Example
A	Maps name to IPv4	myapp.com to 203.45.67.89
AAAA	Maps name to IPv6	myapp.com to 2001:db8::1
CNAME	Maps name to another name	www.myapp.com to myapp.azurewebsites.net
MX	Mail exchange	myapp.com to mail.myapp.com (priority 10)
TXT	Text records (verification, SPF, DKIM)	myapp.com to “v=spf1 include:spf.protection.outlook.com”
NS	Name server delegation	myapp.com to ns1-01.azure-dns.com
SRV	Service location	_sip._tcp.myapp.com to 10 60 5060 sipserver.myapp.com
PTR	Reverse DNS (IP to name)	89.67.45.203.in-addr.arpa to myapp.com

Alias Records

- Special Azure DNS record type
- Points to an Azure resource (App Service, Traffic Manager, Public IP)
- If the resource IP changes, alias record updates automatically
- No manual update needed

Topic 12: Private Link and Private Endpoints

Access Azure PaaS services (Storage, SQL, Key Vault, etc.) over a private IP address inside your VNet — no public internet exposure.

Without Private Endpoint

VM -> Internet -> Storage Account (public endpoint) Traffic goes over the public internet (or Microsoft backbone). Storage account has a public IP.

With Private Endpoint

VM -> Private IP (10.0.1.10) -> Private Endpoint -> Storage Account Traffic stays entirely on the Microsoft backbone. Storage account has NO public endpoint.

How it works

1. Create a Private Endpoint for the PaaS service
2. It gets a private IP from your subnet
3. Create a Private DNS zone for the service (e.g., privatelink.blob.core.windows.net)
4. Link the DNS zone to your VNet
5. VMs access the service using its normal URL, but DNS resolves to the private IP

What services support Private Endpoints

Category	Services
Storage	Blob, File, Queue, Table, Data Lake
Database	Azure SQL, PostgreSQL, MySQL, Cosmos DB, MariaDB
Analytics	Synapse, Data Factory, Databricks
Security	Key Vault, Managed HSM
Integration	Service Bus, Event Grid, Event Hubs
Web	App Service, Static Web Apps
AI	Cognitive Services, OpenAI

Service Endpoints vs Private Endpoints

Feature	Service Endpoint	Private Endpoint
IP type	Public (routed through Azure backbone)	Private (in your VNet)
Access from on-prem	No (only from VNet)	Yes (via ExpressRoute/VPN)
Access from peered VNets	No	Yes
Granular access	Per service (all storage accounts)	Per resource (specific storage account)
Cost	Free	~$7/month per endpoint
DNS	No change needed	Needs Private DNS zone

Always prefer Private Endpoints for production. Service Endpoints are older, less secure (still public IP).

Disable public access

When using Private Endpoints, you should also: - Disable public network access on the PaaS resource - Set “Public network access: Disabled” - This ensures the service is ONLY reachable via Private Endpoint

---

Topic 13: Azure Front Door

A global, scalable entry point for your web applications. Routes traffic to the nearest healthy backend across regions.

When to use Front Door

Scenario	Choose
Single region, single app	Application Gateway
Multi-region, global app	Front Door
Need global load balancing	Front Door
Need CDN/caching	Front Door
Need WAF at global edge	Front Door WAF

SKUs

SKU	Features	Cost
Standard	Global load balancing, health probes, SSL termination	~$35/month + per-GB
Premium	Standard + Private Link to backends, WAF	~$330/month + per-GB

Key Features

Feature	What it Does
Global HTTP/HTTPS load balancing	Route users to nearest healthy backend
SSL termination	Offload SSL at the edge
WAF	Web Application Firewall at global edge
URL-based routing	/app to region 1, /api to region 2
Session affinity	Sticky sessions by cookie
Health probes	Check backend health per region
CDN	Cache static content at edge (Azure CDN integrated)

Front Door vs Application Gateway vs Traffic Manager

Feature	Front Door	Application Gateway	Traffic Manager
Layer	Layer 7 (HTTP/HTTPS)	Layer 7 (HTTP/HTTPS)	Layer 3 (DNS)
Scope	Global (multi-region)	Regional (single region)	Global (DNS-based)
SSL termination	Yes	Yes	No
WAF	Yes	Yes	No
URL routing	Yes	Yes	No
Cookie affinity	Yes	Yes	No
Failover speed	Seconds	N/A	30-120 seconds (DNS TTL)
Protocol	HTTP/HTTPS only	HTTP/HTTPS only	Any (DNS-based)

Traffic Manager is DNS-based: it returns a different IP based on routing method. Client connects directly. Slow failover. No health checking of the actual endpoint.

Front Door is proxy-based: client connects to Front Door, Front Door forwards to backend. Instant failover. Health checks actual endpoint. But only HTTP/HTTPS.

Topic 14: Azure DDoS Protection

Protects Azure resources from Distributed Denial of Service attacks.

