ORACLE

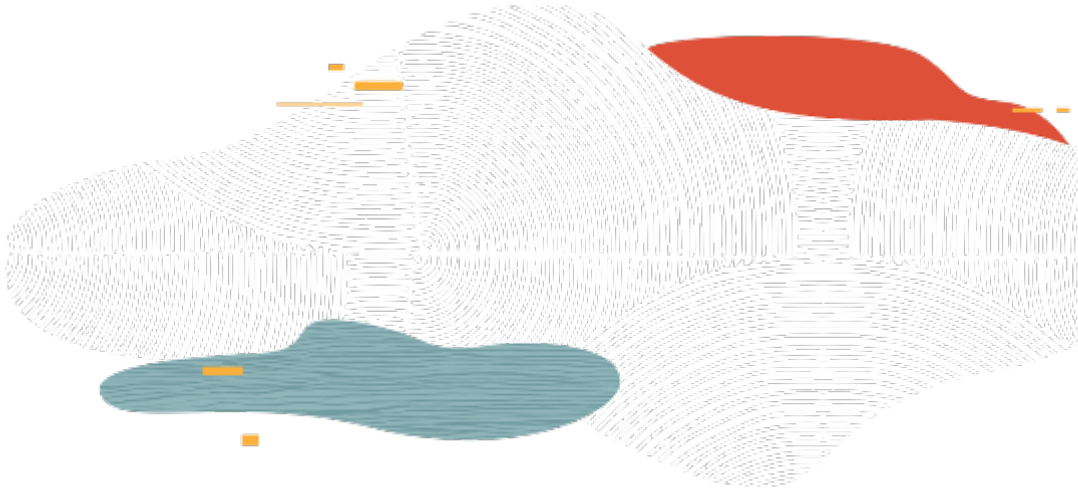

Virtual Cloud Network

Level 100

Rohit Rahi
Oracle Cloud Infrastructure
November, 2019

Safe harbor statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions.

The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Objectives

- Virtual Cloud Network (VCN) basics
- IP addresses
- Gateways and Routing
- Peering
- Transit Routing
- Security
- Putting it all together!

ORACLE

CIDR

CIDR Basics

- CIDR (classless inter-domain routing) notation
  - IP addresses are described as consisting of two groups of bits in the address: the most significant bits are the network prefix, which identifies a whole network (or subnet), and the least significant set forms the host identifier, which specifies a particular interface of a host on that network
  - An IP address has two components, the network address and the host address: <network> <host>
  - A subnet mask separates the IP address into the network and host addresses (<network><host>). Subnetting further divides the host part of an IP address into a subnet and host address (<network><subnet><host>)
  - Subnet Mask is made by setting network bits to all "1"s and setting host bits to all "0"s. Within a given network, two host addresses cannot be assigned to hosts. The "0" address is assigned a network address and "255" is assigned to a broadcast address
  - Notation is constructed from an IP address, a '/' character, and a decimal number. xxx.xxx.xxx.xxx/n, where n is the number of bits used for subnet mask. E.g. 192.168.1.0/24
  - Examples of commonly used netmasks for classed networks are 8-bits (Class A), 16-bits (Class B) and 24-bits (Class C)

CIDR Basics

192.168.1.0/24 would equate to IP range: 192.168.1.0 – 192.168.1.255

- 128 64 32 16 8 4 2 1 -> 2^7 2^6 2^5 2^4 2^3 2^2 2^1 2^0
- 192 is represented as 1 1 0 0 0 0 0 0

192.168.1.0	1 1 0 0 0 0 0 0	1 0 1 0 1 0 0 0	0 0 0 0 0 0 0 1	0 0 0 0 0 0 0 0
/24 subnet mask	1 1 1 1 1 1 1 1	1 1 1 1 1 1 1 1	1 1 1 1 1 1 1 1	0 0 0 0 0 0 0 0
Logical AND	1 1 0 0 0 0 0 0	1 0 1 0 1 0 0 0	0 0 0 0 0 0 0 1	0 0 0 0 0 0 0 0

192.168.1.0/27 would equate to IP range: 192.168.1.0 – 192.168.1.31

- Now same network divided in 8 subnets with 32 hosts each due to the /27 mask (255.255.255.224)

192.168.1.0	1 1 0 0 0 0 0 0	1 0 1 0 1 0 0 0	0 0 0 0 0 0 0 1	0 0 0 0 0 0 0 0
/27 subnet mask	1 1 1 1 1 1 1 1	1 1 1 1 1 1 1 1	1 1 1 1 1 1 1 1	1 1 1 0 0 0 0 0
Logical AND	1 1 0 0 0 0 0 0	1 0 1 0 1 0 0 0	0 0 0 0 0 0 0 1	0 0 0 0 0 0 0 0

- Subnets – 2 x 2 x 2 = 8. Hosts – 2 x 2 x 2 x 2 x 2 = 32
- Subnetworks – 192.168.1.0/27, 192.168.1.32/27, 192.168.1.64/27…

ORACLE

Virtual Cloud Network Intro

Virtual Cloud Network (VCN)

- A private network that you set up in the Oracle data centers, with firewall rules and specific types of communication gateways that you can choose to use
- A VCN covers a single, contiguous IPv4 CIDR block of your choice
- A VCN resides within a single region

IP address range for your VCN

Avoid IP ranges that overlap with other on-premises or other cloud networks

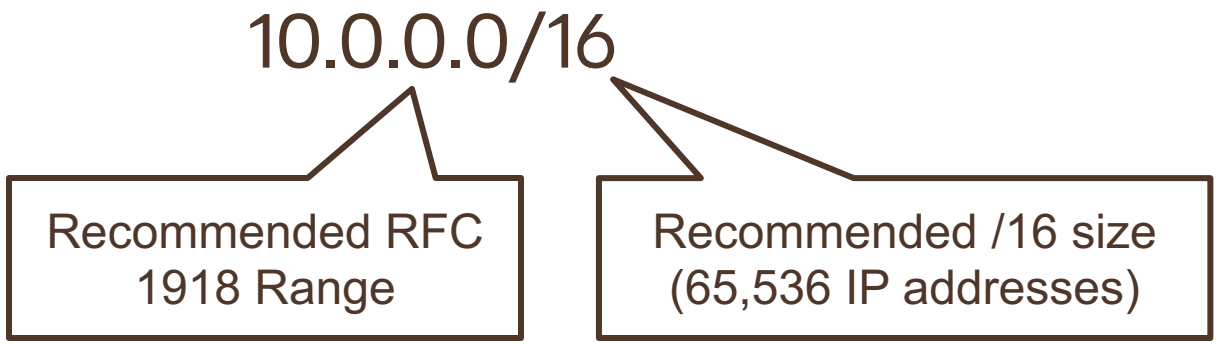

- Use private IP address ranges specified in RFC 1918 (10.0.0.0/8, 172.16/12, 192.168/16)
- Allowable OCI VCN size range is from /16 to /30
- VCN reserves the first two IP addresses and the last one in each subnet's CIDR

