Interconnection of On Premises Datacenters and Public Cloud

High Availability and Security

LACNOG 2022
Wilson Lopes

Motivation

- High demand of migration of on premises workloads to public cloud
- Hybrid Infrastructure – requirement to maintain cloud x on prem connectivity
- Low latency as possible, high availability and security
- Even in cloud, network architecture must follow best practices to be high available and secure
- Use cases – Edge Computing, IoT, CDN’s, Financial and Trade Systems

Cloud Network Infrastructure

Regions and Availability Zones

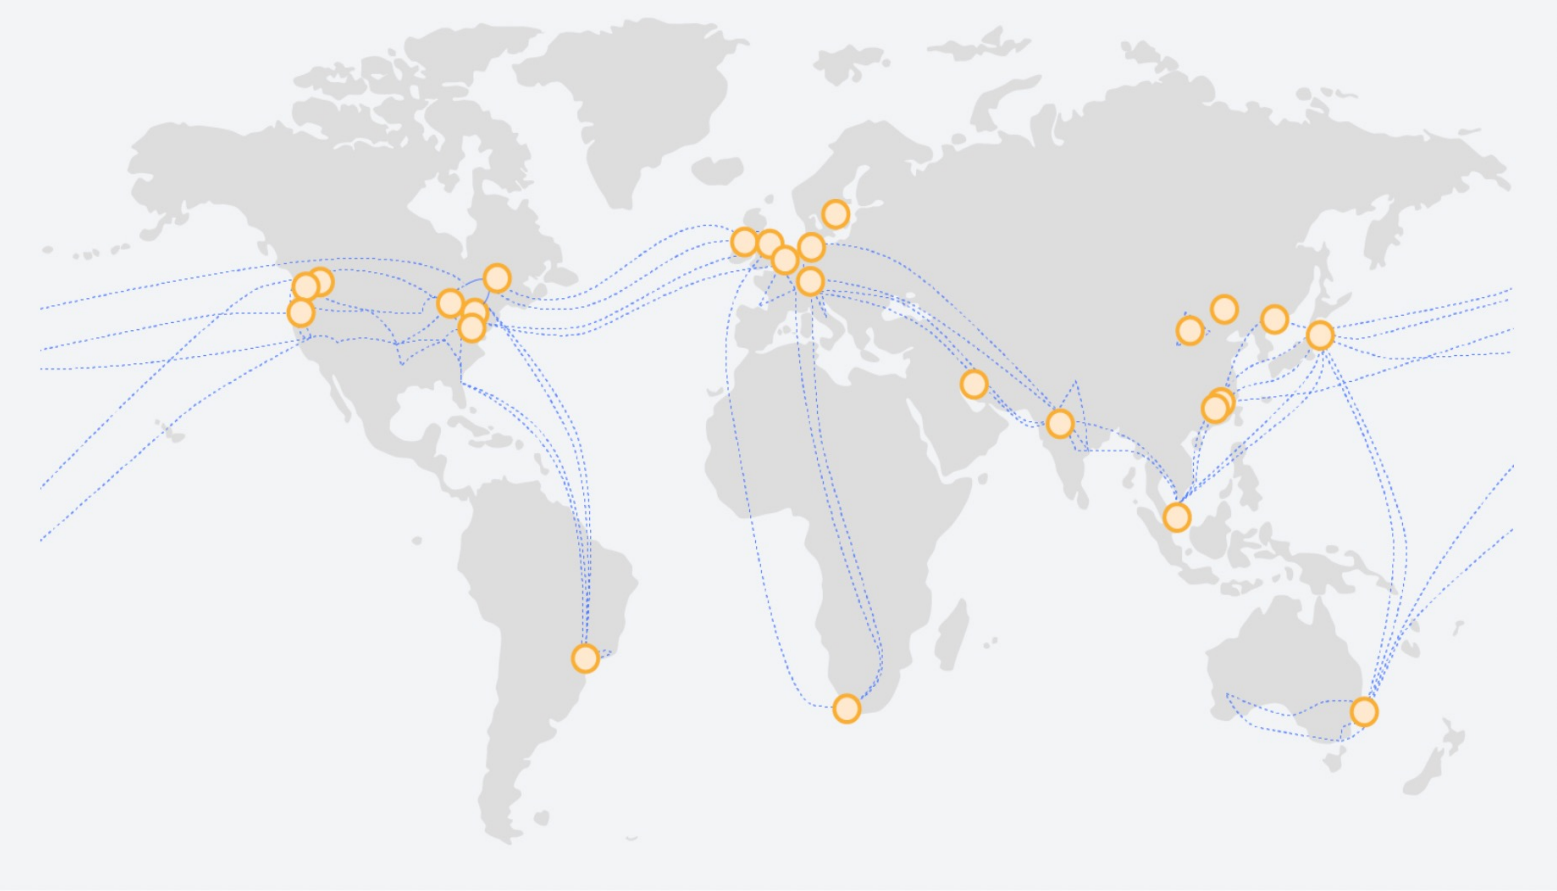

Regions

. Geographic segregated infrastructures around the globe – network, compute, storage

. Services generally are high available and isolated into the region

. Regions are connected using private cloud network infrasctructure

. Generally there are a region only at Sao Paulo in Latin America

Cloud Network Infrastructure

Regions and Availability Zones

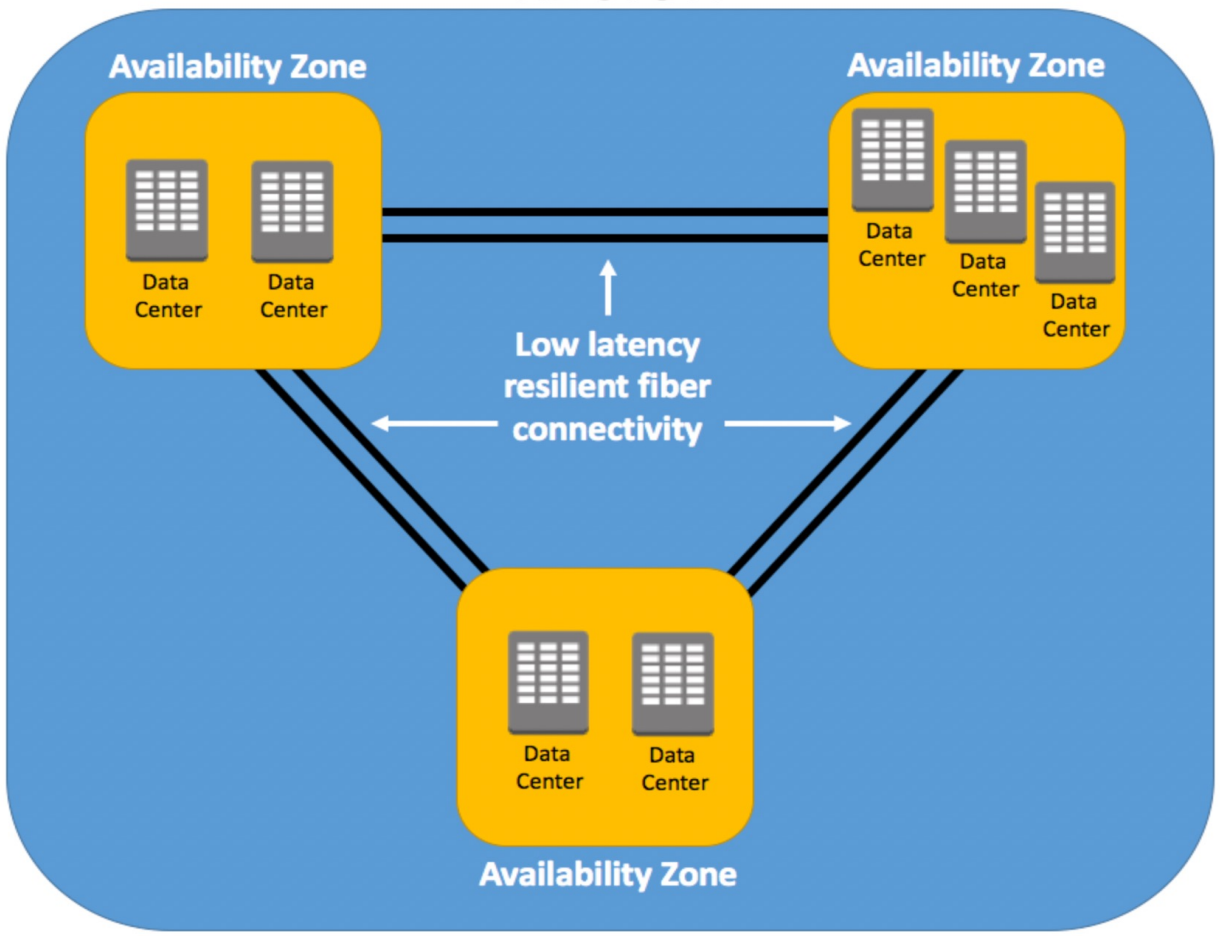

Availabilty Zones - AZs

. One or more Datacenter into the region

. Segregated network, compute and storage services into the region

. To became high available, services must be deployed in more than one AZ

. For example, a service behind a load balancer using 3 AZs, with frontend, backend and database between the 3 AZs

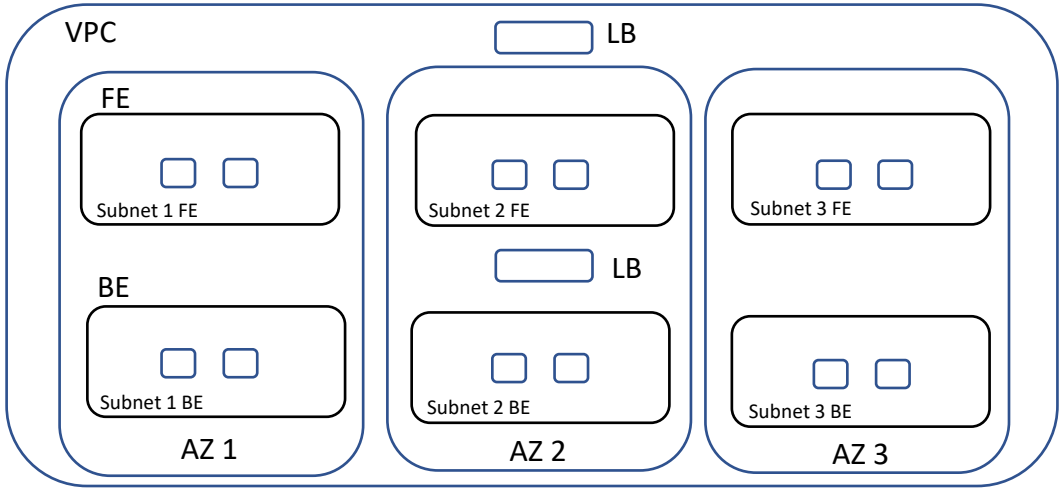

Cloud Network Infrastructure

VPC – Virtual Private Cloud

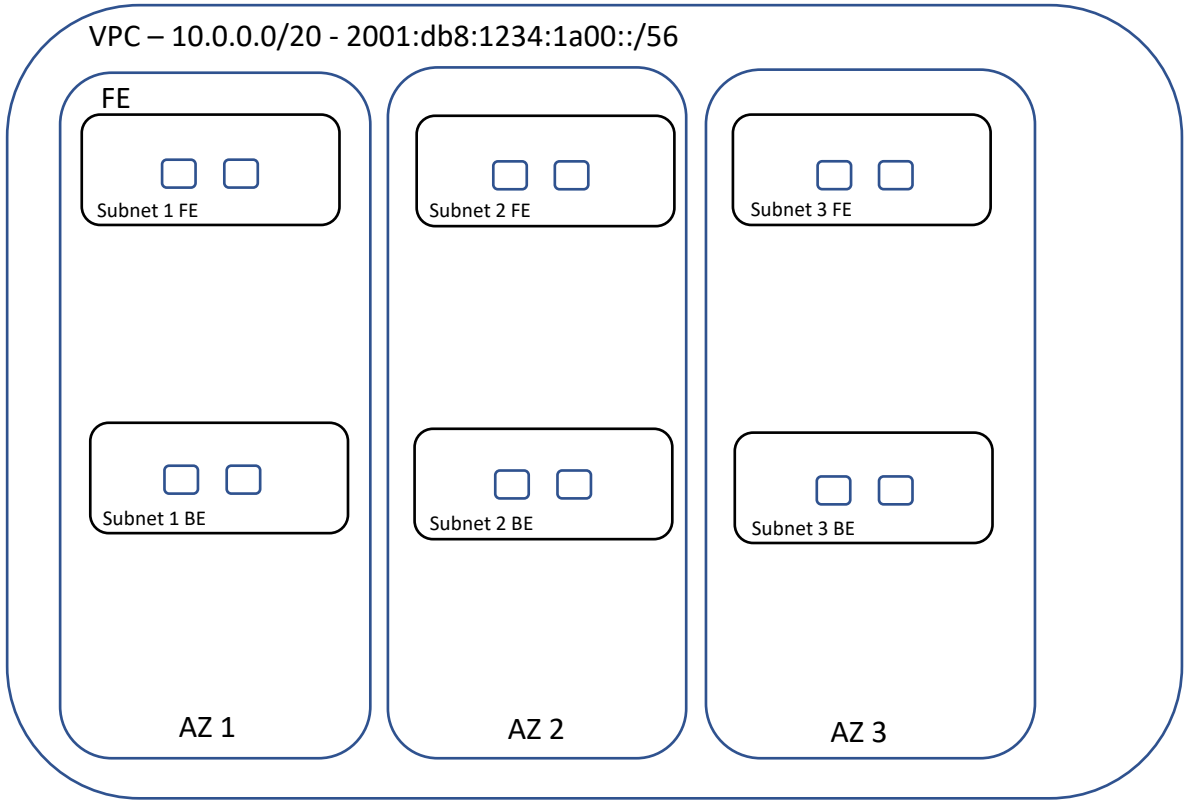

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6

Cloud Network Infrastructure

VPC – Virtual Private Cloud

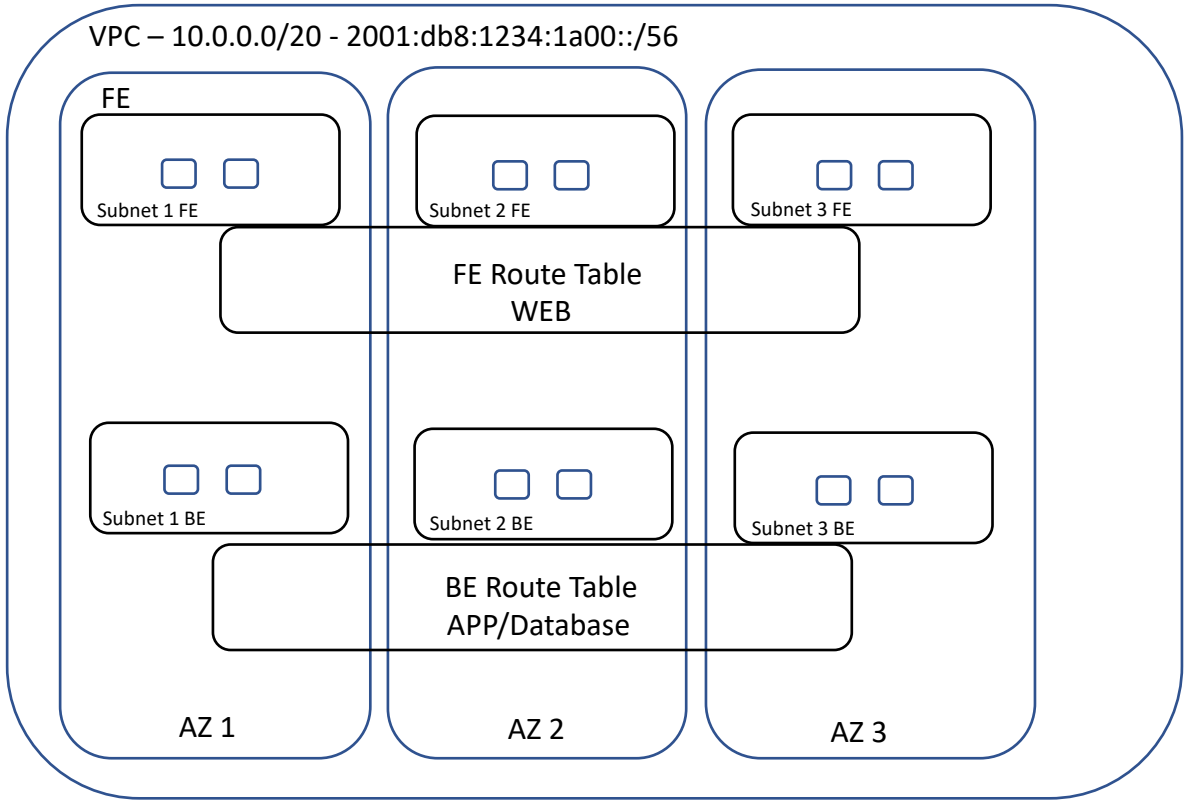

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6
. Route tables – Each subnet can use different route tables