SKUs

SKU	Cost	What it Does
Basic	Free	Always on. Protects Azure infrastructure. No configuration.
Standard	~$2,944/month	Protects your specific resources. Adaptive tuning, metrics, alerts, rapid response support.

Basic protects Azure network. Standard protects YOUR VNets and resources.

What Standard gives you

- Always-on monitoring for your public IP resources
- Adaptive tuning: learns your app normal traffic pattern, adjusts thresholds
- DDoS mitigation metrics and alerts
- Access to DDoS Rapid Response (DRR) team during active attacks
- Cost protection: if you are charged for scale-out during a DDoS attack, Microsoft covers it
- Attack analytics: see attack duration, vectors, drop rates

When to use Standard

- Public-facing web applications (e-commerce, portals)
- Regulated industries (financial, healthcare) with compliance requirements
- Applications that have been attacked before

When Basic is fine

- Internal-only applications
- Dev/test environments
- Non-critical workloads

Topic 15: Network Watcher

A network monitoring and diagnostic service. Enabled automatically when you create or update a VNet in a region.

Key Features

Feature	What it Does	When to Use
IP Flow Verify	Check if a packet is allowed/denied by NSG rules. Tells you which rule blocked it.	Troubleshooting: “Why cannot VM1 reach VM2?”
Next Hop	Shows the next hop for a packet from a VM. Tells you if traffic is going to internet, gateway, NVA, etc.	Troubleshooting: “Where is this traffic actually going?”
NSG Flow Logs	Log all NSG allow/deny decisions. Send to Log Analytics or Storage Account.	Security auditing, compliance, “who talked to whom”
Connection Troubleshoot	Test TCP connectivity from VM to another VM, FQDN, or IP. Shows latency and hop-by-hop path.	Troubleshooting: “Can VM1 reach the database?”
Packet Capture	Capture packets on a VM NIC (like tcpdump). Save to Storage Account.	Deep network troubleshooting
Topology	Visual map of resources in a VNet and their connections.	Understanding network layout
Network Performance Monitor	Monitor network performance between Azure and on-prem, or between Azure regions.	Proactive monitoring for latency/ packet loss

NSG Flow Logs — most important for security

Setting	Options
Version	Version 1 (5-tuple only) or Version 2 (adds flow tuple, throughput)
Retention	Store in Storage Account (days) or send to Log Analytics
Traffic Analytics	Process flow logs to show top talkers, traffic patterns, visualization

Traffic Analytics shows: - Which VMs are talking most - Which protocols/ ports are used - Traffic between subnets, VNets, regions - Top talkers - Network security group effectiveness

---

Topic 16: Azure Virtual WAN

A unified networking service that brings all networking connectivity (VPN, ExpressRoute, site-to-site, point-to-site) into a single operational interface.

When to use Virtual WAN instead of individual gateways

Scenario	Choose
1-3 branch offices, simple connectivity	VPN Gateway + ExpressRoute Gateway
10+ branch offices, complex global network	Virtual WAN
Need optimized routing between branches	Virtual WAN
Need SD-WAN integration	Virtual WAN

Virtual WAN Types

Type	What it Does
Basic	S2S VPN and P2S VPN only
Standard	Everything: S2S VPN, P2S VPN, ExpressRoute, inter-hub, VNet-to-VNet

Key Concepts

Component	What it Is
Virtual WAN	The top-level resource. Can have multiple hubs.
Hub	A Microsoft-managed virtual network. Contains gateways (VPN, ExpressRoute).
Spoke VNet	Your VNets connected to the hub via VNet connection.
Branch/Site	On-prem location connected via VPN or ExpressRoute.
Hub routing	All branches and VNets can communicate through the hub. Transit routing built-in.

Virtual WAN vs Hub-Spoke (DIY)

Feature	DIY Hub-Spoke	Virtual WAN
Branch-to-VNet	Manual peering + UDRs	Automatic
Branch-to-branch	Manual BGP + routing	Automatic
Transit routing	Custom (UDRs, NVAs)	Built-in
Management	Multiple gateways, configs	Single pane of glass
Scale	Limited by gateway SKUs	Designed for hundreds of branches
Cost	Individual gateway costs	Hub hourly + egress

For enterprise-scale with many branches, Virtual WAN is the right choice. For simpler setups, DIY hub-spoke is fine.

Topic 17: Azure Route Server

Simplifies dynamic routing between your NVAs (network virtual appliances) and Azure virtual networks.

Why Route Server?

- If you deploy a third-party NVA (Palo Alto, Fortinet, Cisco) in your VNet
- The NVA needs to exchange routes with Azure
- Without Route Server: you manually configure UDRs on every subnet
- With Route Server: NVA exchanges BGP routes with Azure automatically

How it works

1. Deploy Route Server in a dedicated subnet (RouteServerSubnet, min /27)
2. Configure BGP peering between Route Server and your NVA

3. Route Server learns routes from NVA and programs them into Azure routing table
4. Azure VNet routes are also shared with the NVA

Key points

- Supports BGP only
- Works with any NVA that supports BGP
- Max 4 BGP peers
- Free (no additional cost)

Topic 18: Service Endpoints (Legacy — prefer Private Endpoints)

Brings Azure service traffic onto the Azure backbone instead of the public internet.