Subnet

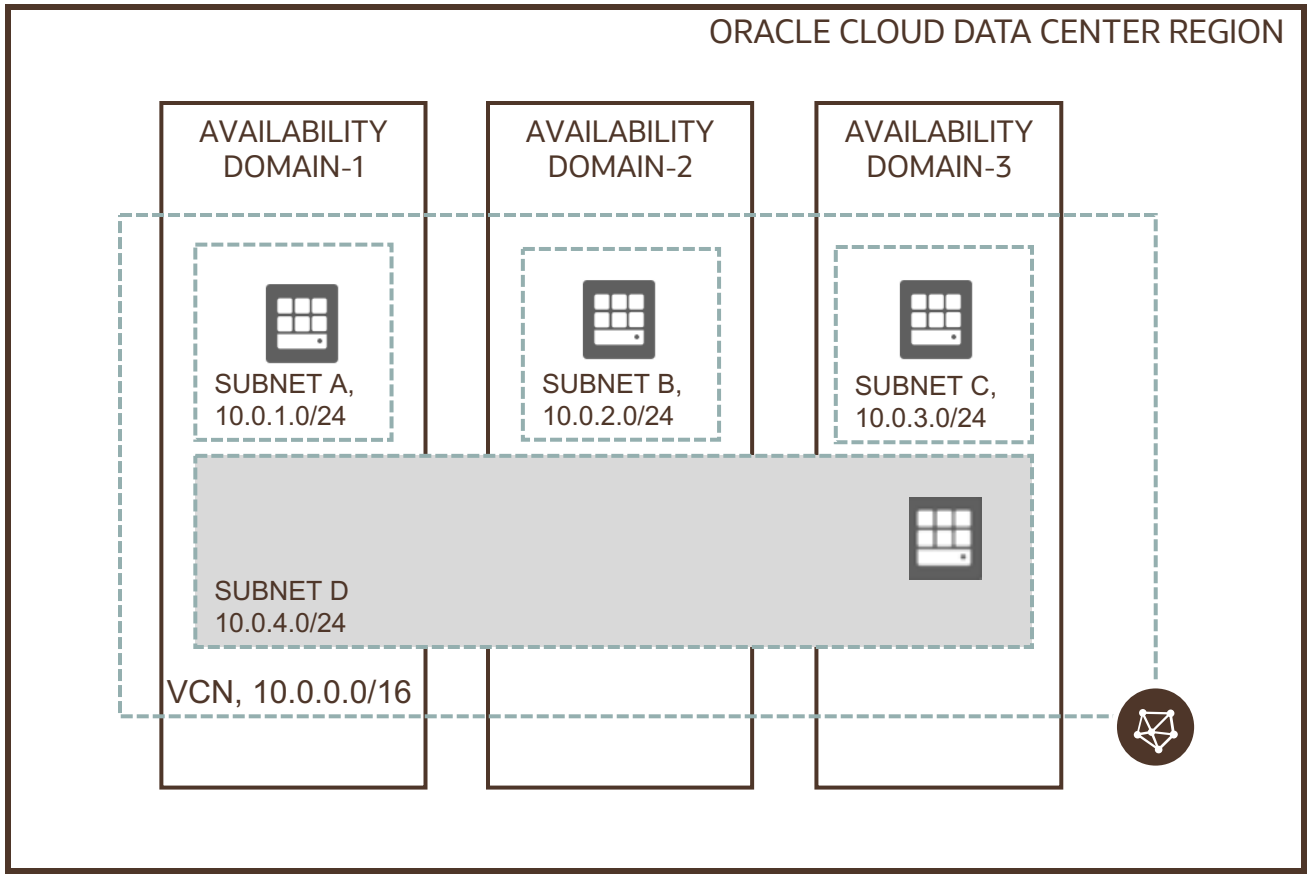

Each VCN network is subdivided into subnets

- Each subnet can be AD-specific or **Regional (recommended)**
  - AD specific subnet is contained within a single AD in a multi-AD region
  - Regional subnet spans all three ADs in a multi-AD region
- Each subnet has a contiguous range of IPs, described in CIDR notation. Subnet IP ranges cannot overlap

Subnet

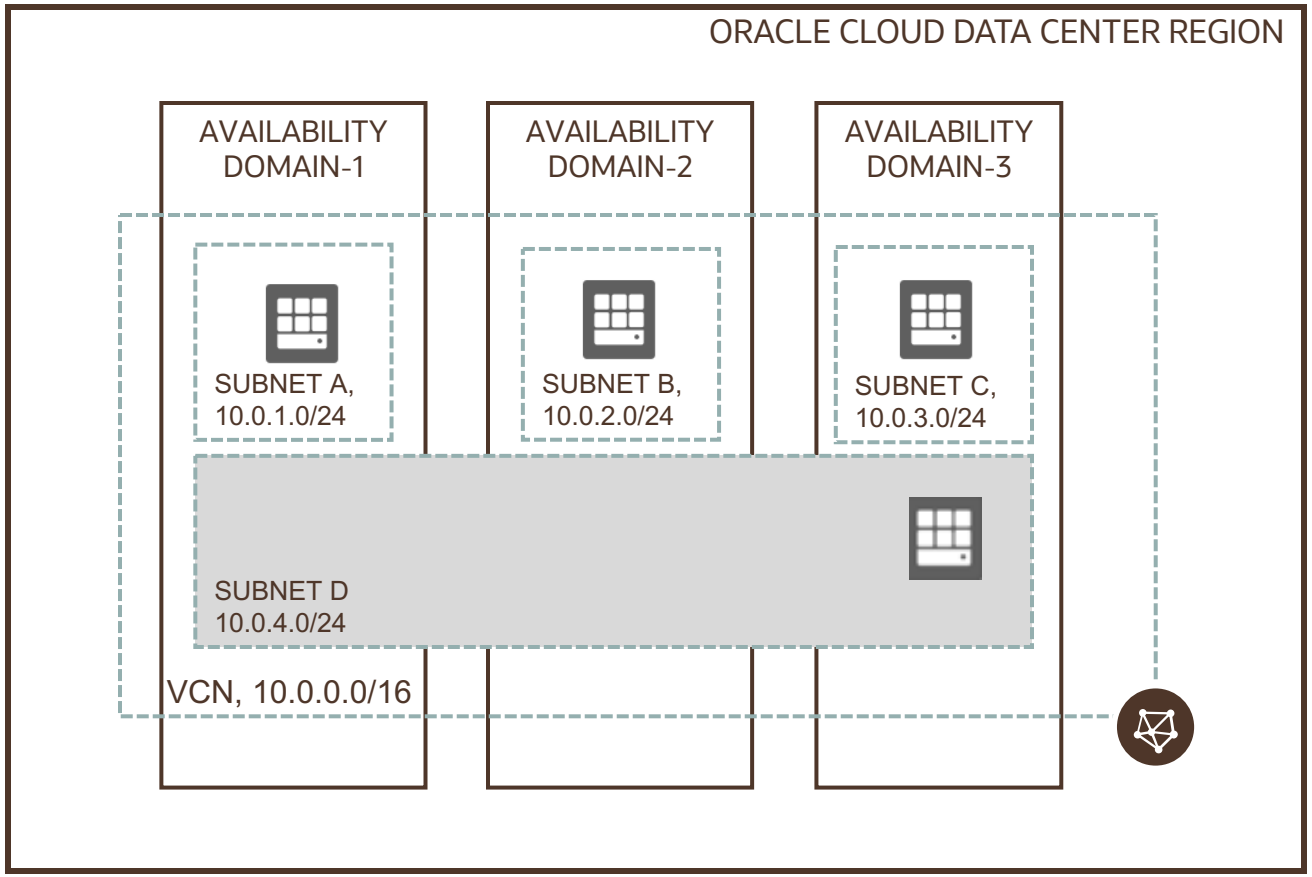

- Instances are placed in subnets and draw their internal IP address and network configuration from their subnet
- Subnets can be designated as either
  - **Private** (instances contain private IP addresses assigned to VNICs)
  - **Public** (contain both private and public IP addresses assigned to VNICs)
- VNIC is a component that enables a compute instance to connect to a VCN. The VNIC determines how the instance connects with endpoints inside and outside the VCN

ORACLE

IP Addresses

Private IP Addresses

- Each instance in a subnet has at least one primary private IP address
- Instances ≥ 2 VNICs (additional VNICs called secondary VNICs)
- Each VNIC has one primary private IP; can have additional private IPs called secondary private IPs
- A private IP can have an optional public IP assigned to it