Cloud Network Infrastructure

VPC – Virtual Private Cloud

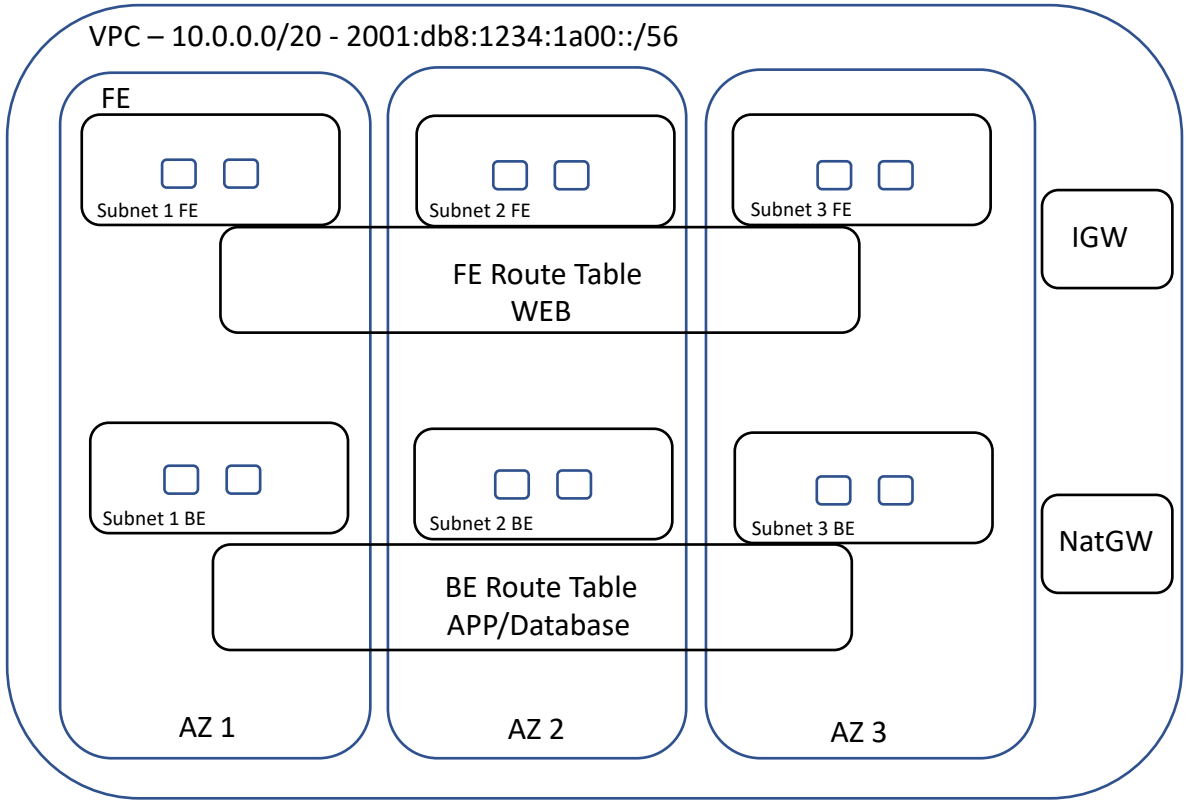

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6
. Route tables – Each subnet can use different route tables
. IGW – Internet Gateway – connect “Public Subnets” to internet – ingress/egress
. NatGW – Nat Gateway – connect “Private Subnets” to internet – egress Only

Cloud Network Infrastructure

VPC – Virtual Private Cloud

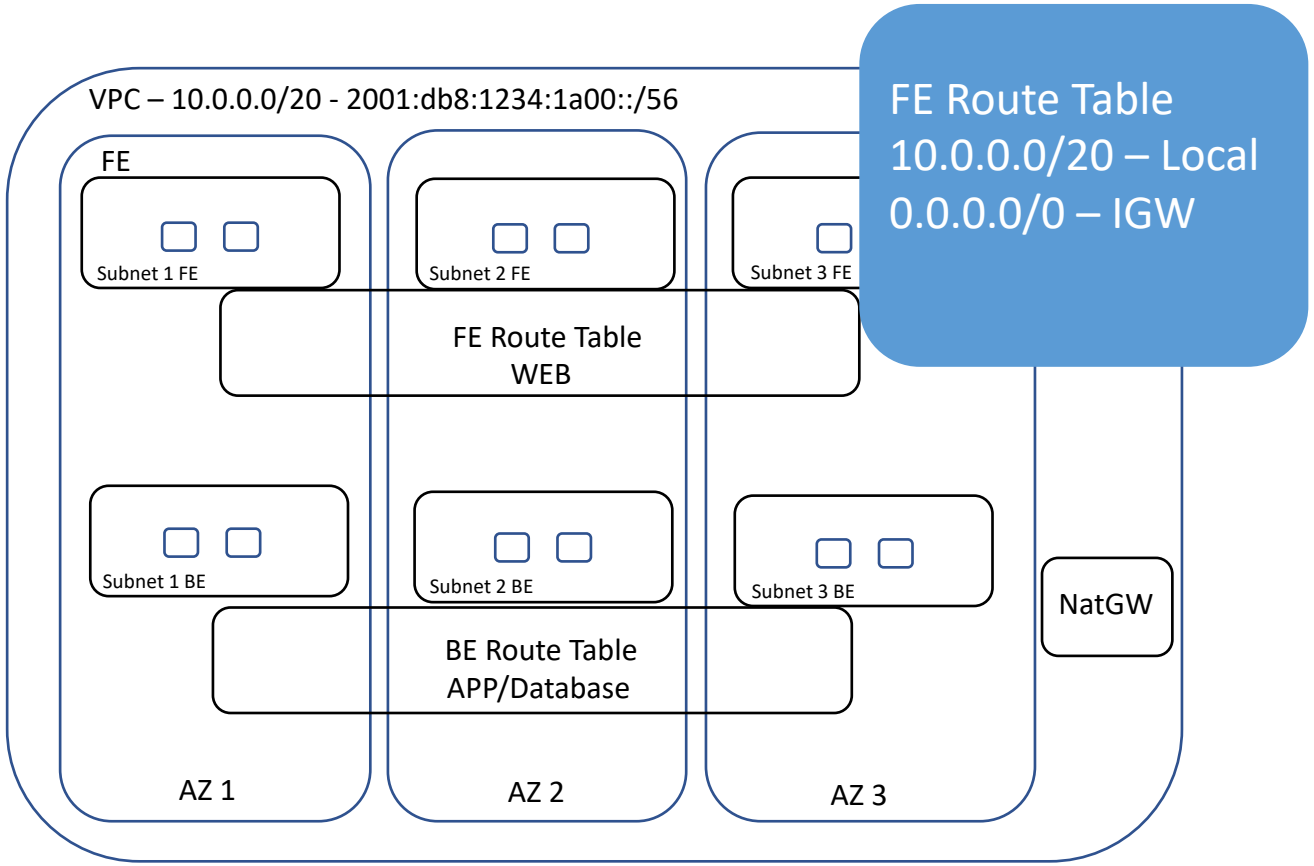

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6
. Route tables – Each subnet can use different route tables
. IGW – Internet Gateway – connect “Public Subnets” to internet – ingress/egress
. NatGW – Nat Gateway – connect “Private Subnets” to internet – egress Only

Cloud Network Infrastructure

VPC – Virtual Private Cloud

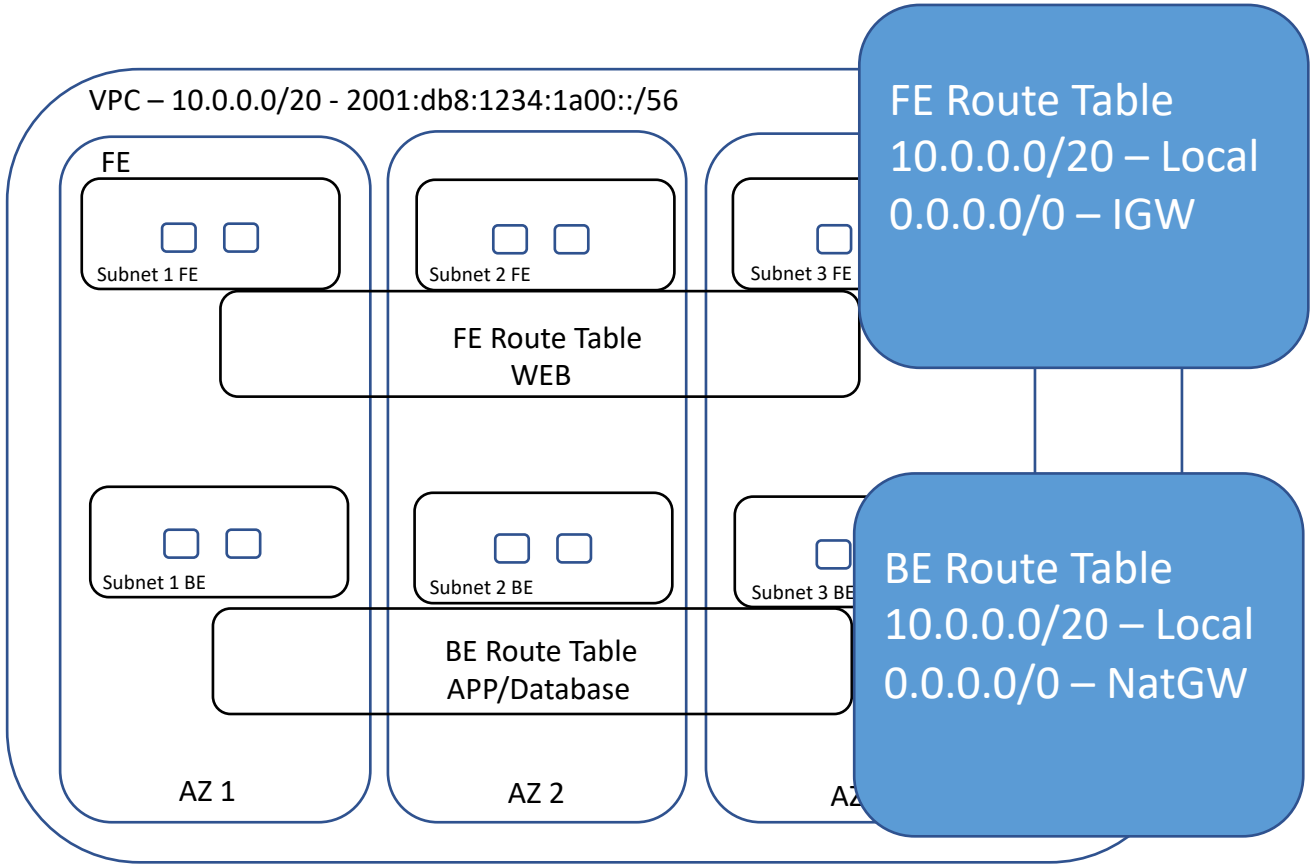

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6
. Route tables – Each subnet can use different route tables
. IGW – Internet Gateway – connect “Public Subnets” to internet – ingress/egress
. NatGW – Nat Gateway – connect “Private Subnets” to internet – egress Only

Cloud Network Infrastructure

VPC – Virtual Private Cloud

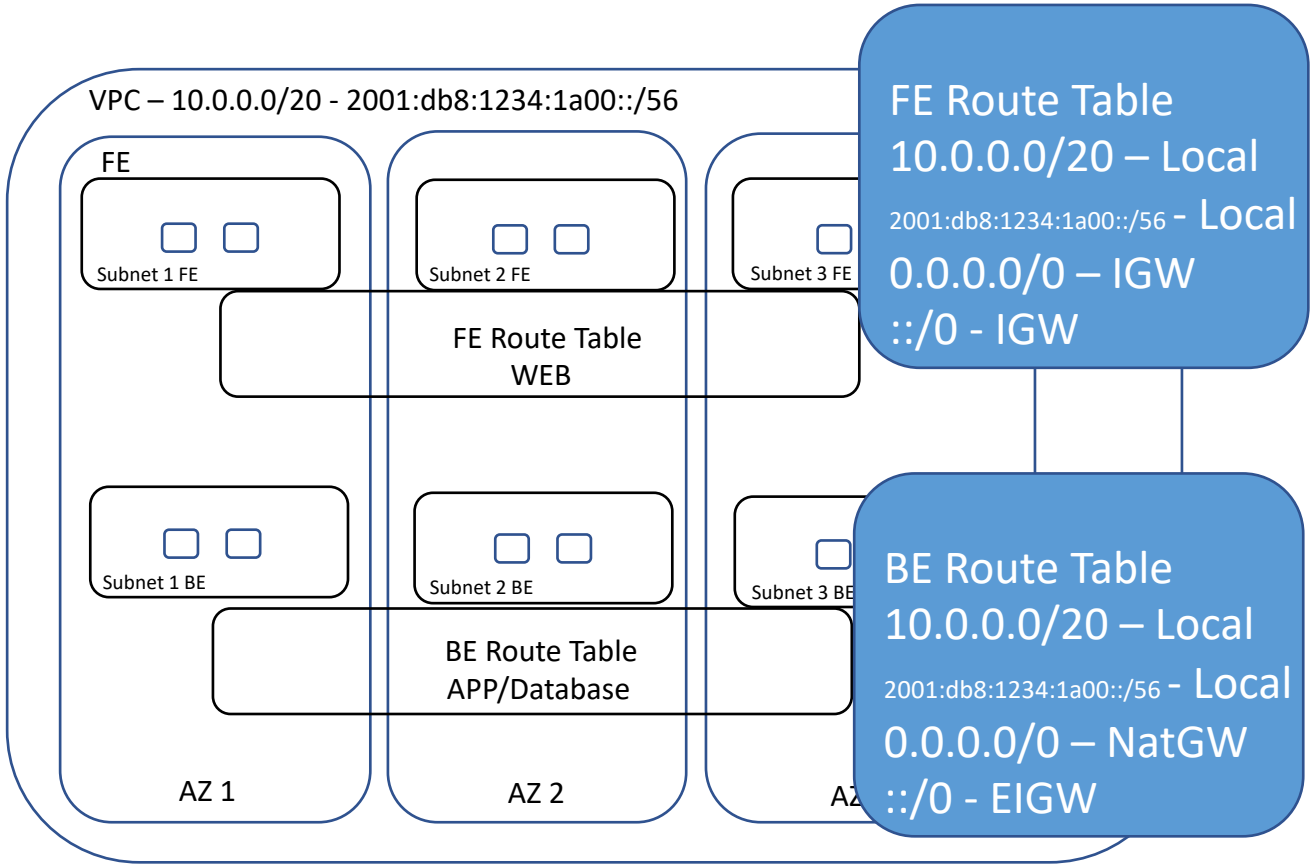

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6
. Route tables – Each subnet can use different route tables
. IGW – Internet Gateway – connect “Public Subnets” to internet – ingress/egress
. NatGW – Nat Gateway – connect “Private Subnets” to internet – egress Only
. IPv6 – Internet Gateway and Egress Only Internet Gateway