How it works

- Enable Microsoft.Storage service endpoint on a subnet
- Traffic to Storage from that subnet goes over Azure backbone, not internet
- Storage still has a public endpoint — it just verifies the traffic came from your VNet

Limitations (why Private Endpoints are preferred)

- Not reachable from on-prem (only from the VNet)
- Not reachable from peered VNets
- Not per-resource — applies to ALL storage accounts in the region
- Still uses public IP (just routed through backbone)

When you might still use Service Endpoints

- Budget is very tight (free vs ~$7/month per Private Endpoint)
- Simple cloud-only workloads where on-prem access is not needed
- Bulk access to a service type (all storage accounts)

Topic 19: Azure NAT Gateway

Managed outbound NAT service for VMs that need to connect to the internet.

Why NAT Gateway?

- By default, VMs access internet via outbound SNAT through the load balancer or default public IP
- Default outbound SNAT is unpredictable — the public IP can change, and there are port limitations

- NAT Gateway gives you stable, predictable, scalable outbound connectivity

How it works

1. Create a NAT Gateway with one or more static public IPs
2. Attach it to a subnet
3. All VMs in that subnet use the NAT Gateway for outbound internet access
4. Outbound traffic appears to come from the NAT Gateway static public IPs

Key Properties

Property	Details
Public IPs	1-16 static public IPs per NAT Gateway
Ports per IP	64,000 concurrent outbound connections per IP
Scaling	Automatically scales to handle more connections
Timeout	TCP idle timeout: 4 minutes (configurable up to 120)
Cost	~$32/month + $0.045/GB processed

NAT Gateway vs Default Outbound

Feature	Default Outbound	NAT Gateway
Public IP	Dynamic, unpredictable	Static, predictable
Port exhaustion	Possible under high load	64K ports per IP, add more IPs
SNAT port allocation	Per-VM, limited	Dynamic, shared pool
Whitelisting	Hard (IP changes)	Easy (static IPs)
SLA	None	99.9%

When to use NAT Gateway

- VMs need to reach internet and you need a stable source IP (for whitelisting by external services)
- High outbound connection volume (default SNAT runs out of ports)
- Compliance requires predictable, auditable outbound IPs

Topic 20: Network Virtual Appliances (NVAs)

Third-party firewalls/routers deployed as VMs in Azure. Examples: Palo Alto, Fortinet, Cisco, Check Point, F5.

Why NVAs instead of Azure Firewall?

Feature	Azure Firewall	NVA (e.g., Palo Alto)
Management	Fully managed (PaaS)	You manage OS, patches, HA config
Features	Good for most use cases	Advanced features (deep packet inspection, app-aware policies, VPN to other vendors)
Cost	~$500-900/month	VM cost + license (varies widely)
HA	Built-in (zone-redundant)	You configure (active/passive or active/active)
TLS inspection	Premium SKU only	Most NVAs support it
Vendor support	Microsoft	Appliance vendor

Feature	Azure Firewall	NVA (e.g., Palo Alto)
Learning curve	Lower	Higher (vendor-specific)

NVA HA Architecture

Internet -> Azure Load Balancer -> NVA-1 (Active) -> Backend -> NVA-2 (Passive) -> Backend

- Deploy 2 NVAs in an Availability Set
- Internal LB distributes traffic to the active NVA
- If active NVA fails, LB routes to passive
- Both NVAs sync state (session tables, config)

When to choose NVA

- Organization already uses a specific vendor and wants consistency (same policies, same management console)
- Need advanced features Azure Firewall does not offer
- Compliance requires a specific vendor firewall

When to choose Azure Firewall

- Want fully managed, no OS patching
- Do not need vendor-specific features
- Prefer native Azure integration

Summary: Azure Networking Services by Use Case

Need	Service
Private network in Azure	VNet + Subnets
Firewall per subnet	NSG
Group VMs for NSG rules	ASG
Distribute traffic (TCP/UDP)	Azure Load Balancer
Distribute HTTP/HTTPS traffic	Application Gateway
Web Application Firewall	App Gateway WAF or Front Door WAF
Centralized network firewall	Azure Firewall
Connect two VNets	VNet Peering
Connect on-prem (private, high bandwidth)	ExpressRoute
Connect on-prem (over internet)	VPN Gateway
Secure RDP/SSH access	Azure Bastion
Custom routing	UDR / Route Tables
DNS hosting (public)	Azure DNS
DNS hosting (private/internal)	Azure Private DNS
Access PaaS privately	Private Link / Private Endpoints
Global load balancing	Front Door
DDoS protection	DDoS Protection Standard
Network monitoring/diagnostics	Network Watcher
Global networking (many branches)	Virtual WAN
Dynamic routing with NVAs	Route Server
Stable outbound internet IP	NAT Gateway

Need	Service
Third-party firewall	NVA