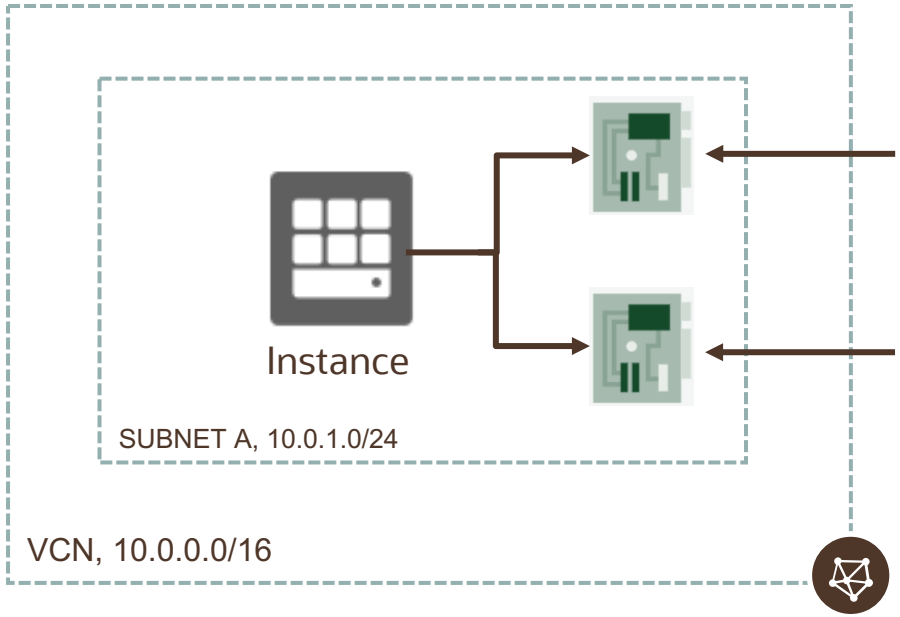

Primary VNIC
- Primary private IP address
- Secondary private IP, #1, #2…#31

Secondary VNIC
- Primary private IP address
- Secondary private IP, #1, #2…#31

Multiple VNICs on virtual machines

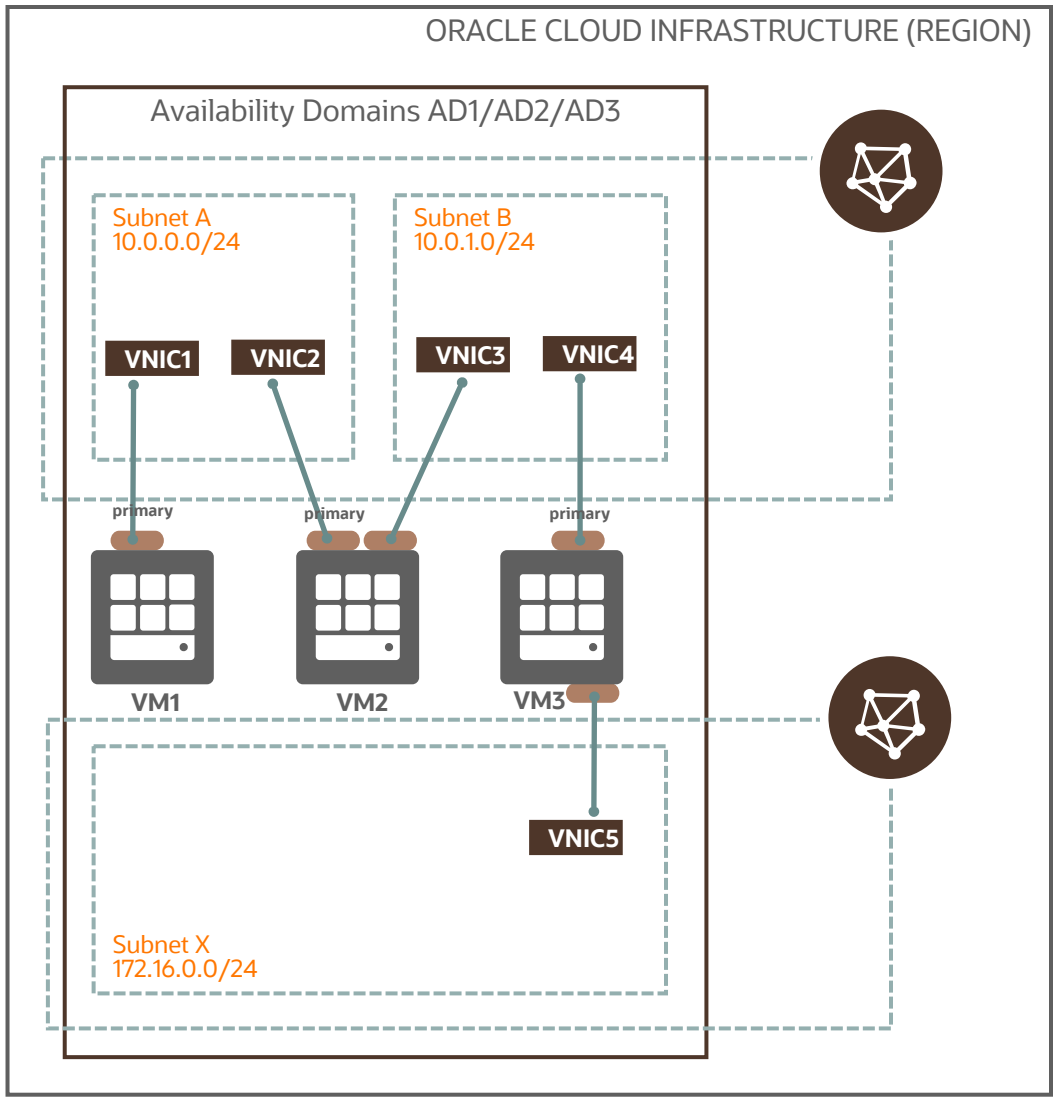

- Every VM has one primary VNIC created at launch, and a corresponding Ethernet device on the instance with the IP address configuration of the primary VNIC
- When a secondary VNIC is added, new Ethernet device is added and is recognized by the instance OS
  - VM1 - single VNIC instance
  - VM2 - connected to two VNICs from two subnets within the same VCN. Used for virtual appliance scenarios
  - VM3 - connected to two VNICs from two subnets from separate VCNs. Used to connect instances to a separate management network for isolated access

Public IP

- Public IP address is an IPv4 address that is reachable from the internet; assigned to a private IP object on the resource (Instance, load balancer)
- Possible to assign a given resource multiple public IPs across one or more VNICs

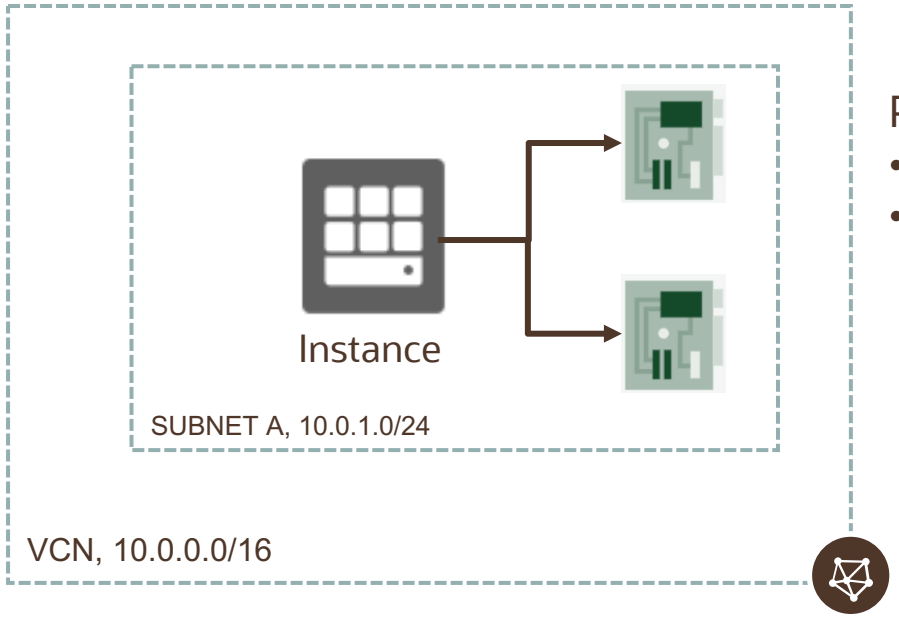

Primary VNIC
- Primary private IP address, public IP address
- Secondary private IP, #1, public IP address