Cloud Network Infrastructure

VPC – Virtual Private Cloud

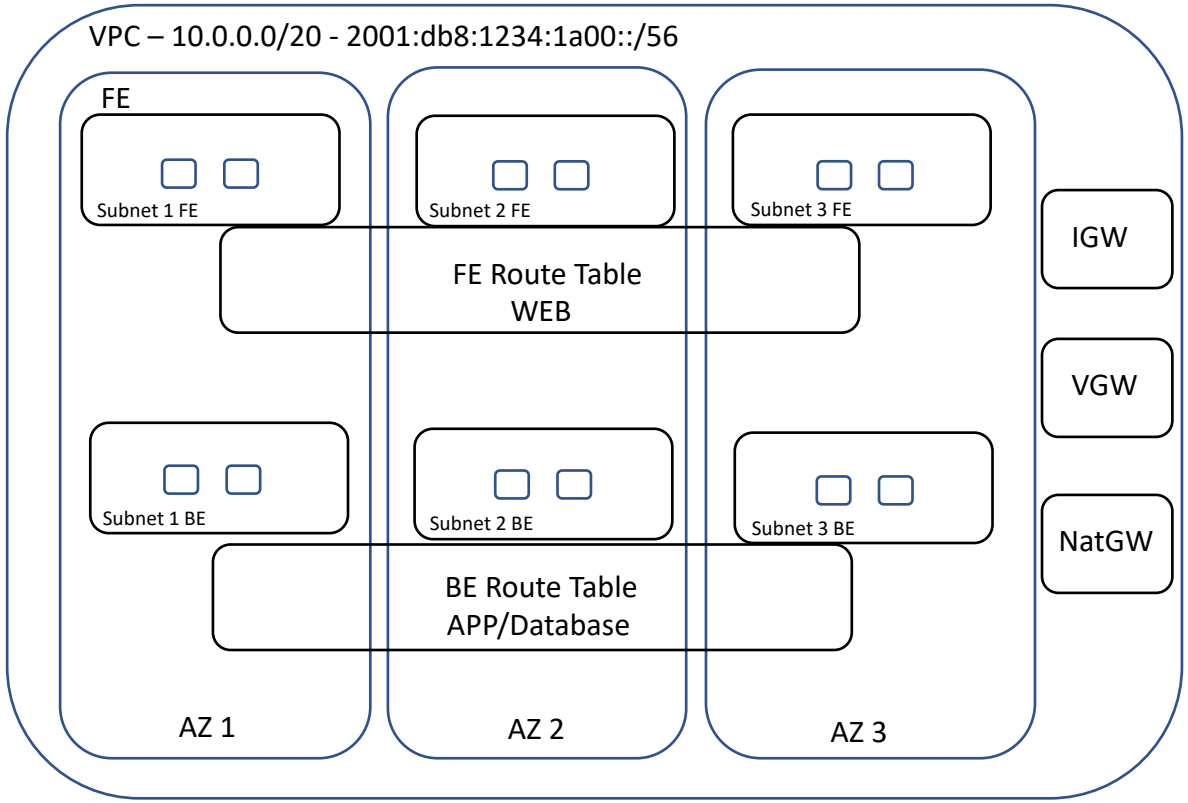

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6
. Route tables – Each subnet can use different route tables
. IGW – Internet Gateway – connect “Public Subnets” to internet – ingress/egress
. NatGW – Nat Gateway – connect “Private Subnets” to internet – egress Only
. IPv6 – Internet Gateway and Egress Only Internet Gateway
. VGW – Virtual Private Gateway – connect to external networks via VPN or private connections – should connect only “Private Subnets”. BGP and static routes.

Cloud Network Infrastructure

VPC – Virtual Private Cloud

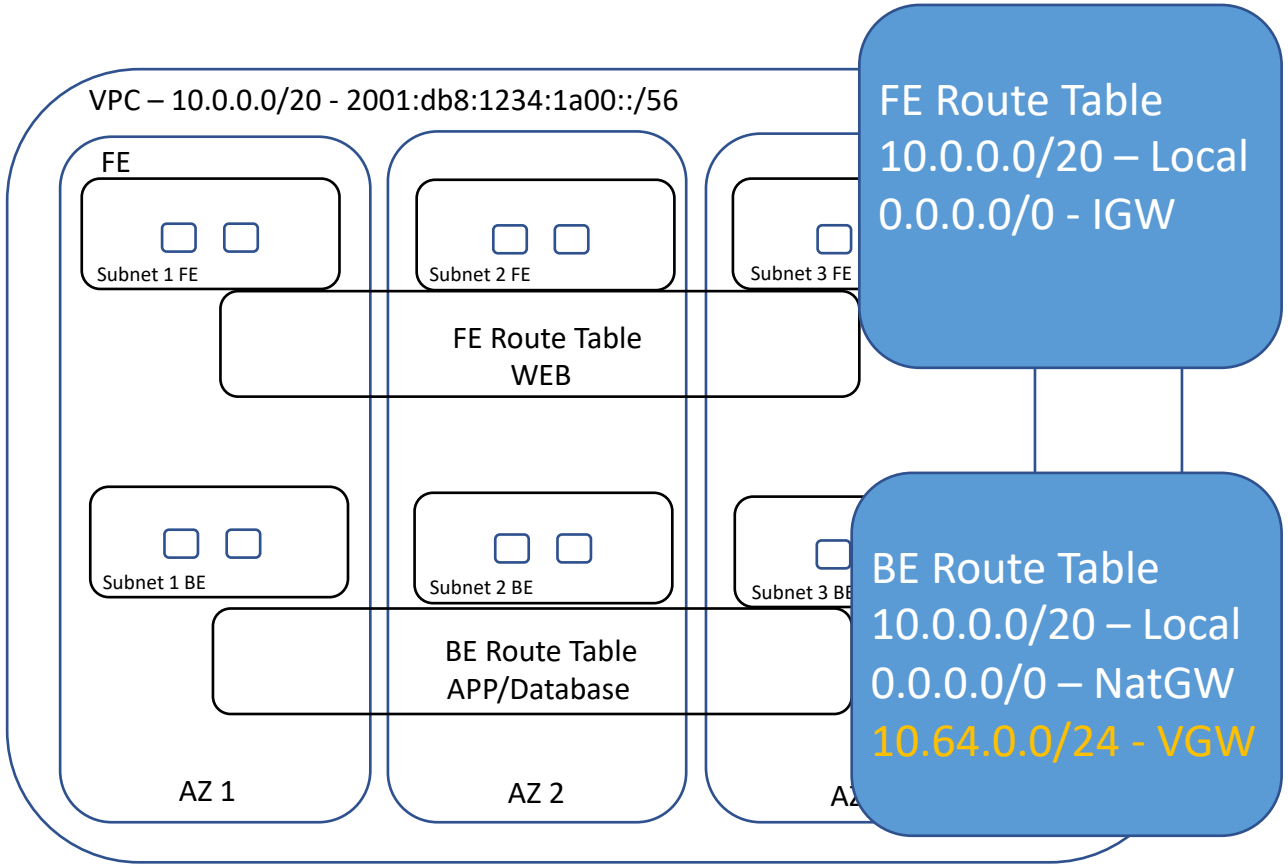

VPC

. Isolated virtual network into the region
. Segregated in subnets. Subnets are per AZ
. Ipv4/Ipv6/Ipv6 only
. Bring your own ip – public ipv4 and ipv6
. Route tables – Each subnet can use different route tables
. IGW – Internet Gateway – connect “Public Subnets” to internet – ingress/egress
. NatGW – Nat Gateway – connect “Private Subnets” to internet – egress Only
. IPv6 – Internet Gateway and Egress Only Internet Gateway
. VGW – Virtual Private Gateway – connect to external networks via VPN or private connections – should connect only “Private Subnets”. BGP and static routes.

Cloud Network Infrastructure

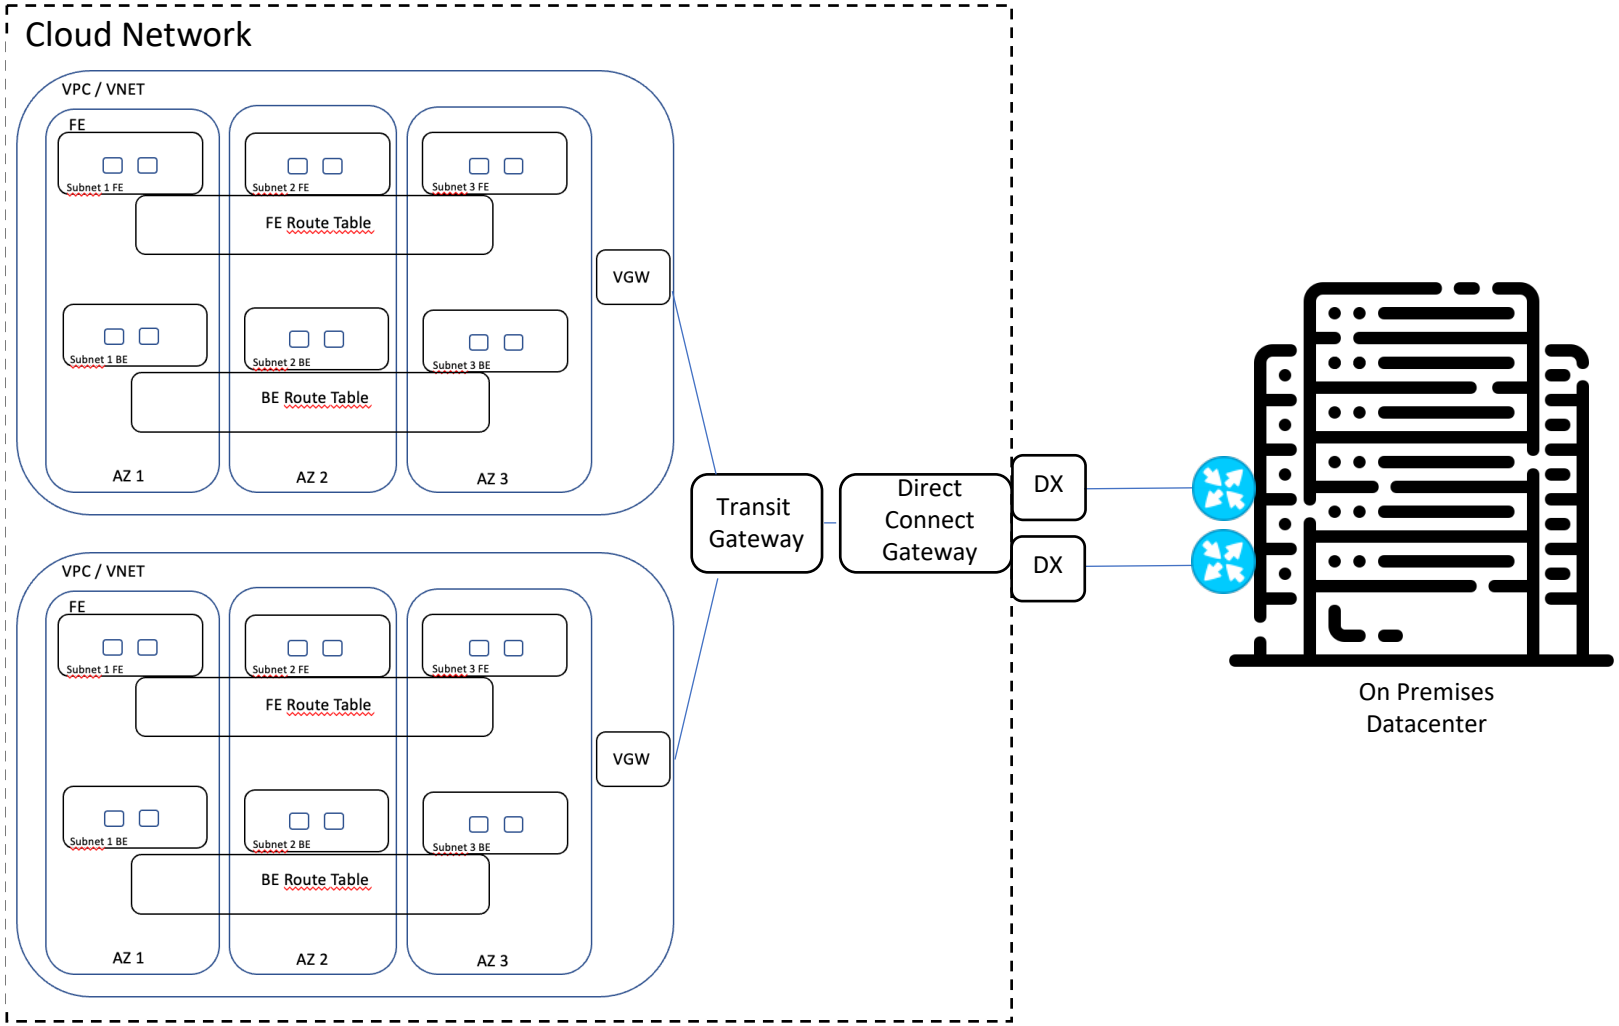

Cloud Network Infrastructure

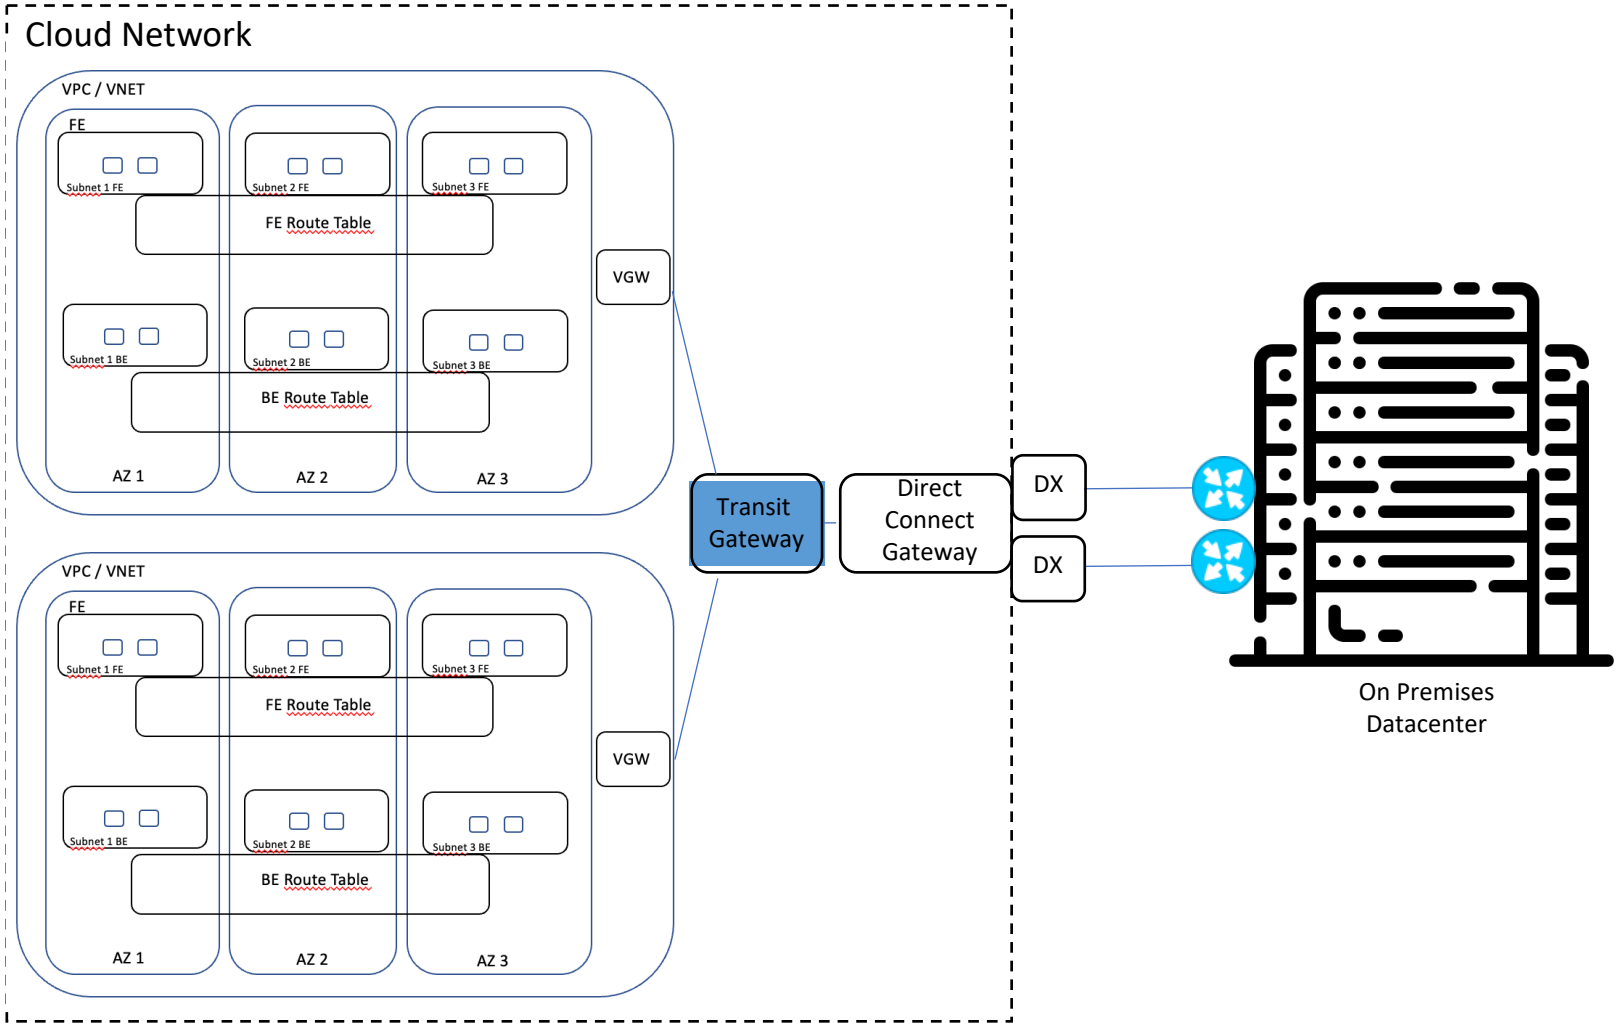

Transit Gateway

. Interconnection Hub – connect to VGWs

. Per region resource

. Direct connect and VPN support

. Route received from direct connect are prefered

. Static and dynamic bgp routing

. Limits – 10k static routes, 1k dynamic routes

Cloud Network Infrastructure

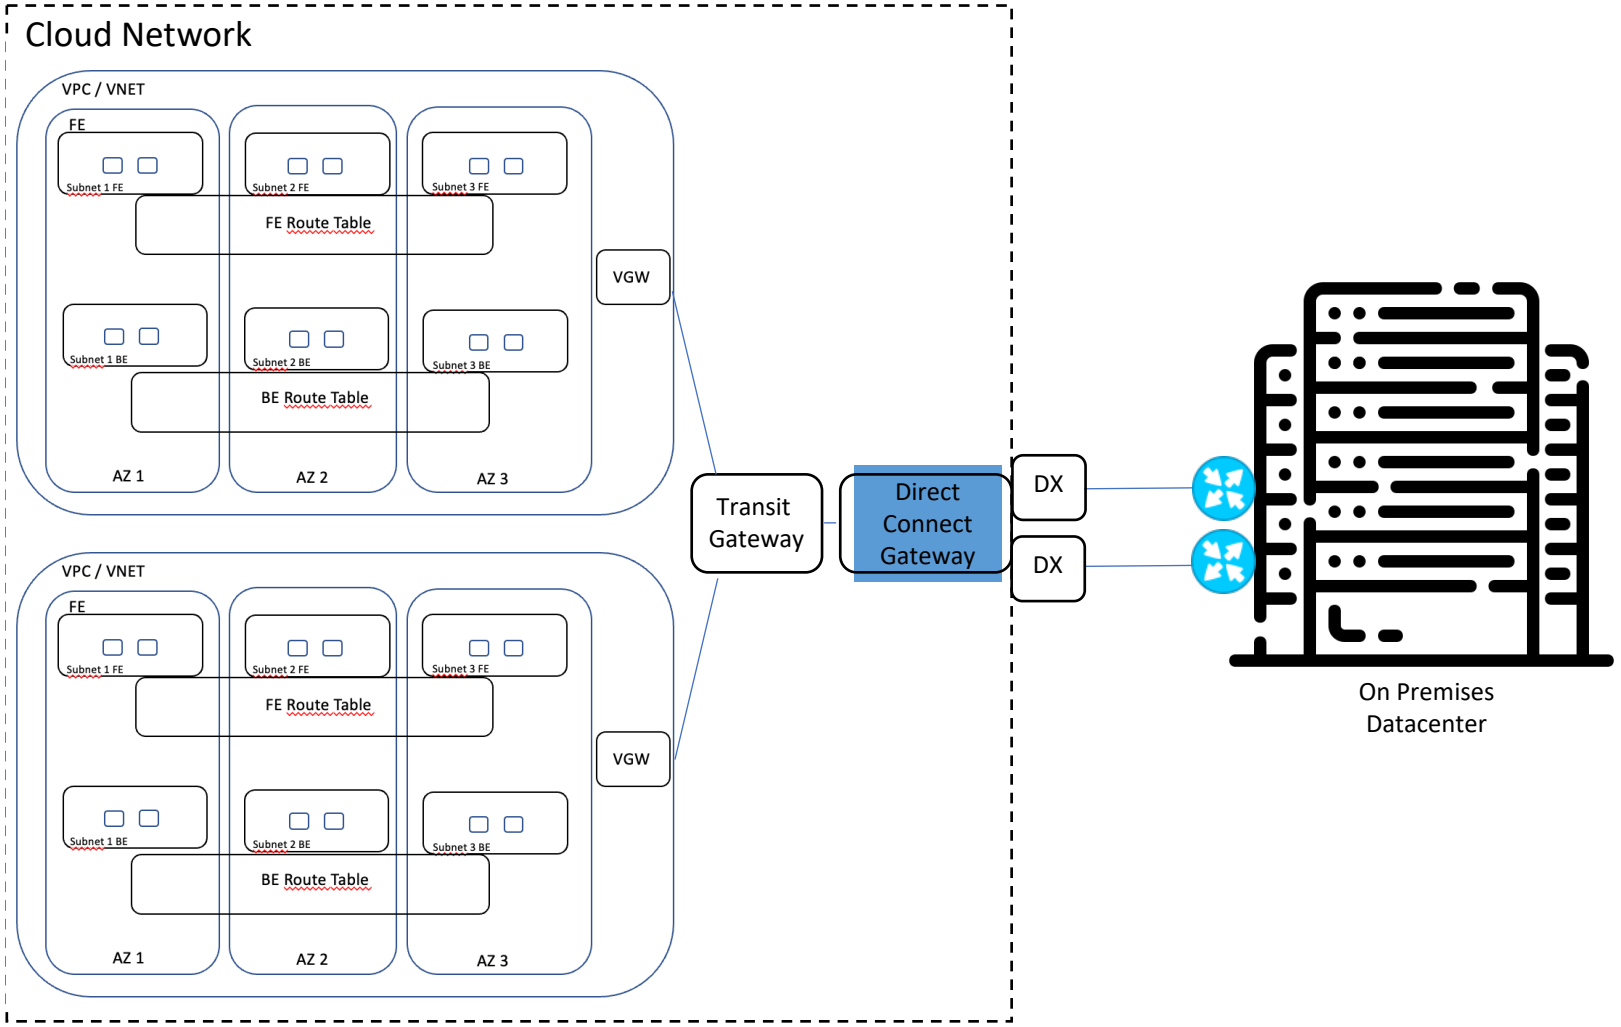

Transit Gateway

. Interconnection Hub

. Per region resource

. Direct connect and VPN support

. Route received from direct connect are prefered

. Static and dynamic bgp routing

. Limits – 10k static routes, 1k dynamic routes

Direct Connect Gateway

. Global resource

. Attached to transit gateway or virtual private gateway

Cloud Network Infrastructure

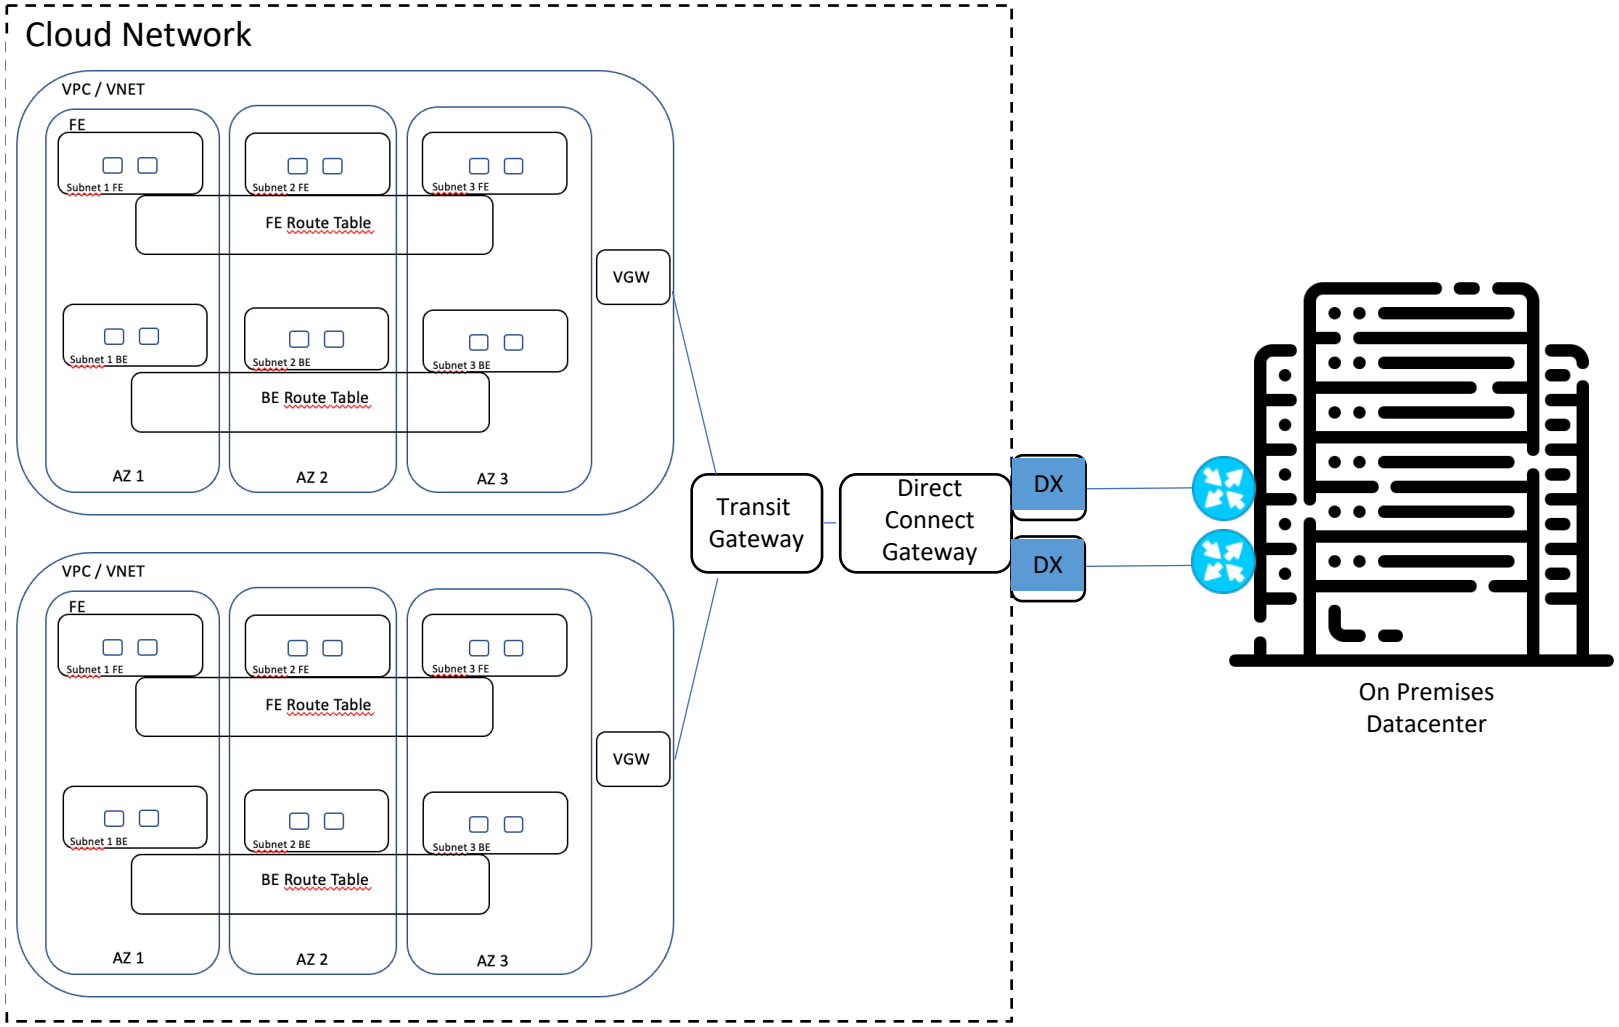

Transit Gateway

. Interconnection Hub
. Per region resource
. Direct connect and VPN support
. Route received from direct connect are prefered
. Static and dynamic bgp routing
. Limits – 10k static routes, 1k dynamic routes

Direct Connect Gateway

. Global resource
. Attached to transit gateway or virtual private gateway

Direct Connect Locations - DX

. Physical location of cloud facility
. Connections can be dedicated or hosted by partners