Secondary VNIC
- Primary private IP address
- Secondary private IP, #1, #2…#31

Public IP Addresses

- Public IP types: Ephemeral and Reserved
  - Ephemeral: temporary and existing for the lifetime of the instance
  - Reserved: Persistent and existing beyond the lifetime of the instance it's assigned to (can be unassigned and then reassigned to another instance)
  - Ephemeral IP can be assigned to primary private IP only (hence, only 1 per VNIC v/s a max 32 for Reserved IP)
- No charge for using Public IP, including when the Reserved public IP addresses are unassociated
- Public IP assigned to
  - Instance (not recommended in most cases)
  - Oracle provided; cannot choose/edit, but can view
    - OCI Public Load Balancer, NAT Gateway, DRG - IPSec tunnels, OKE master/worker
  - Oracle provided; cannot choose/edit/view
    - Internet Gateway, Autonomous Database

ORACLE

Gateways and Routing

Internet Gateway

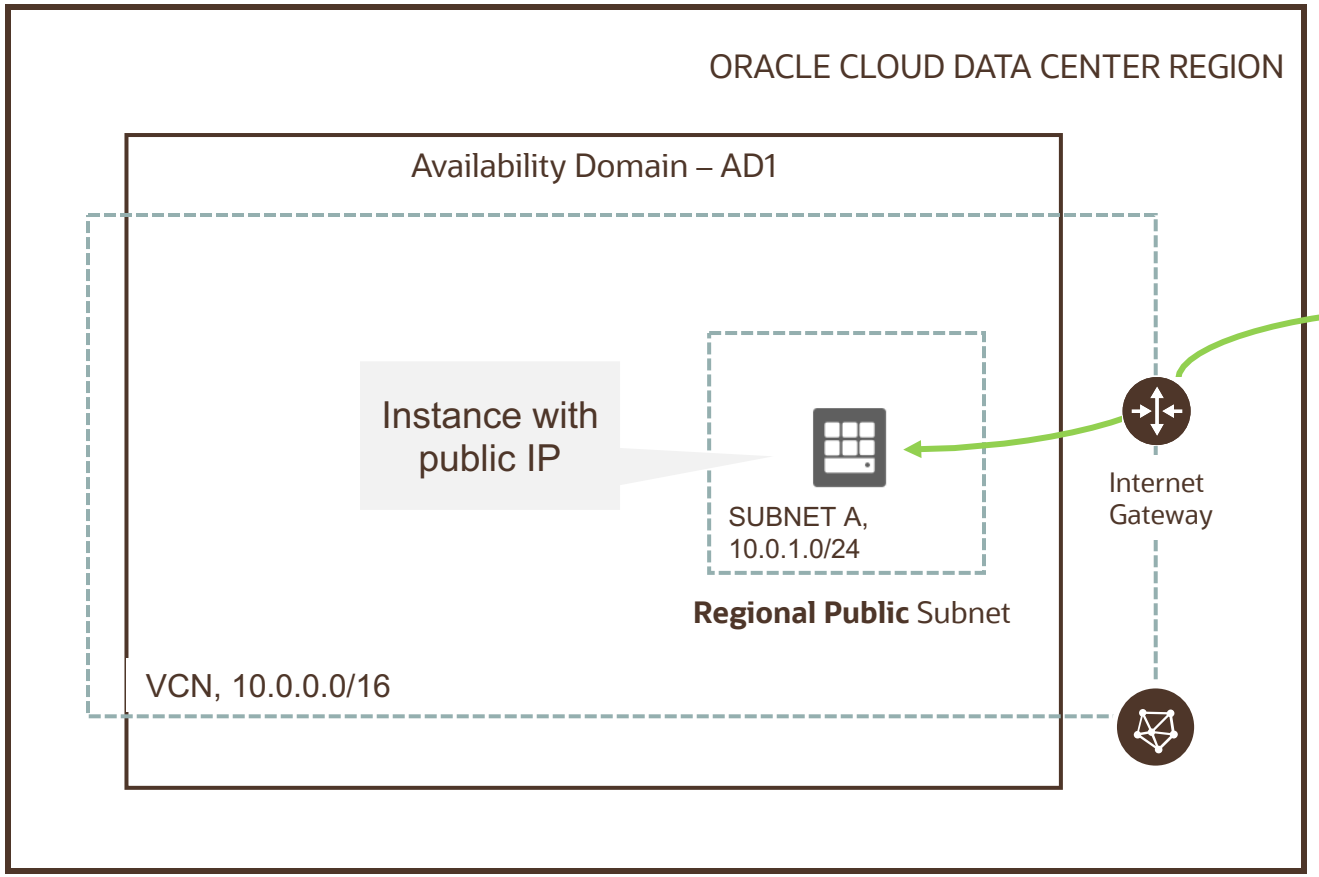

Internet gateway provides a path for network traffic between your VCN and the internet

You can have only one internet gateway for a VCN

After creating an internet gateway, you must add a route for the gateway in the VCN's Route Table to enable traffic flow