Cloud Network Infrastructure

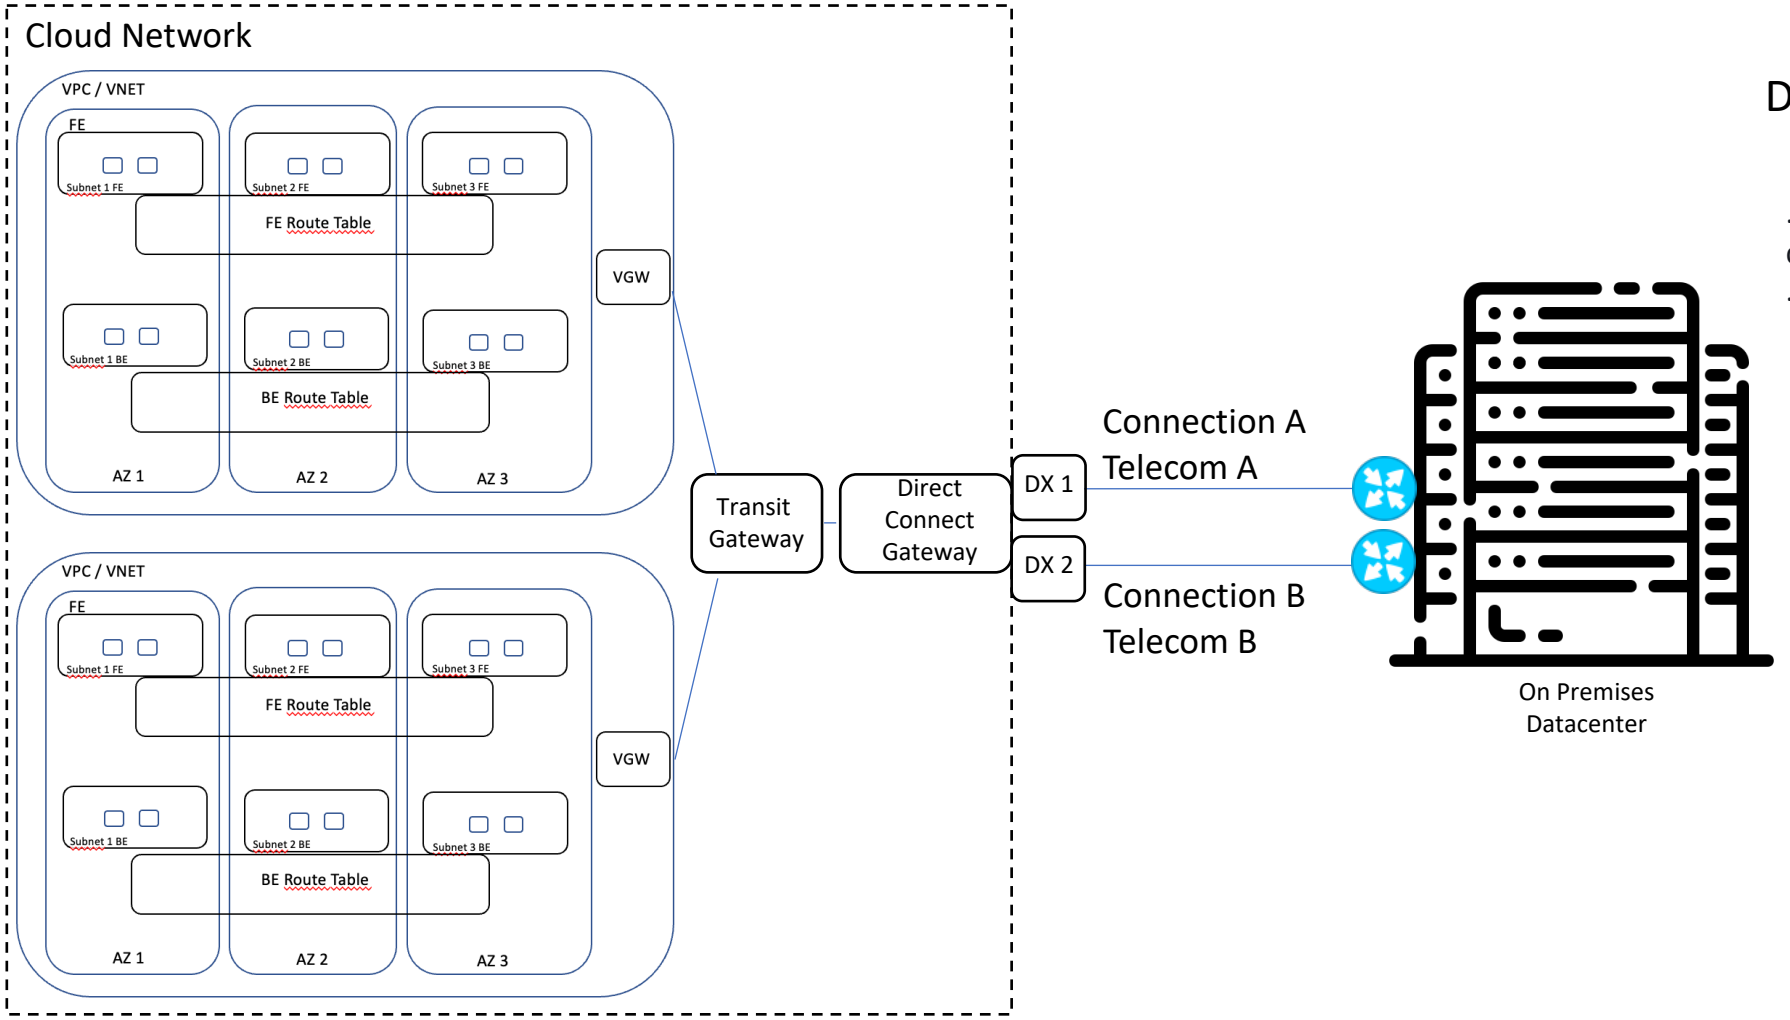

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection

Cloud Network Infrastructure

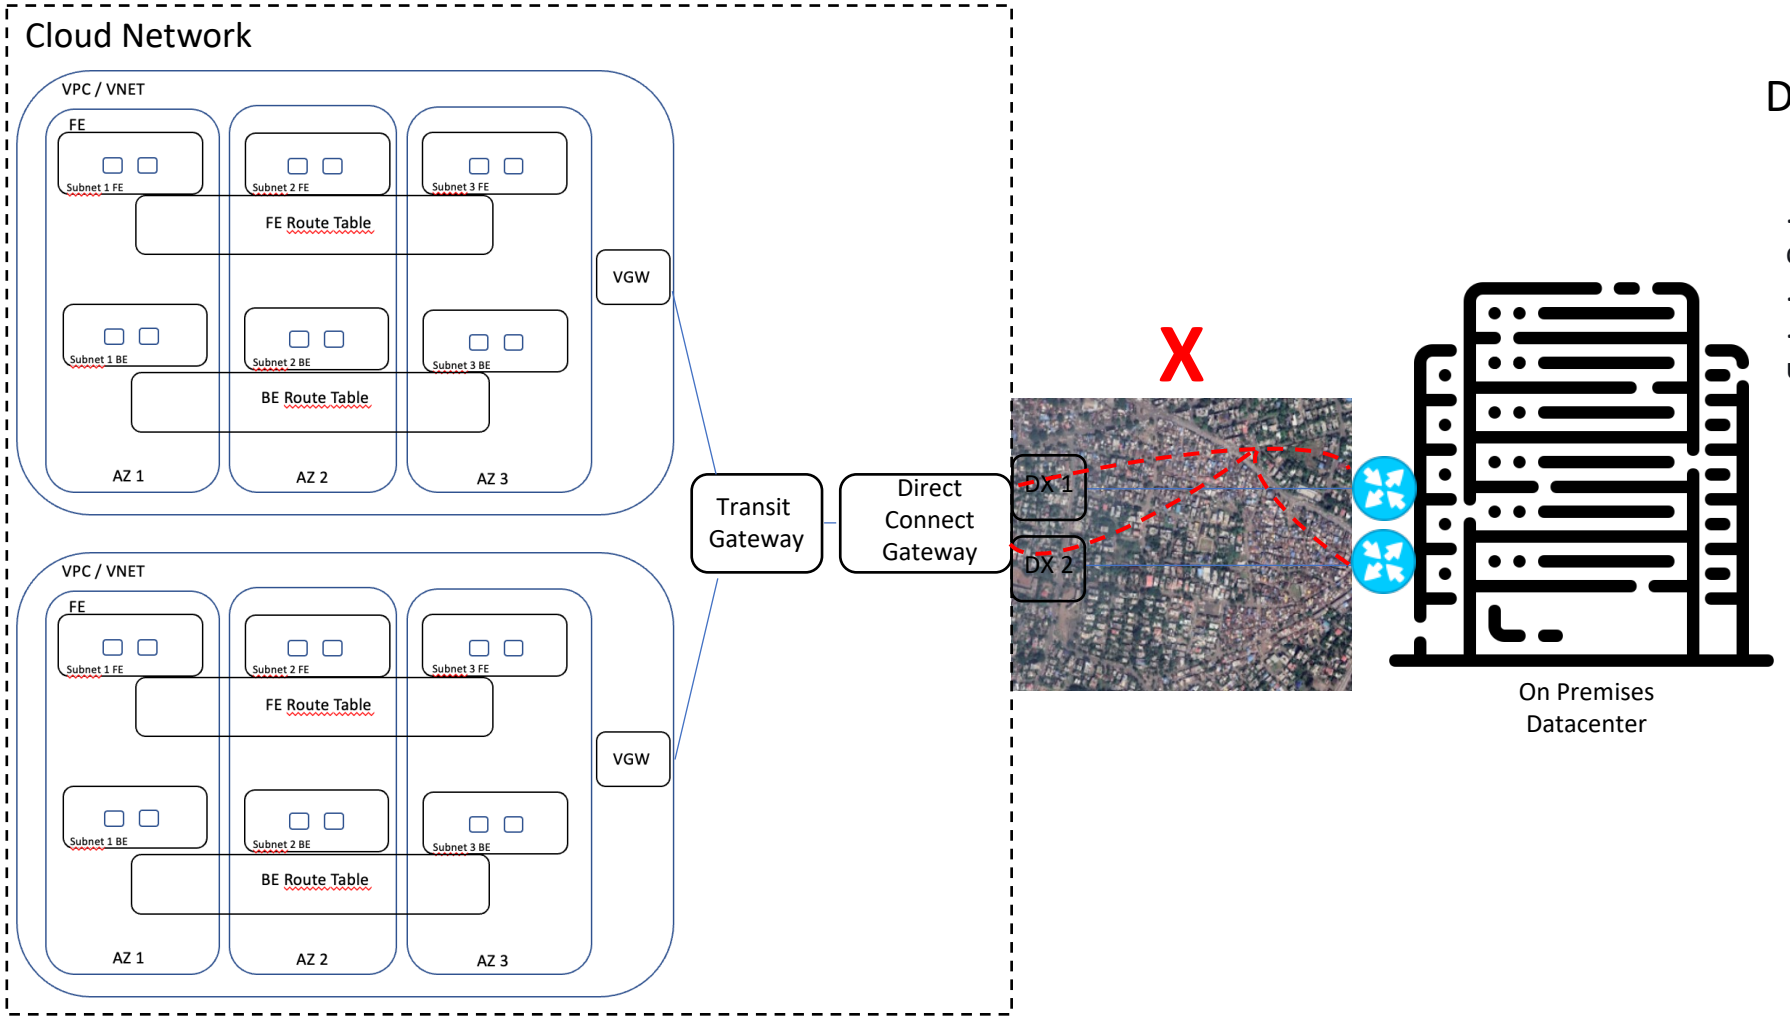

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection
. Analyze the physical path of each connection to not use same facility/last mile

Cloud Network Infrastructure

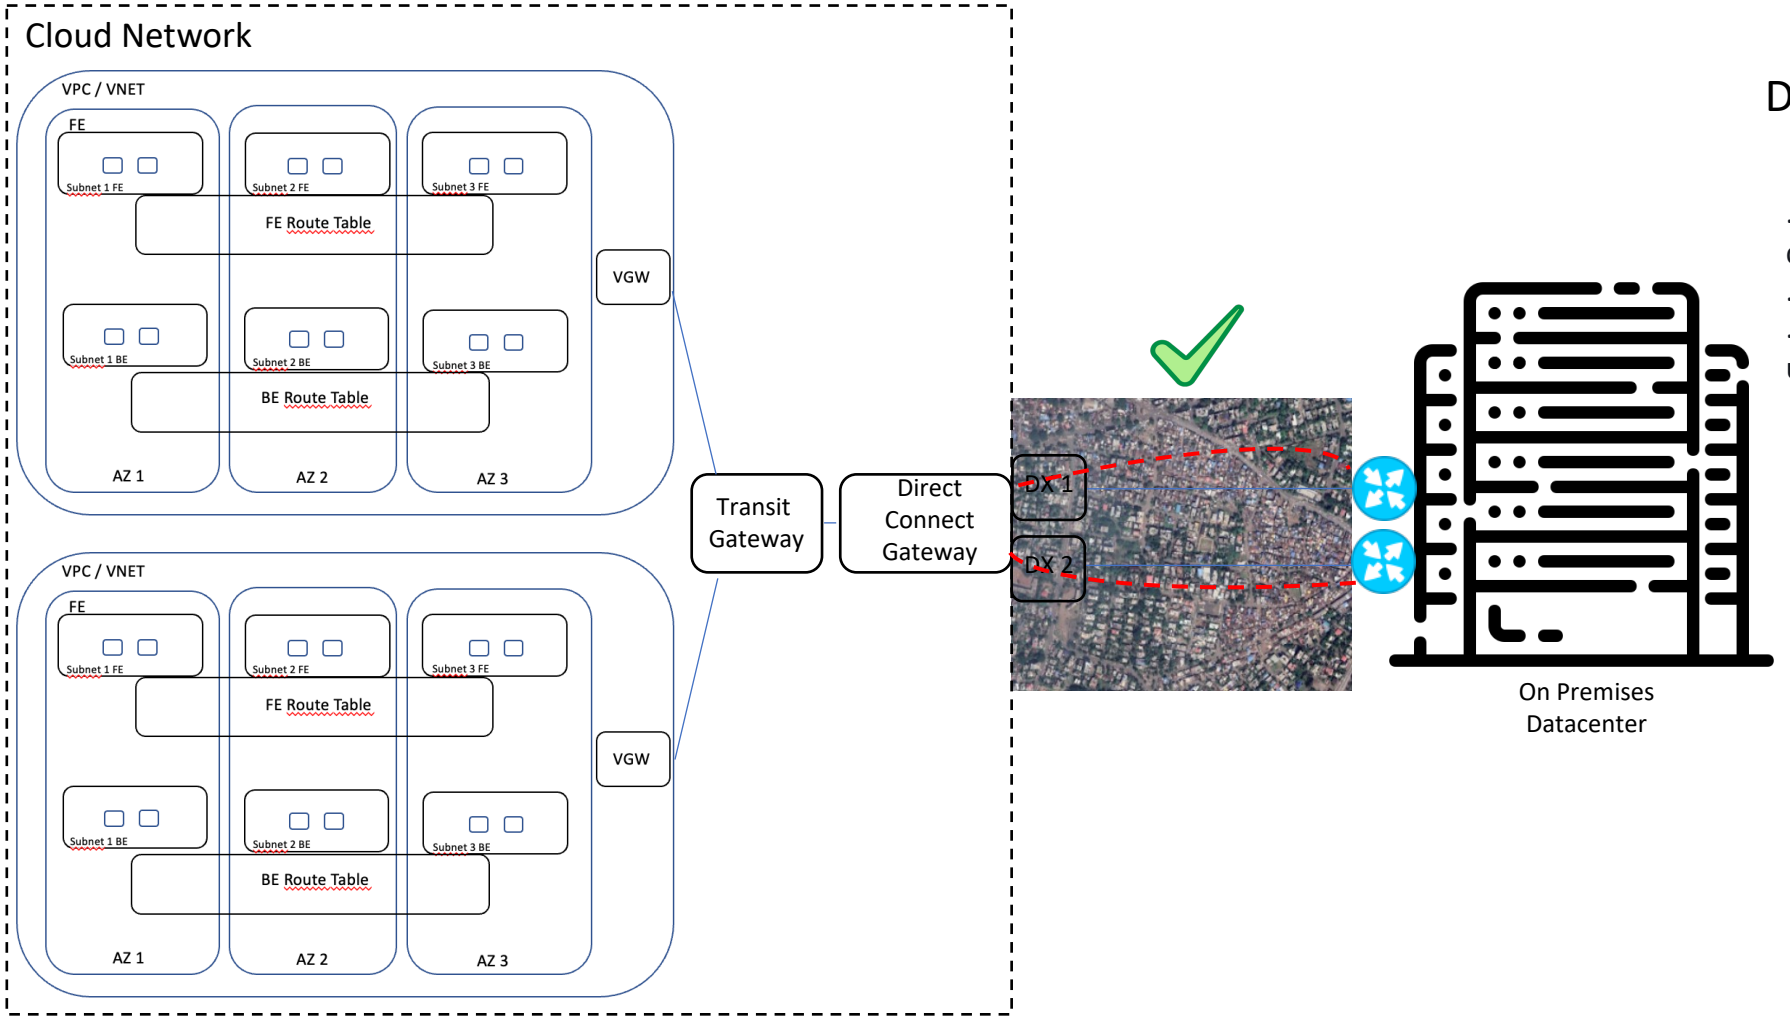

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection
. Analyze the physical path of each connection to not use same facility/last mile