Route Table

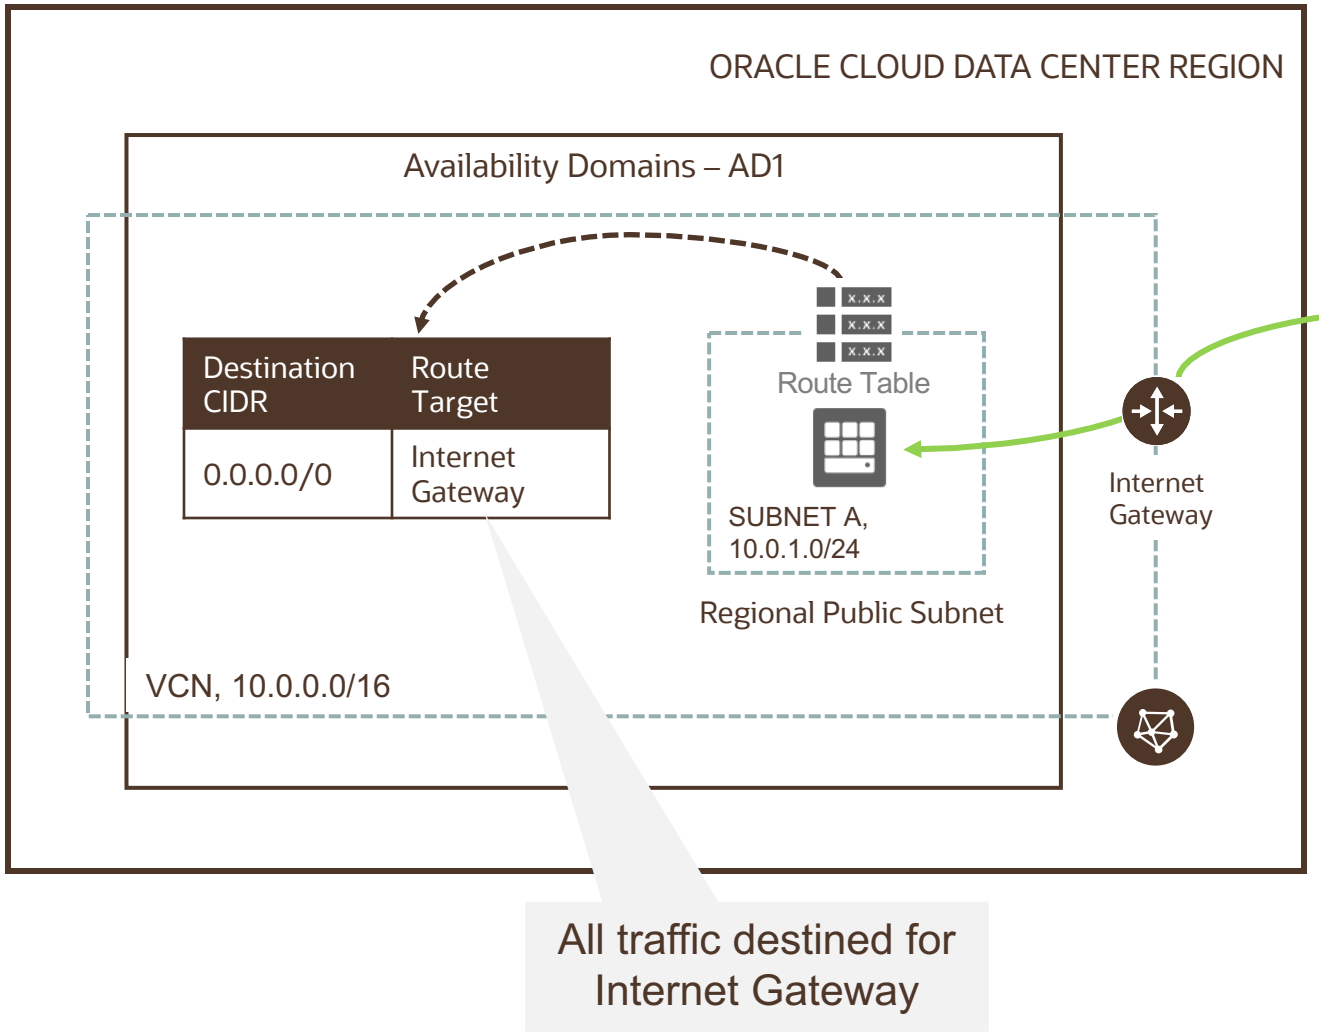

Route Table is used to send traffic out of the VCN

Consists of a set of route rules; each rule specifies

- Destination CIDR block
- Route Target (the next hop) for the traffic that matches that CIDR

Route Table

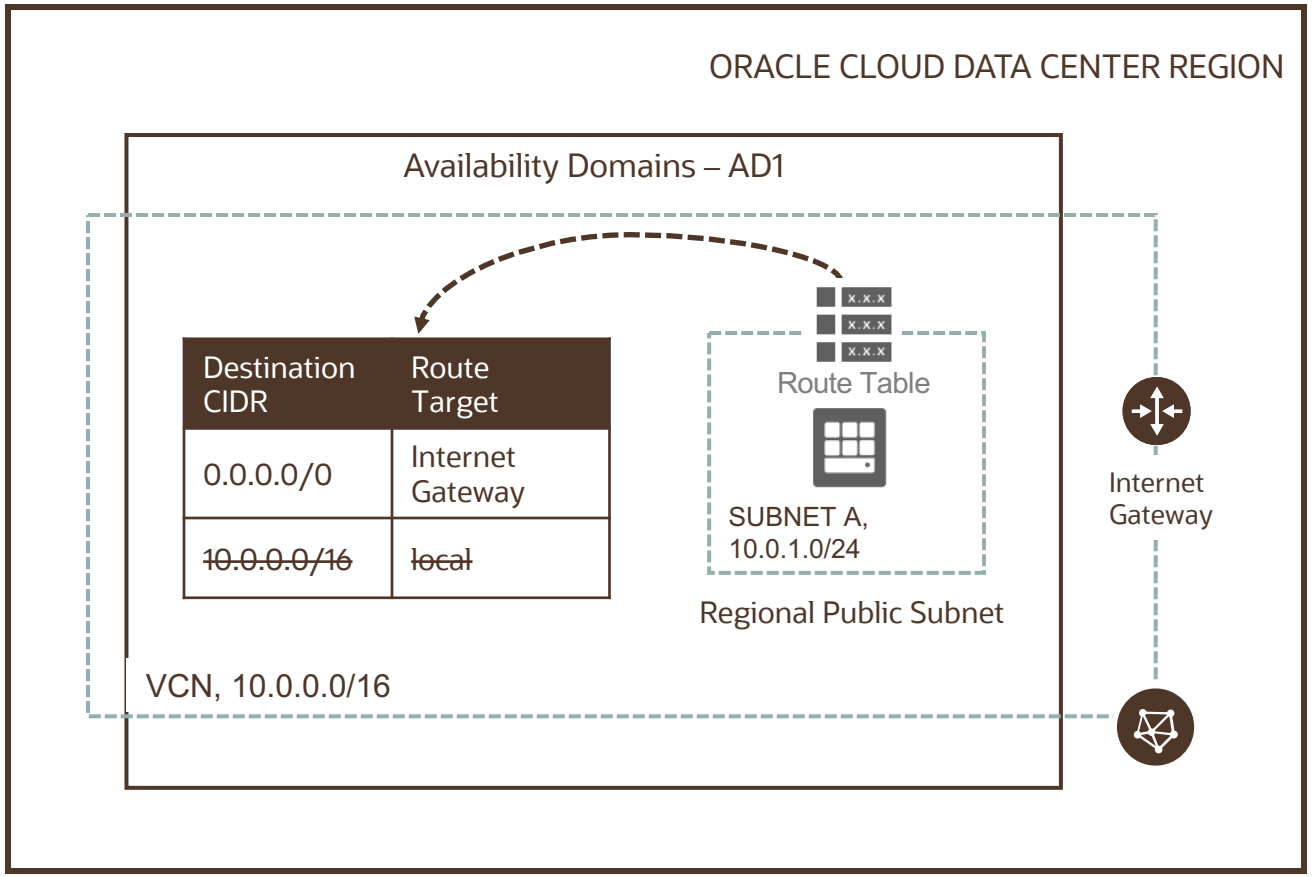

- Each subnet uses a single route table specified at time of subnet creation, but can be edited later
- Route table is used only if the destination IP address is not within the VCN's CIDR block
- No route rules are required in order to enable traffic within the VCN itself
- When you add an internet gateway, NAT gateway, service gateway, dynamic routing gateway or a peering connection, you must update the route table for any subnet that uses these gateways or connections