Cloud Network Infrastructure

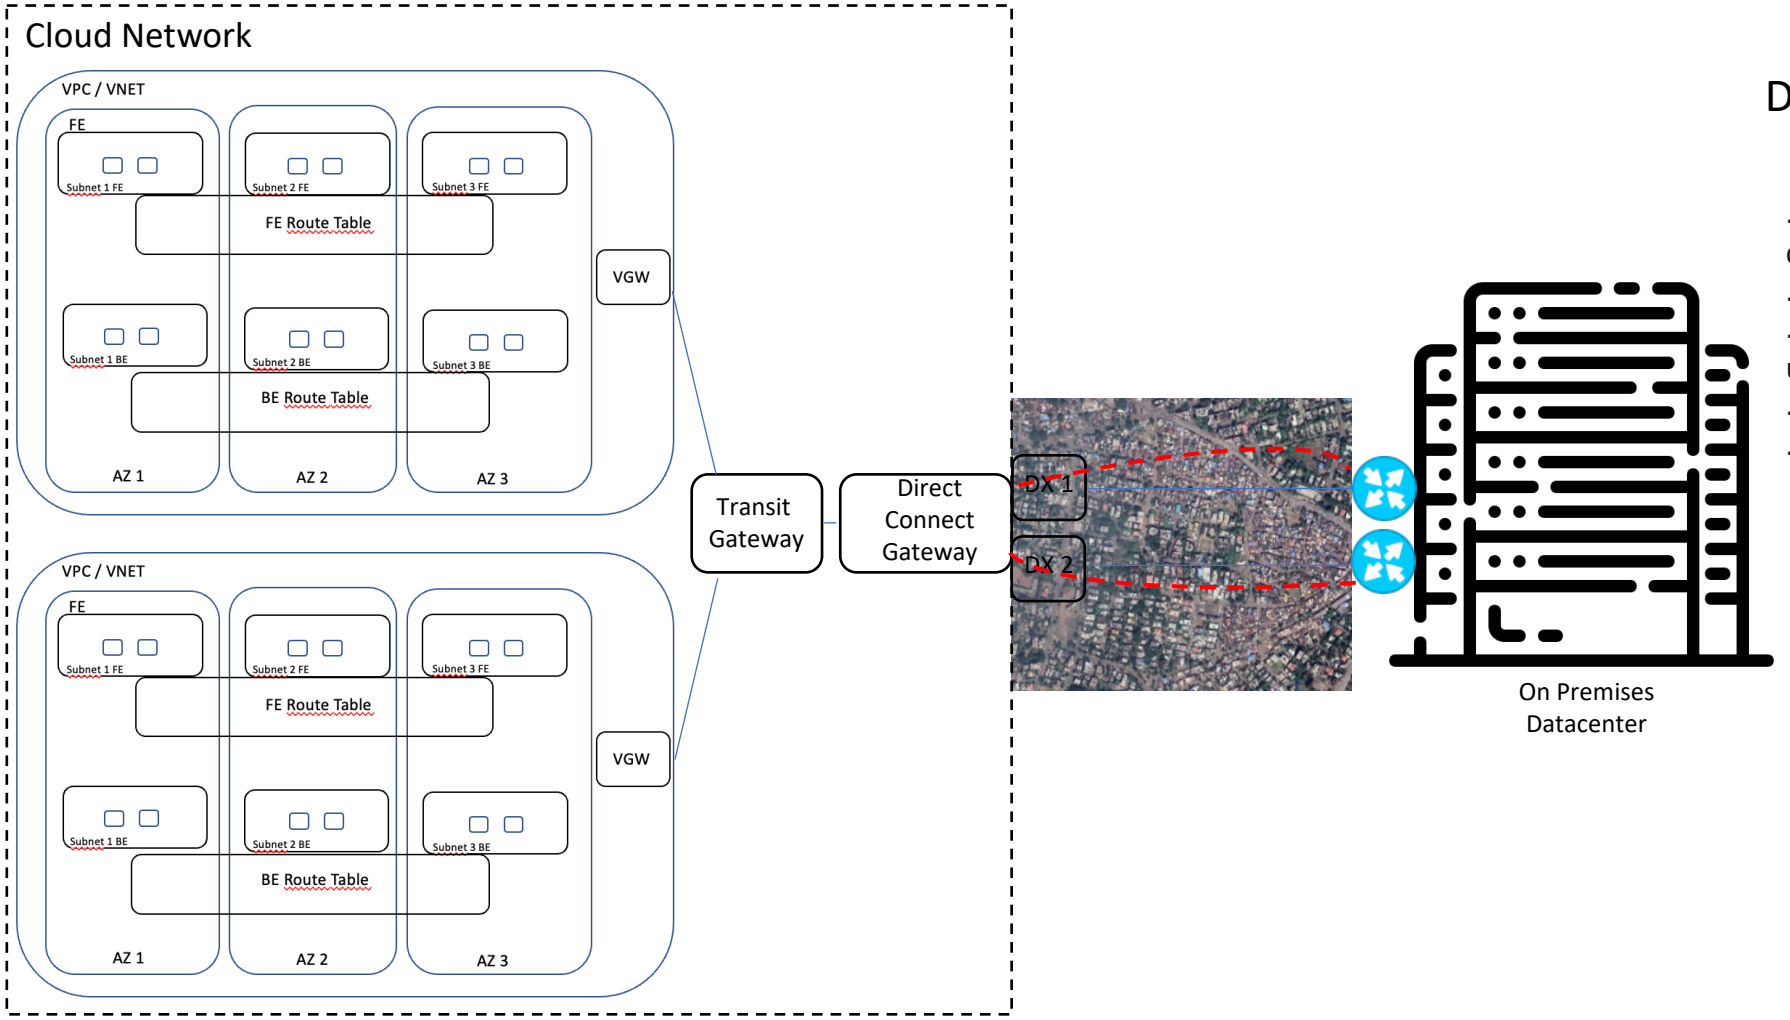

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection
. Analyze the physical path of each connection to not use same facility/last mile
. 1/10/100Gbps for dedicated connections
. Lower than 1Gbps for hosted connections

Cloud Network Infrastructure

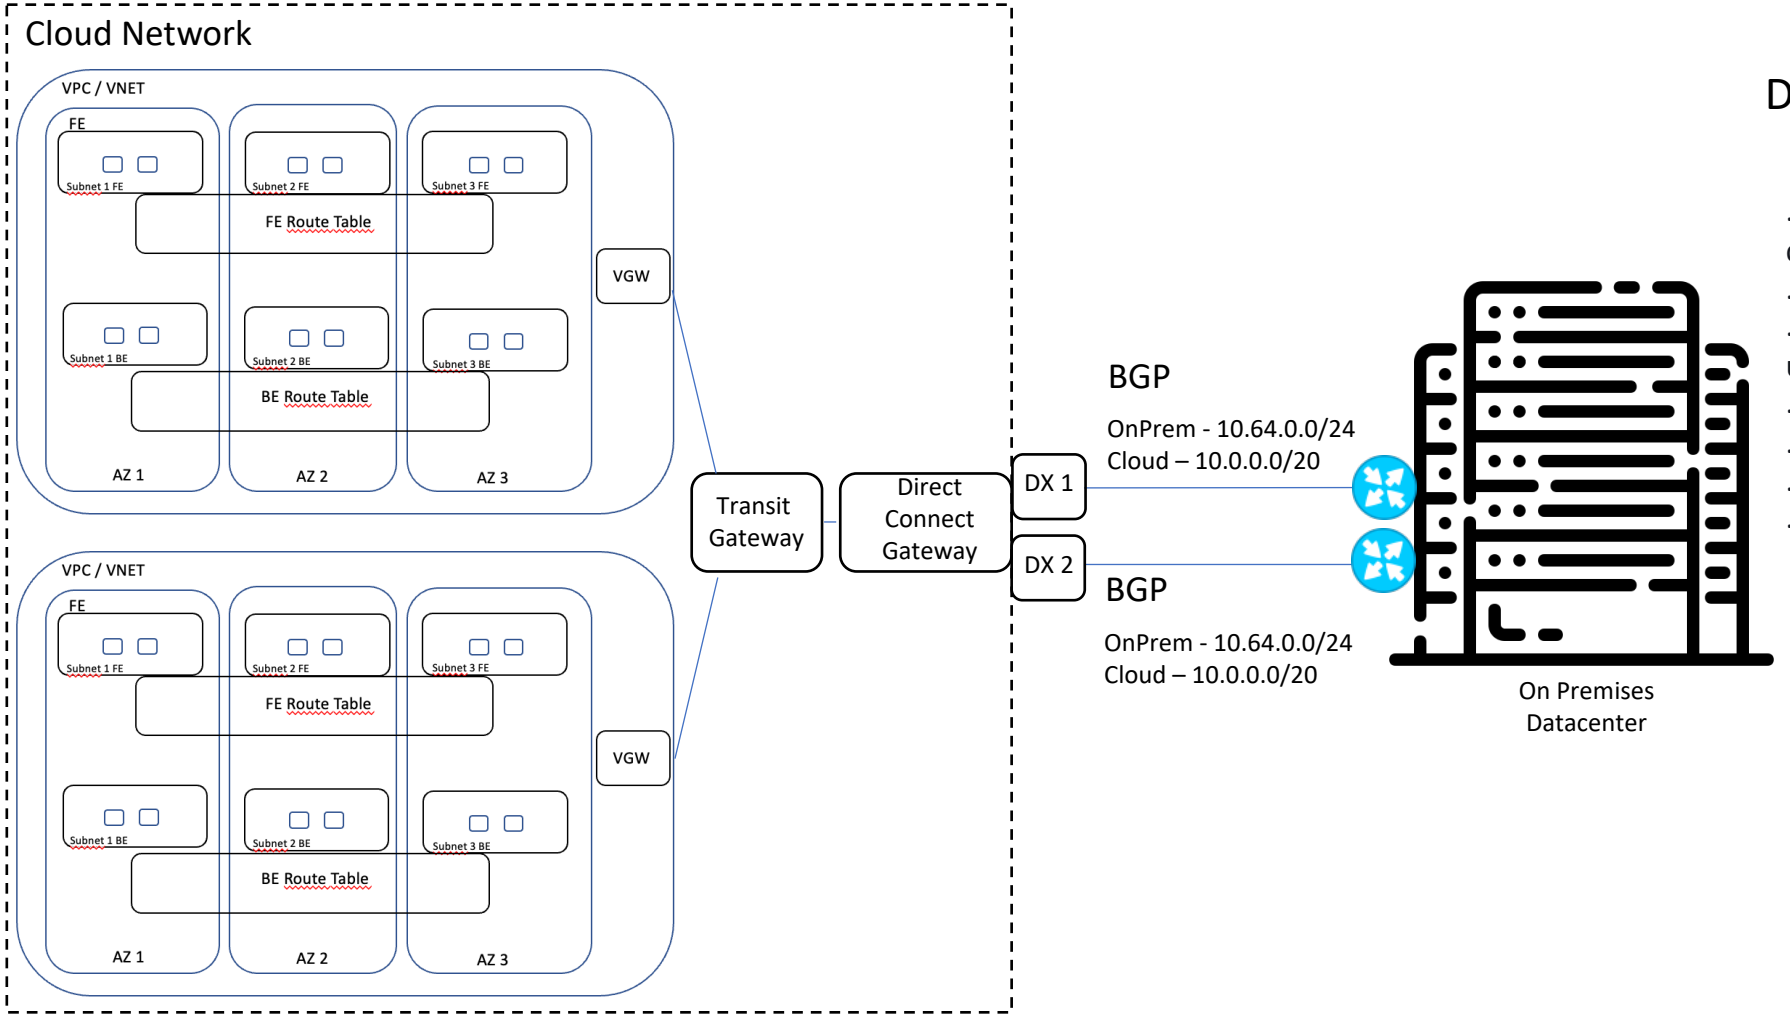

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection
. Analyze the physical path of each connection to not use same facility/last mile
. 1/10/100Gbps for dedicated connections
. Lower than 1Gbps for hosted connections
. BGP – Dynamic routing and convergence
. Can use private ASN or your own

Cloud Network Infrastructure

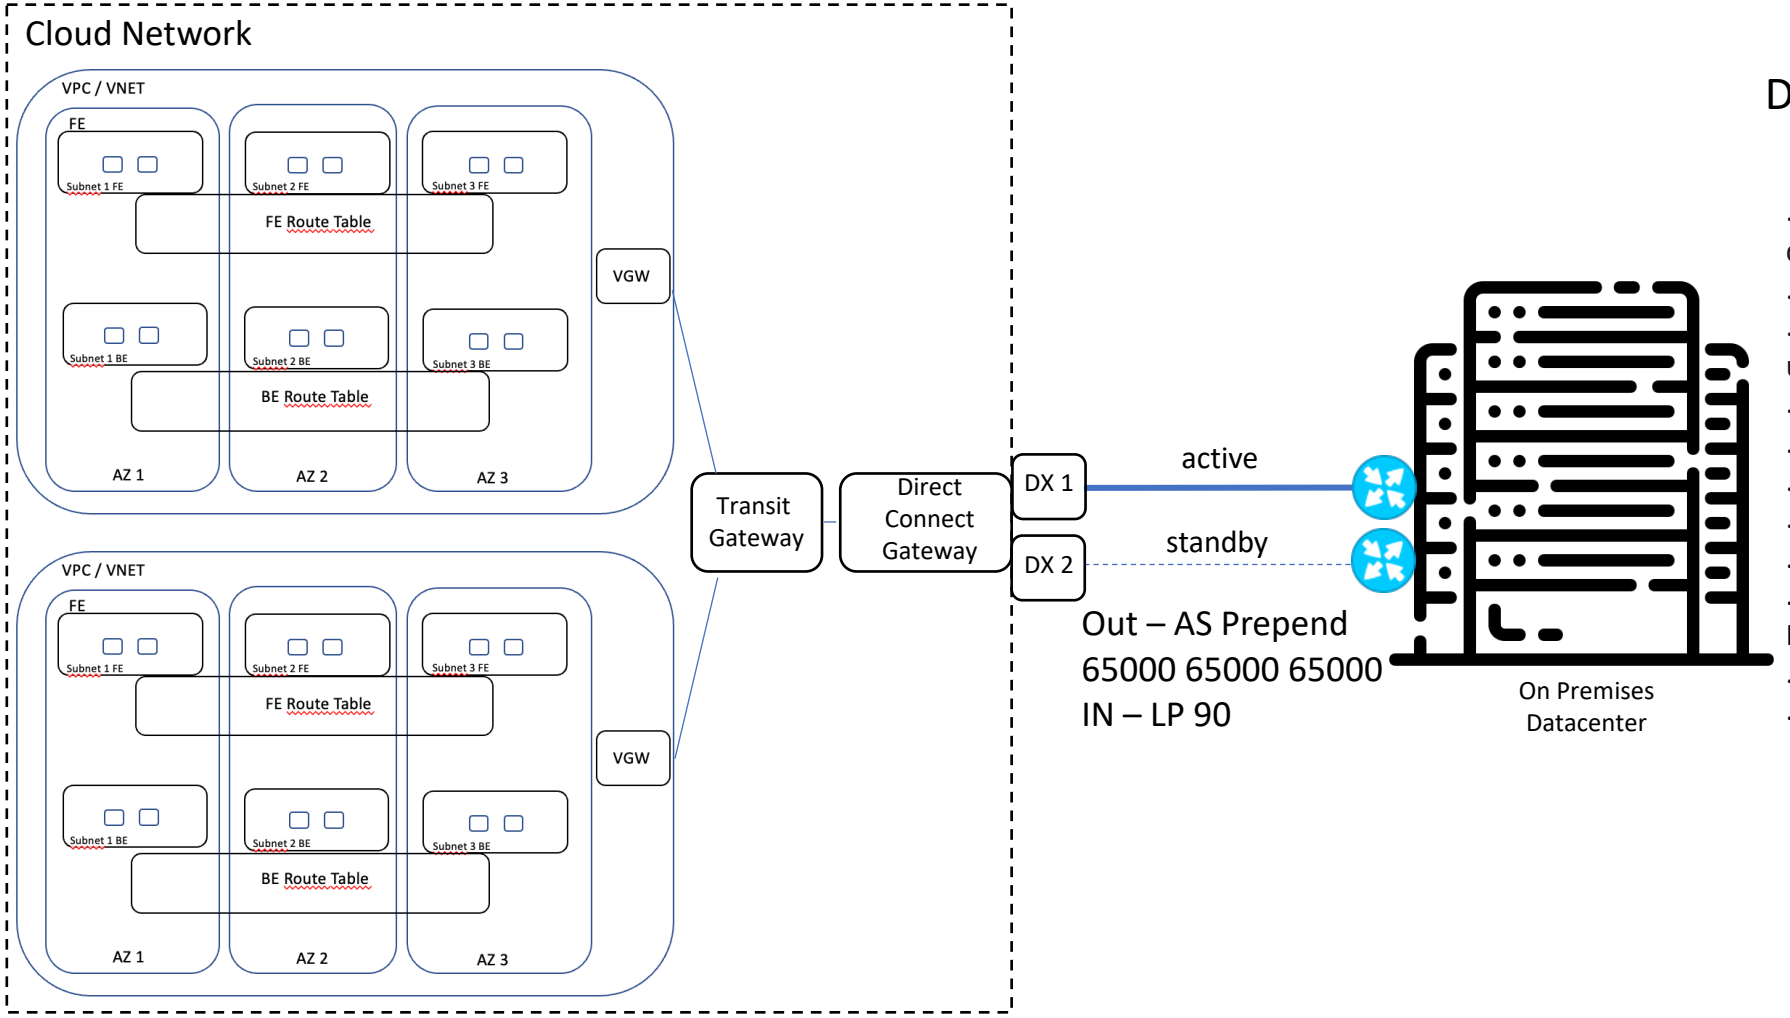

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection
. Analyze the physical path of each connection to not use same facility/last mile
. 1/10/100Gbps for dedicated connections
. Lower than 1Gbps for hosted connections
. BGP – Dynamic routing and convergence
. Can use private ASN or your own
. Cloud provider use ECMP for same route metric
. Active/Standby path can be implemented using AS Prepend and Local Preference
. There are short limits of advertised routes
. Summarization is mandatory to scale