NAT Gateway

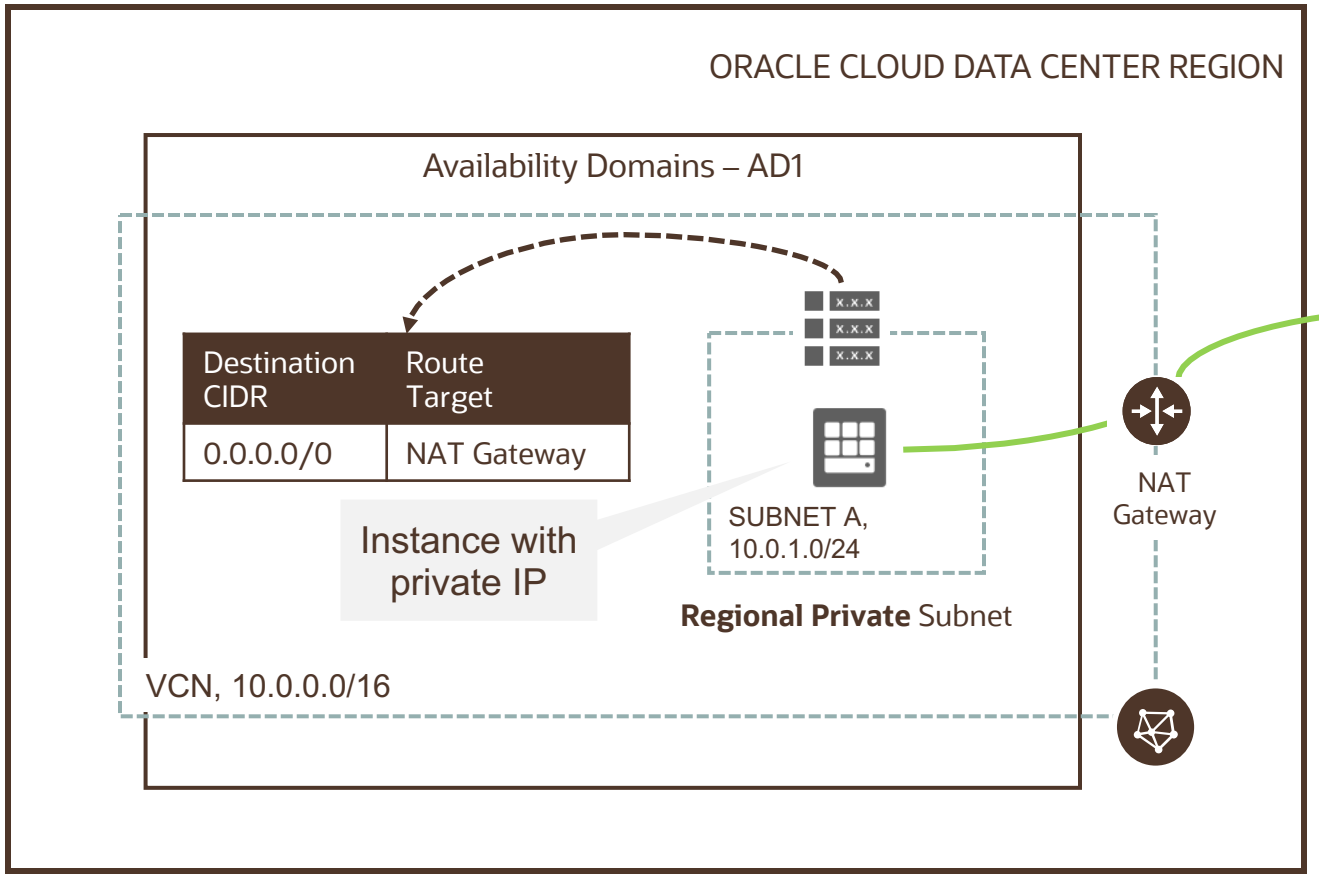

NAT gateway gives an entire private network access to the internet without assigning each host a public IP address

Hosts can initiate outbound connections to the internet and receive responses, but not receive inbound connections initiated from the internet. Use case: updates, patches)

You can have more than one NAT gateway on a VCN, though a given subnet can route traffic to only a single NAT gateway

Service Gateway

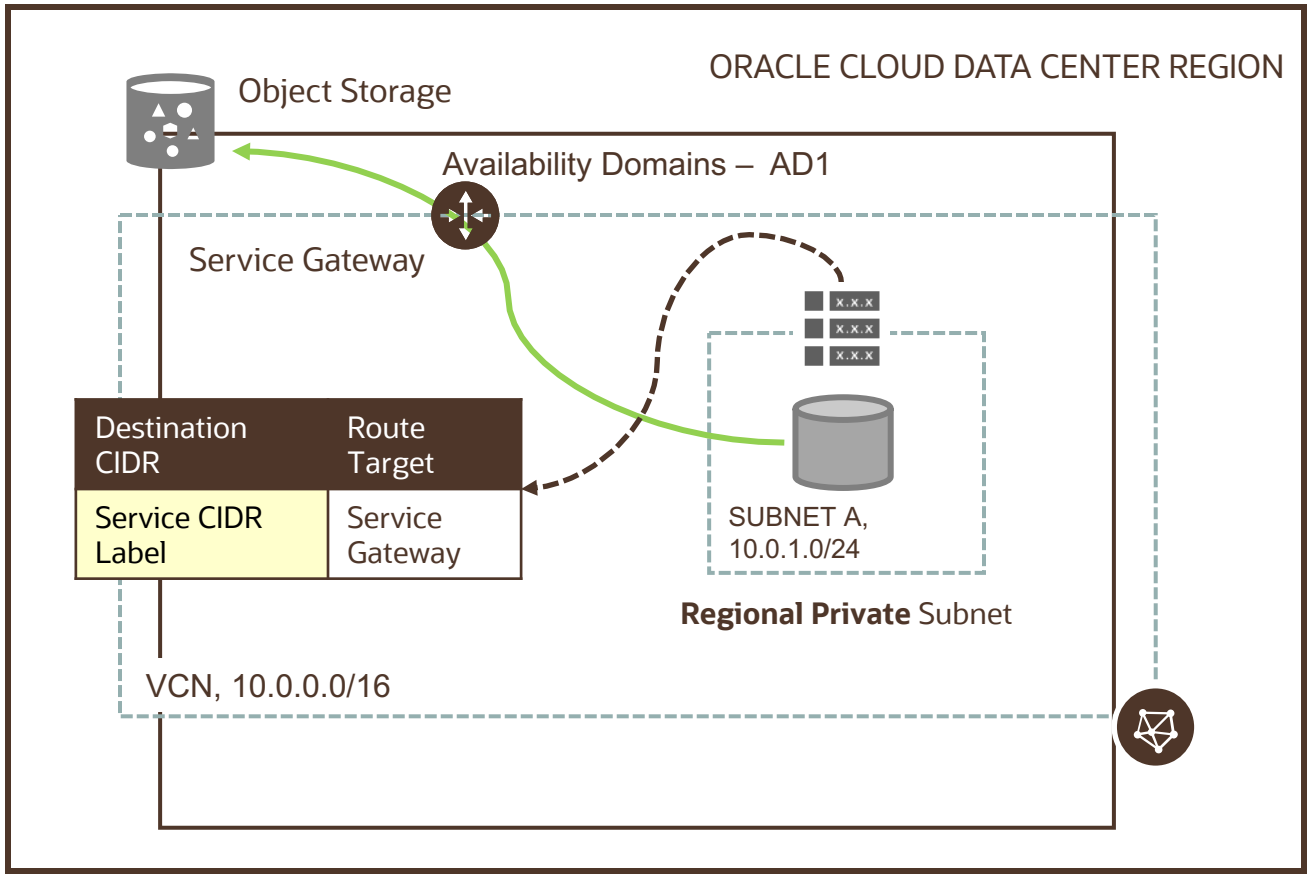

Service gateway lets resources in VCN access public OCI services such as Object Storage, but without using an internet or NAT gateway

Any traffic from VCN that is destined for one of the supported OCI public services uses the instance's private IP address for routing, travels over OCI network fabric, and never traverses the internet. Use case: back up DB Systems in VCN to Object Storage)

Service CIDR labels represent all the public CIDRs for a given Oracle service or a group of Oracle services. E.g.

- OCI <region> Object Storage
- All <region> Services

Dynamic Routing Gateway

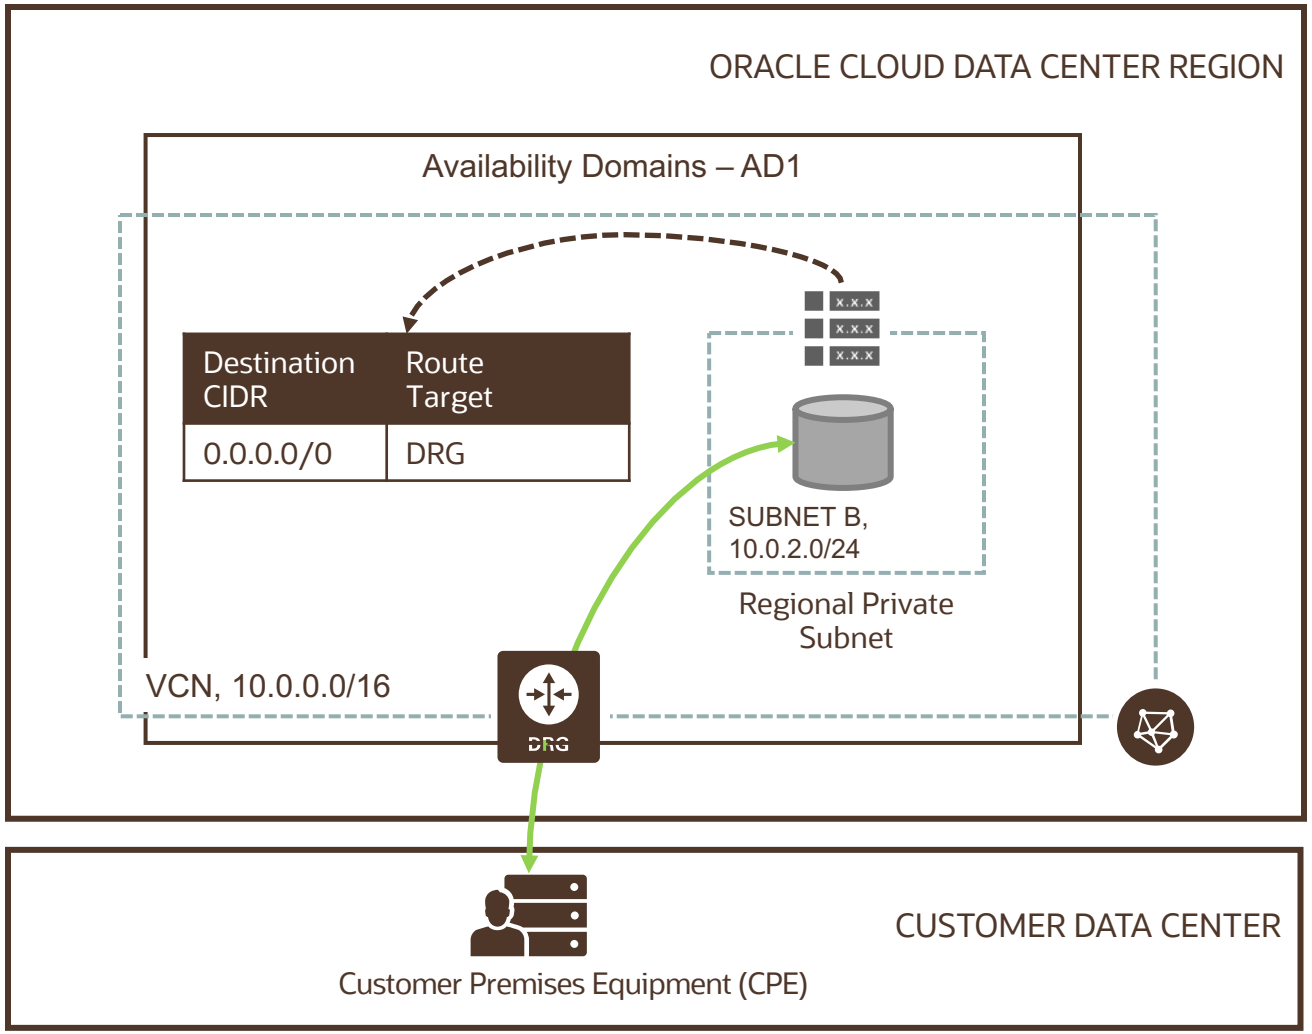

A virtual router that provides a path for private traffic between your VCN and destinations other than the internet

You can use it to establish a connection with your on-premises network via IPsec VPN or FastConnect (private, dedicated connectivity)

After attaching a DRG, you must add a route for the DRG in the VCN's route table to enable traffic flow

DRG is a standalone object. You must attach it to a VCN. VCN and DRG have a 1:1 relationship

ORACLE

Peering

Local Peering (within region)

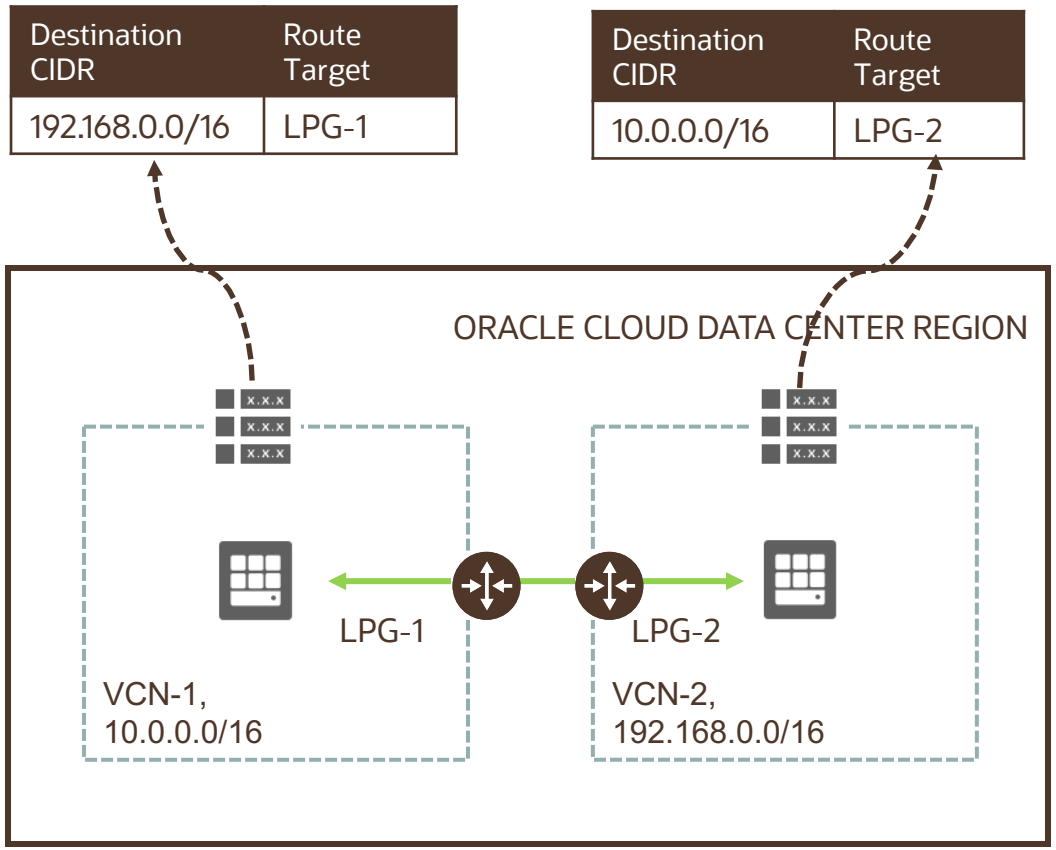

VCN peering is the process of connecting multiple VCNs

Local VCN peering is the process of connecting two VCNs in the **same region** so that their resources can communicate using private IP addresses

A local peering gateway (LPG) is a component on a VCN for routing traffic to a locally peered VCN

The two VCNs in the peering relationship shouldn't have overlapping CIDRs

Remote Peering (across region)