Cloud Network Infrastructure

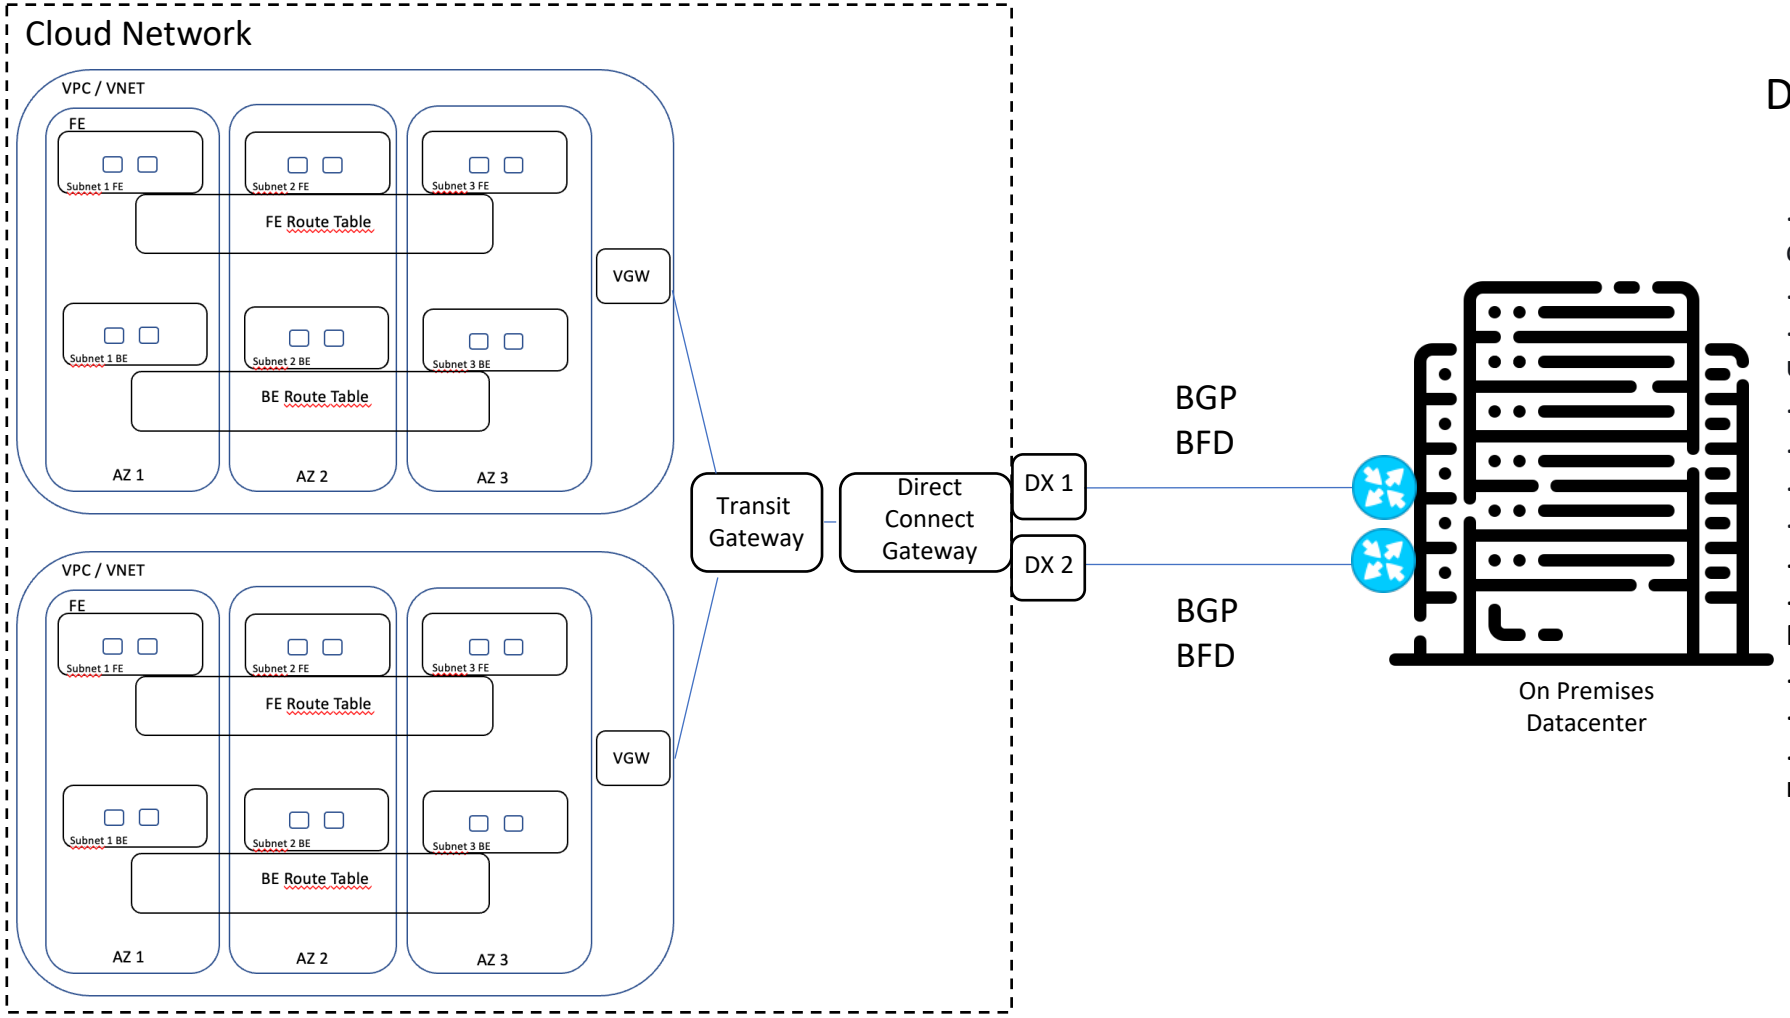

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection
. Analyze the physical path of each connection to not use same facility/last mile
. 1/10/100Gbps for dedicated connections
. Lower than 1Gbps for hosted connections
. BGP – Dynamic routing and convergence
. Can use private ASN or your own
. Cloud provider use ECMP for same route metric
. Active/Standby path can be implemented using AS Prepend and Local Preference
. There are short limits of advertised routes
. Summarization is mandatory to scale
. BFD – Minimize the convergence time to miliseconds

Cloud Network Infrastructure

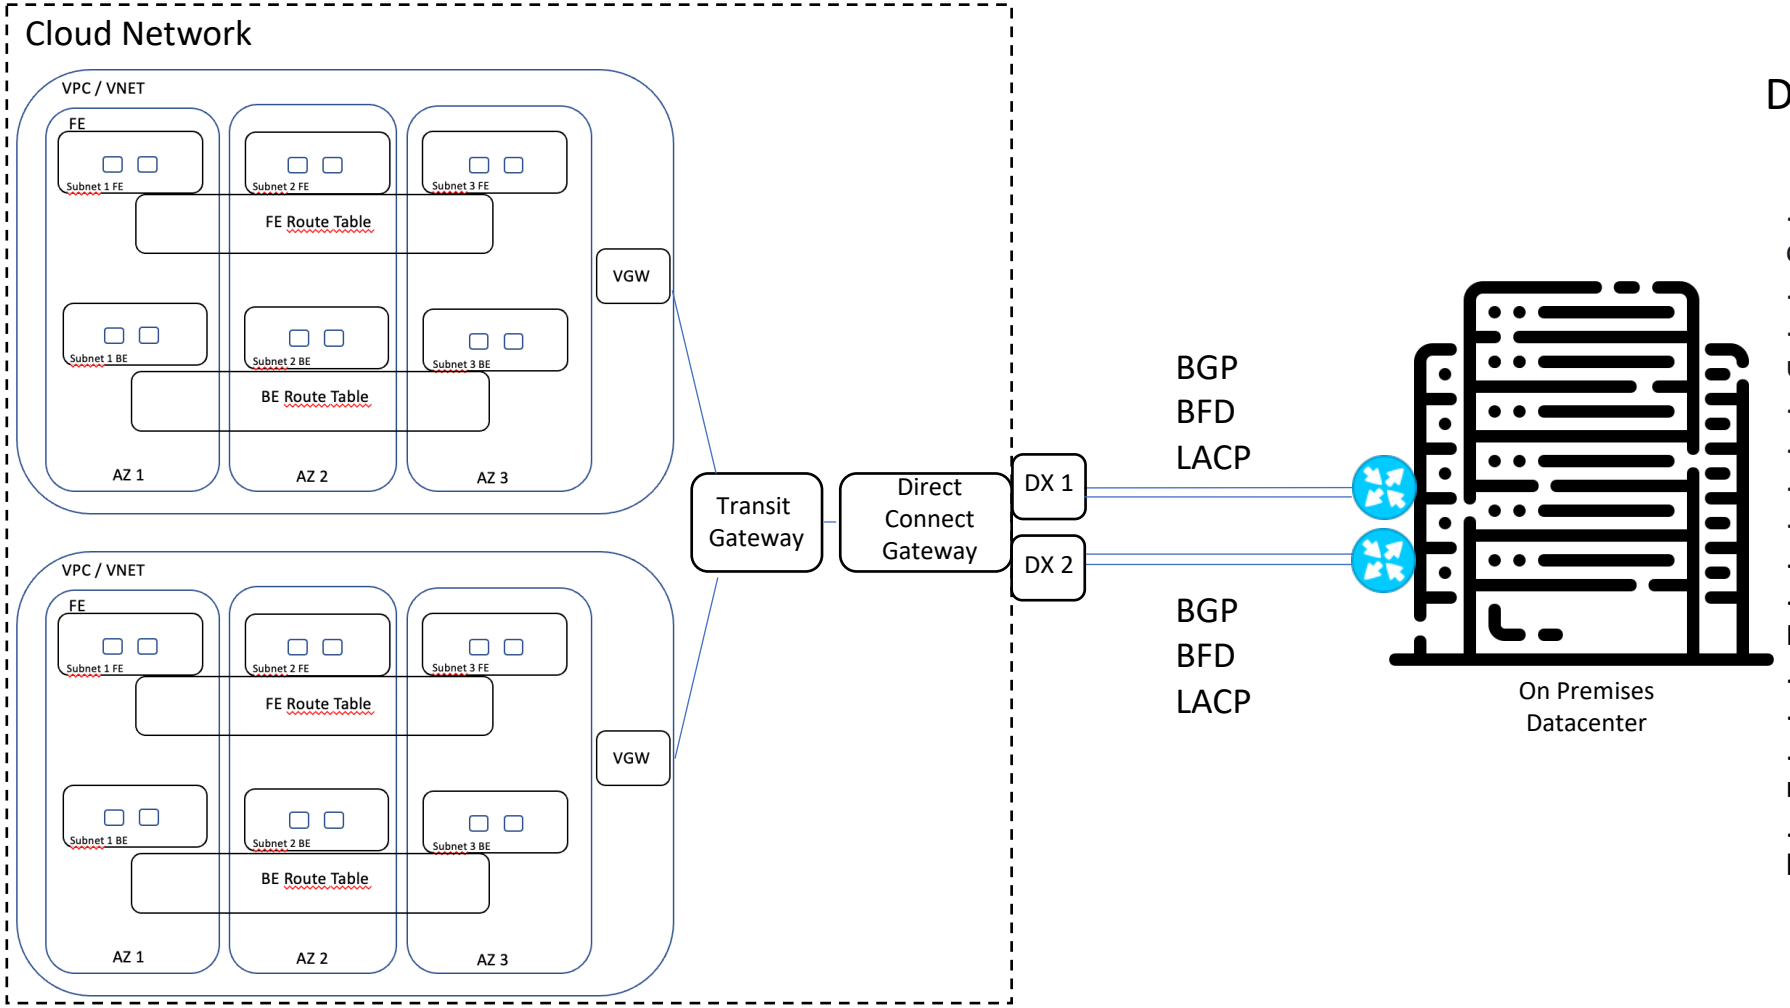

Direct Connect – Private Connections

. To improve high availability, connect to at least two different direct connect locations
. Use different telecom operators per connection
. Analyze the physical path of each connection to not use same facility/last mile
. 1/10/100Gbps for dedicated connections
. Lower than 1Gbps for hosted connections
. BGP – Dynamic routing and convergence
. Can use private ASN or your own
. Cloud provider use ECMP for same route metric
. Active/Standby path can be implemented using AS Prepend and Local Preference
. There are short limits of advertised routes
. Summarization is mandatory to scale
. BFD – Minimize the convergence time to miliseconds
. LACP – aggregate links and high availability per DX location

Cloud Network Infrastructure

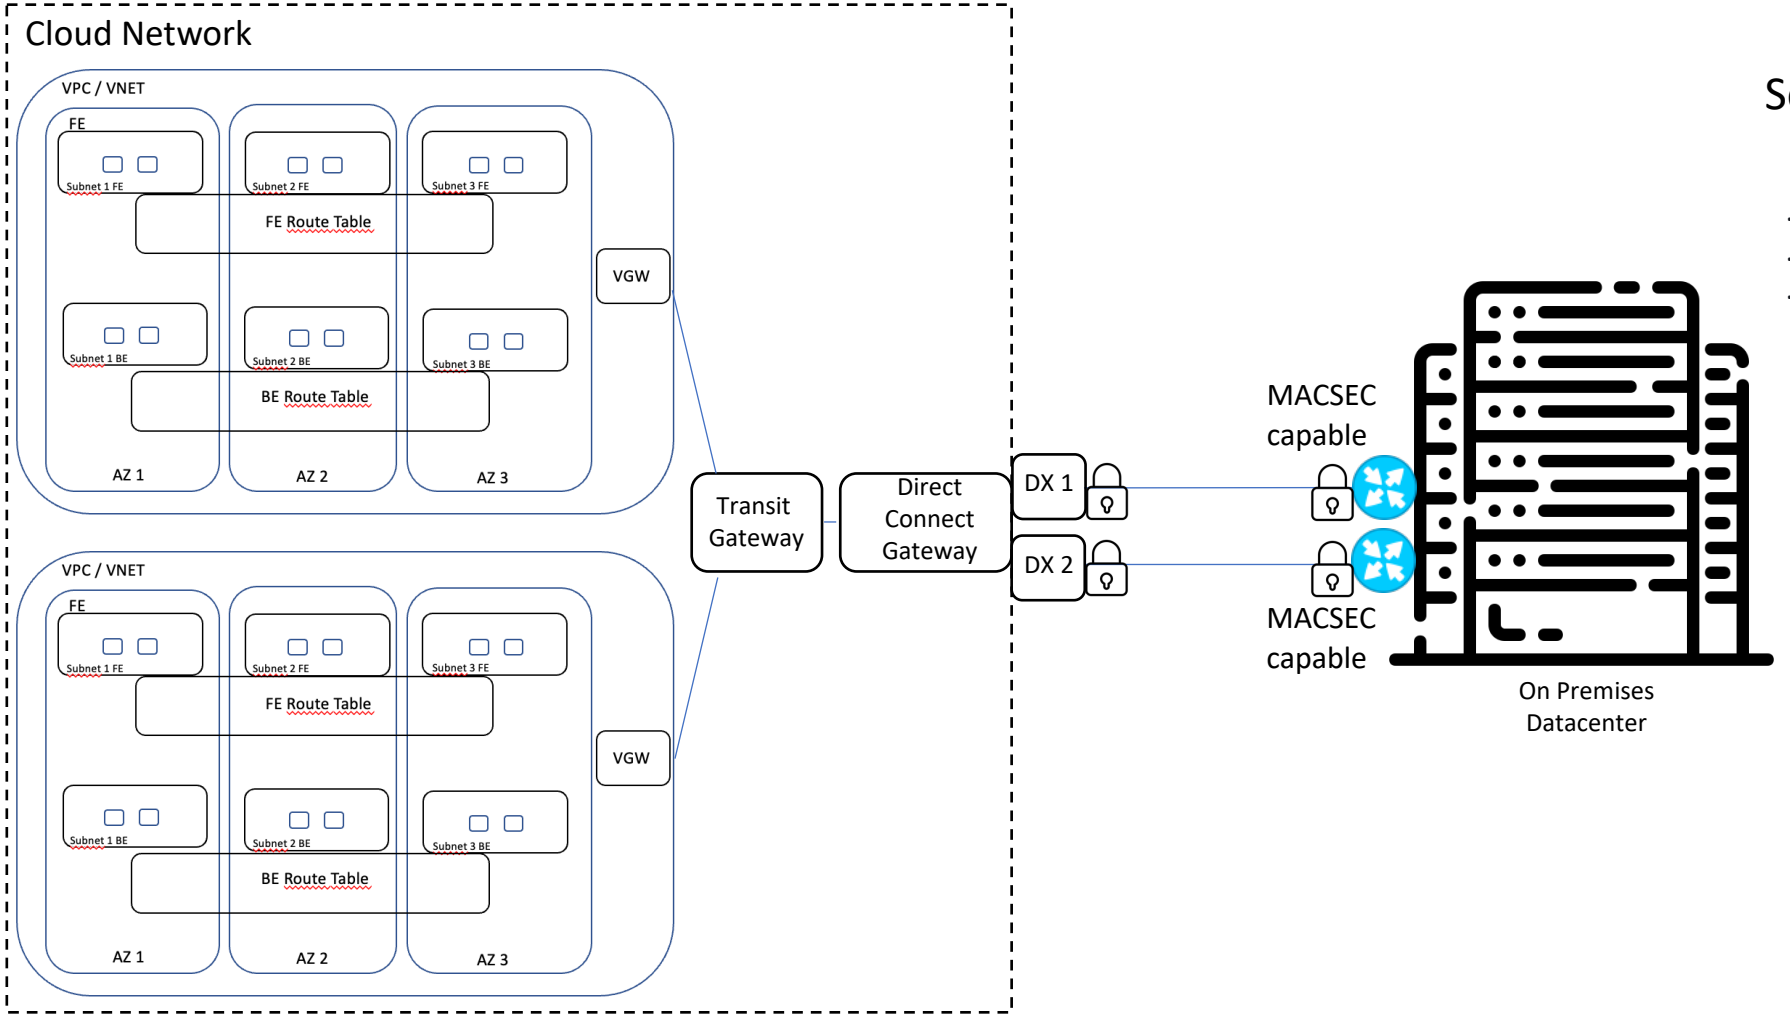

Secure Data in Transit - MACSEC

. MACSEC – L2 encryption
. On prem device must support in hardware
. Entire path must accept L2 MACSEC datagram

Cloud Network Infrastructure

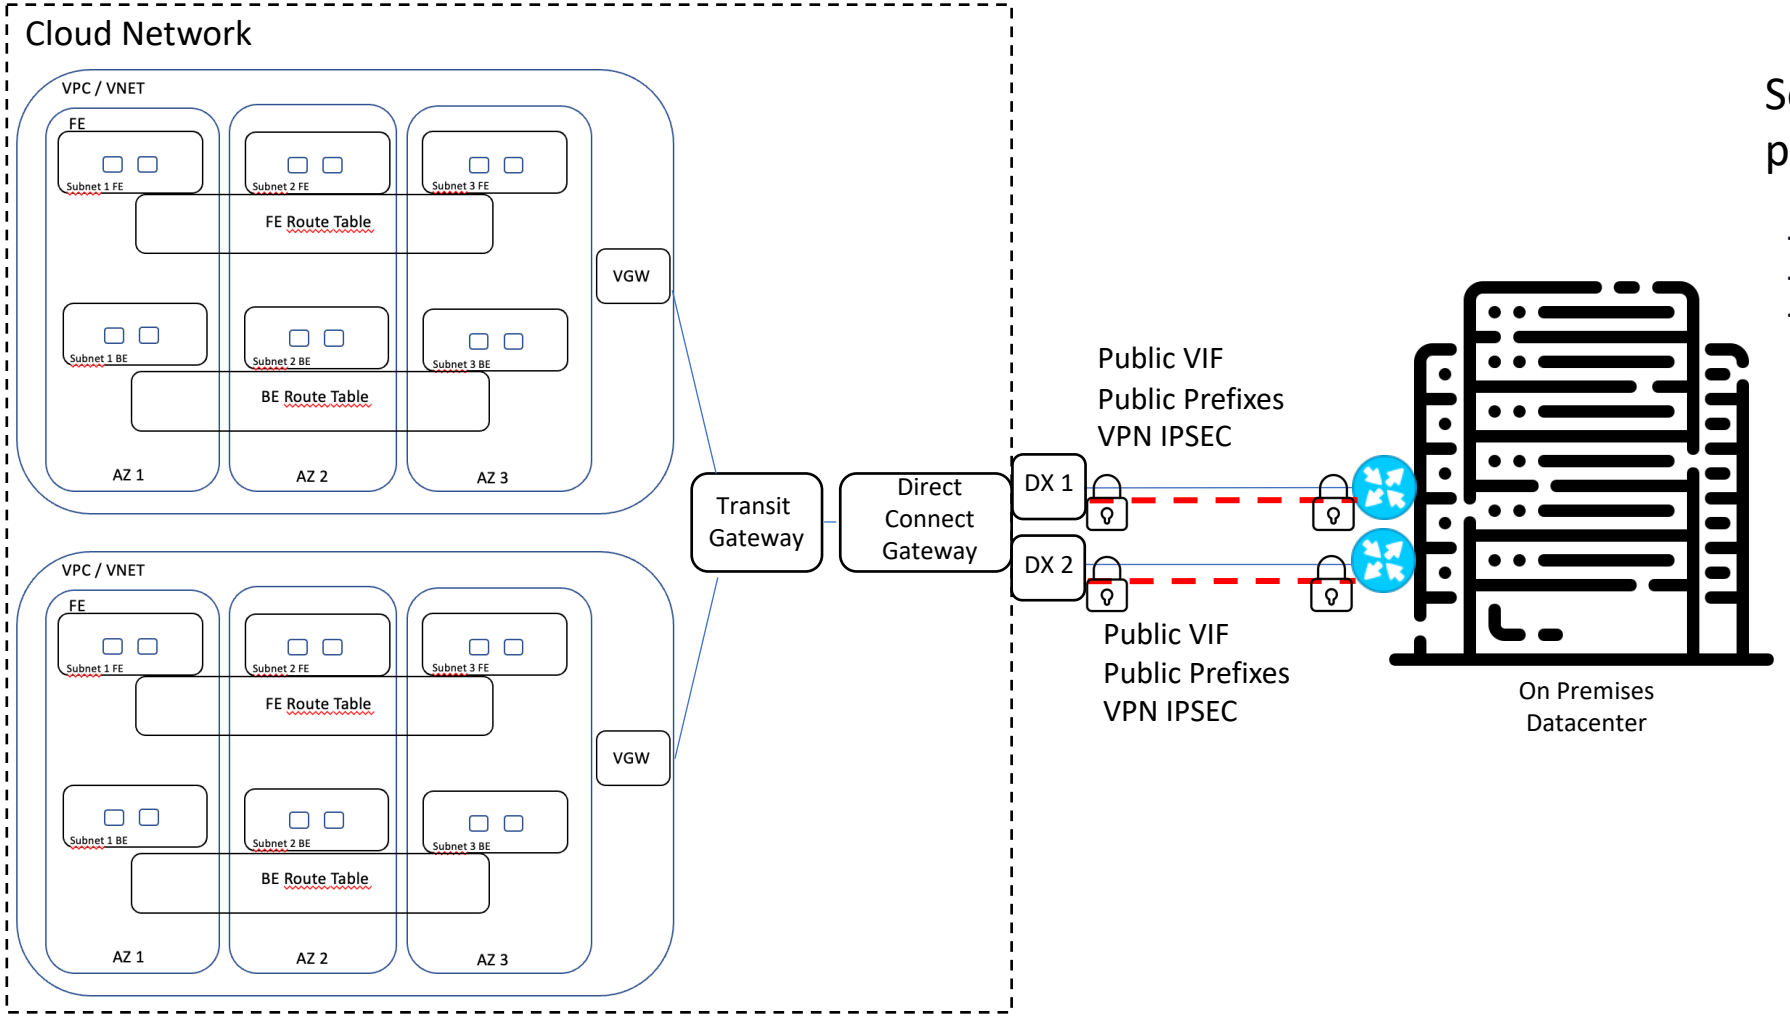

Secure Data in Transit – VPN over private connections

. VPN IPSEC over direct connect - public vif
. Receive public prefixes from cloud network
. On prem must have public ASN/Prefixes

Cloud Network Infrastructure

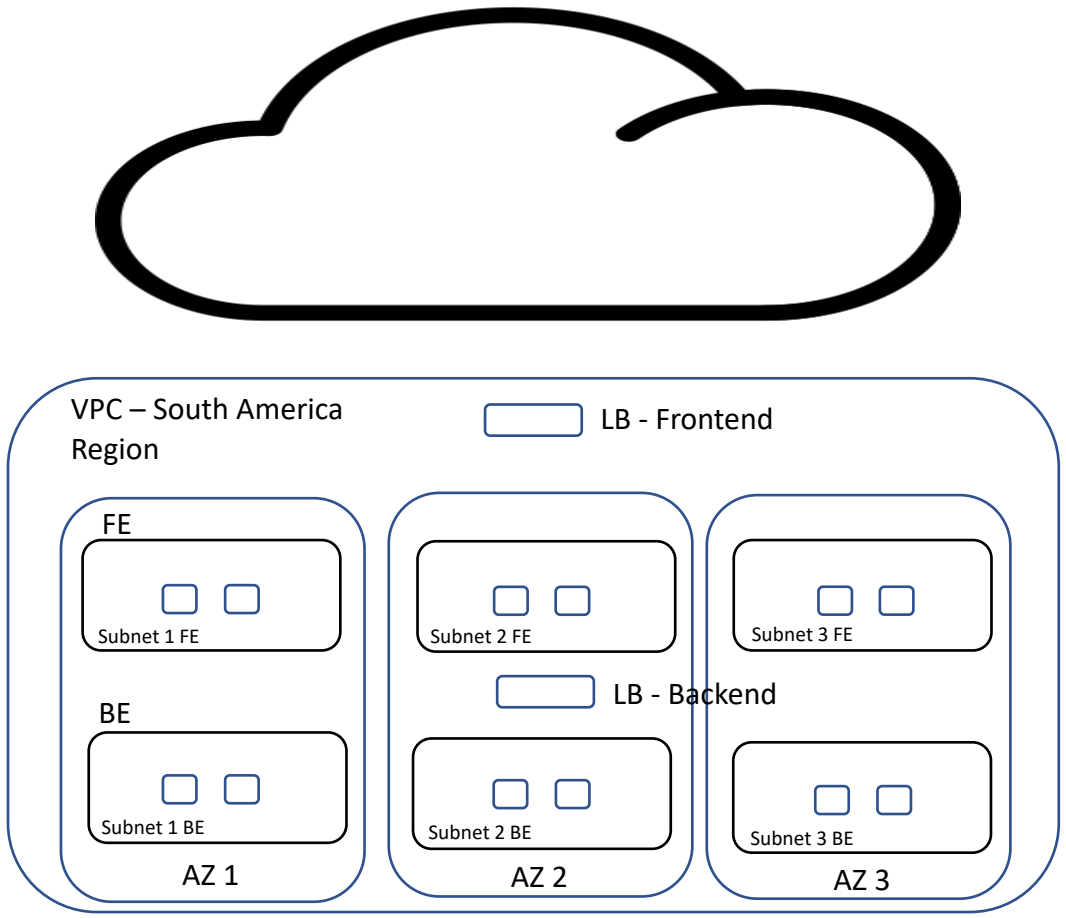

Internet Applications

. Applications generally are hosted in one region
. Latency increased to applications accessed from users globally

Cloud Network Infrastructure

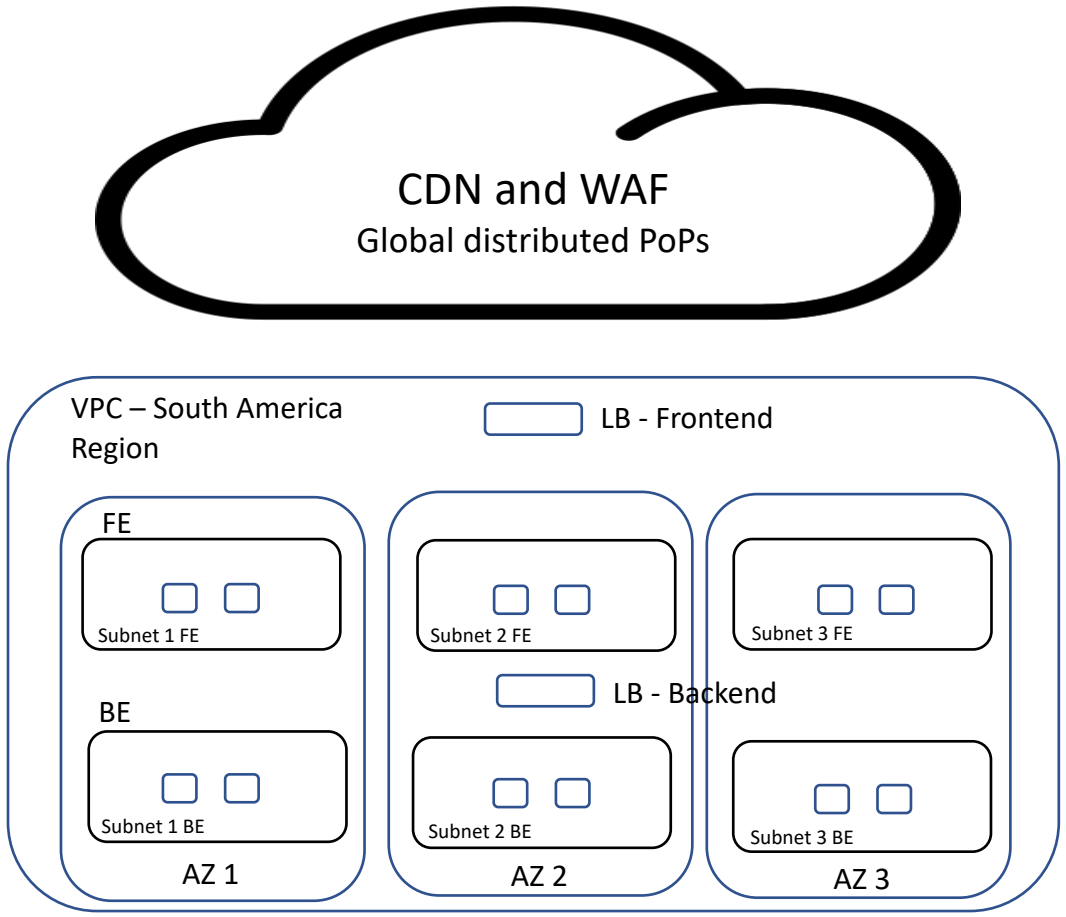

Internet Applications

. Applications generally are hosted in one region only
. Latency increased to applications accessed from users globally

. CDN – Content Delivery Network
  . Delivery the content from the PoP closest to the user
  . Cache – decrease the hit on origin servers, accelerate access
  . DDoS Protection
  . Can be used the cloud provider or other CDN company

. WAF – Web Application Firewall
  . Protect the infrastrucure from application layer attacks

Conclusion

- Even in cloud, network architecture must follow best practices to be high available and secure
- Latency between Cloud x On Prem can be a problem in Latin America
- The network engineer must be close to the application requirements to plan the right architecture – from L1 to L7
- Automation mindset starting from day one – reliability and scale

Author

Wilson Lopes

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

- 18 years of experience in the network engineering and unix systems administration areas. I started my career at NIC.BR, working after in telecomunications companies and financial institutions of Brazil.
- Dedicated in the last 4 years in the design and implementation of hybrid datacenter infrastructures, cloud datacenter migration and network automation.