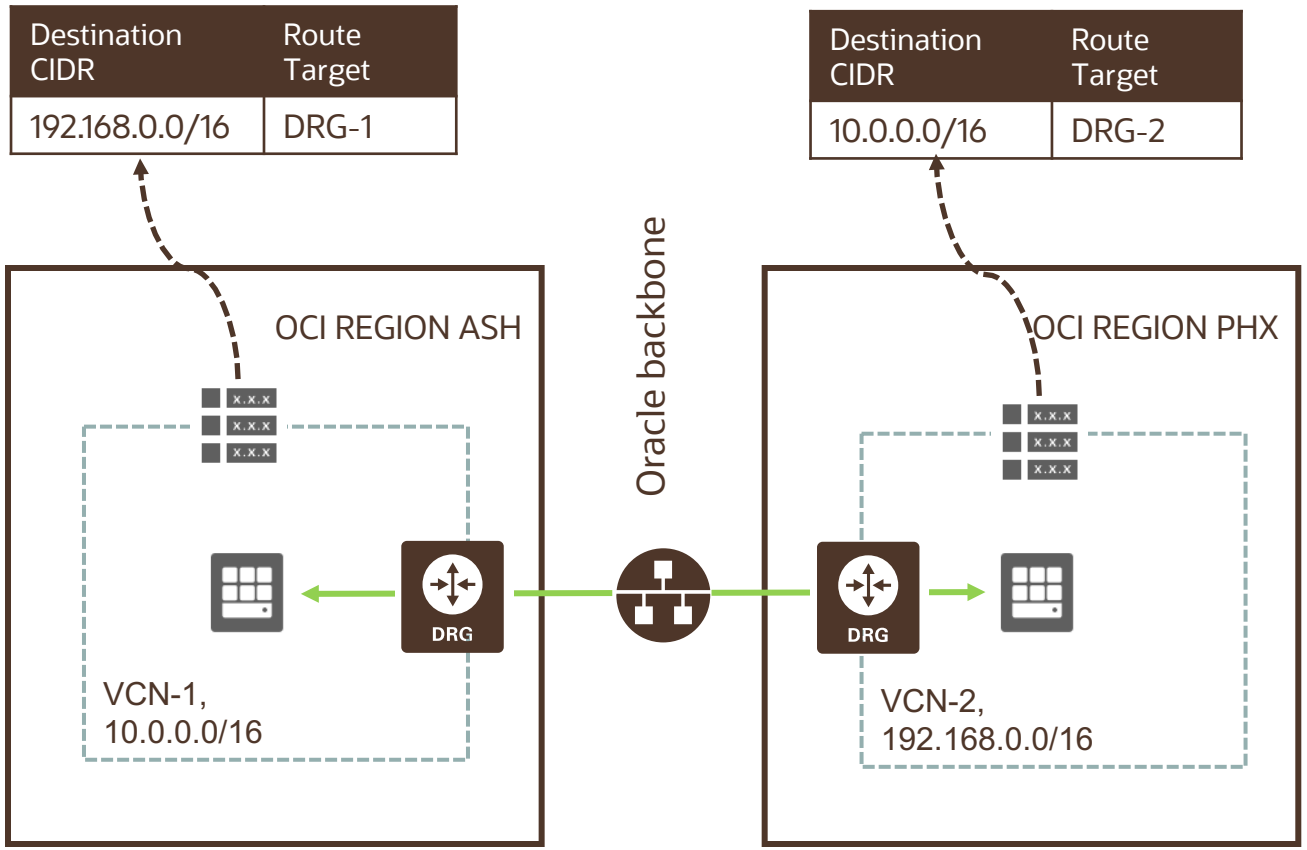

Remote VCN peering is the process of connecting two VCNs in **different regions** so that their resources can communicate using private IP addresses

Requires a remote peering connection (RPC) to be created on the DRGs. RPC's job is to act as a connection point for a remotely peered VCN

The two VCNs in the peering relationship must not have overlapping CIDRs

Summary of OCI network connectivity options

Scenario	Solution
Let instances connect to the Internet, and receive connections from it	Internet Gateway
Let instances reach the Internet without receiving connections from it	NAT Gateway
Let VCN hosts privately connect to object storage, bypassing the internet	Service Gateway
Make an OCI extend an on-premise network, with easy connectivity in both directions	IPsec VPN FastConnect
Privately connect two VCNs in a region	Local Peering Gateway
Privately connect two VCNs in different regions	Remote Peering Connection (DRG)

ORACLE

Transit Routing

Transit Routing: Hub and Spoke

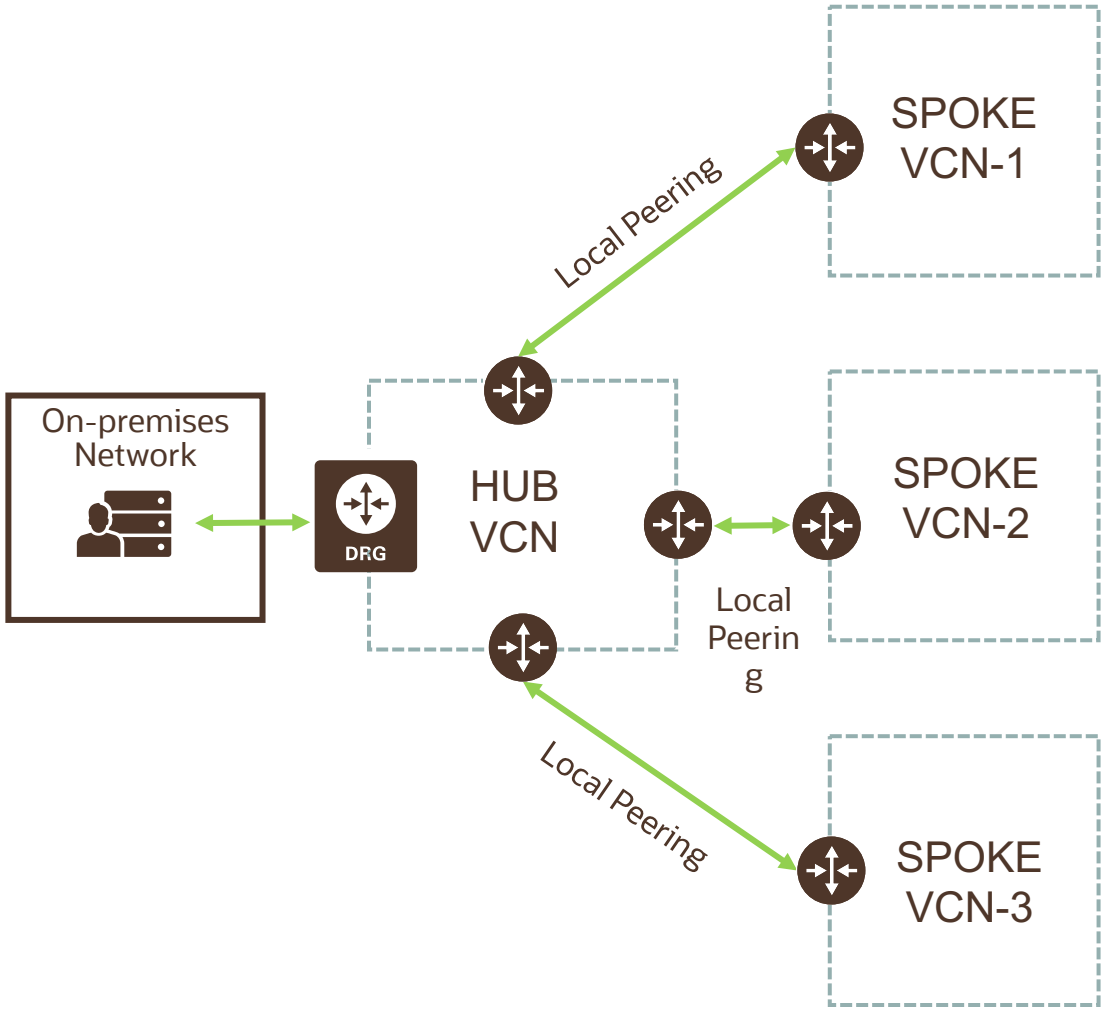

Transit Routing refers to set up in which an on-premises network uses a connected VCN to reach Oracle resources or services beyond that VCN. Two scenarios:

- Access to multiple VCNs in the same region
- Private access to Oracle services

One of the VCNs acts as the Hub and connects to on-premises network. The other VCNs are locally peered with the Hub VCN. The traffic between the on-premises network and the peered VCNs transits through the Hub VCN.

The VCNs must be in the same region but can be in different tenancies.

Transit Routing: Hub and Spoke

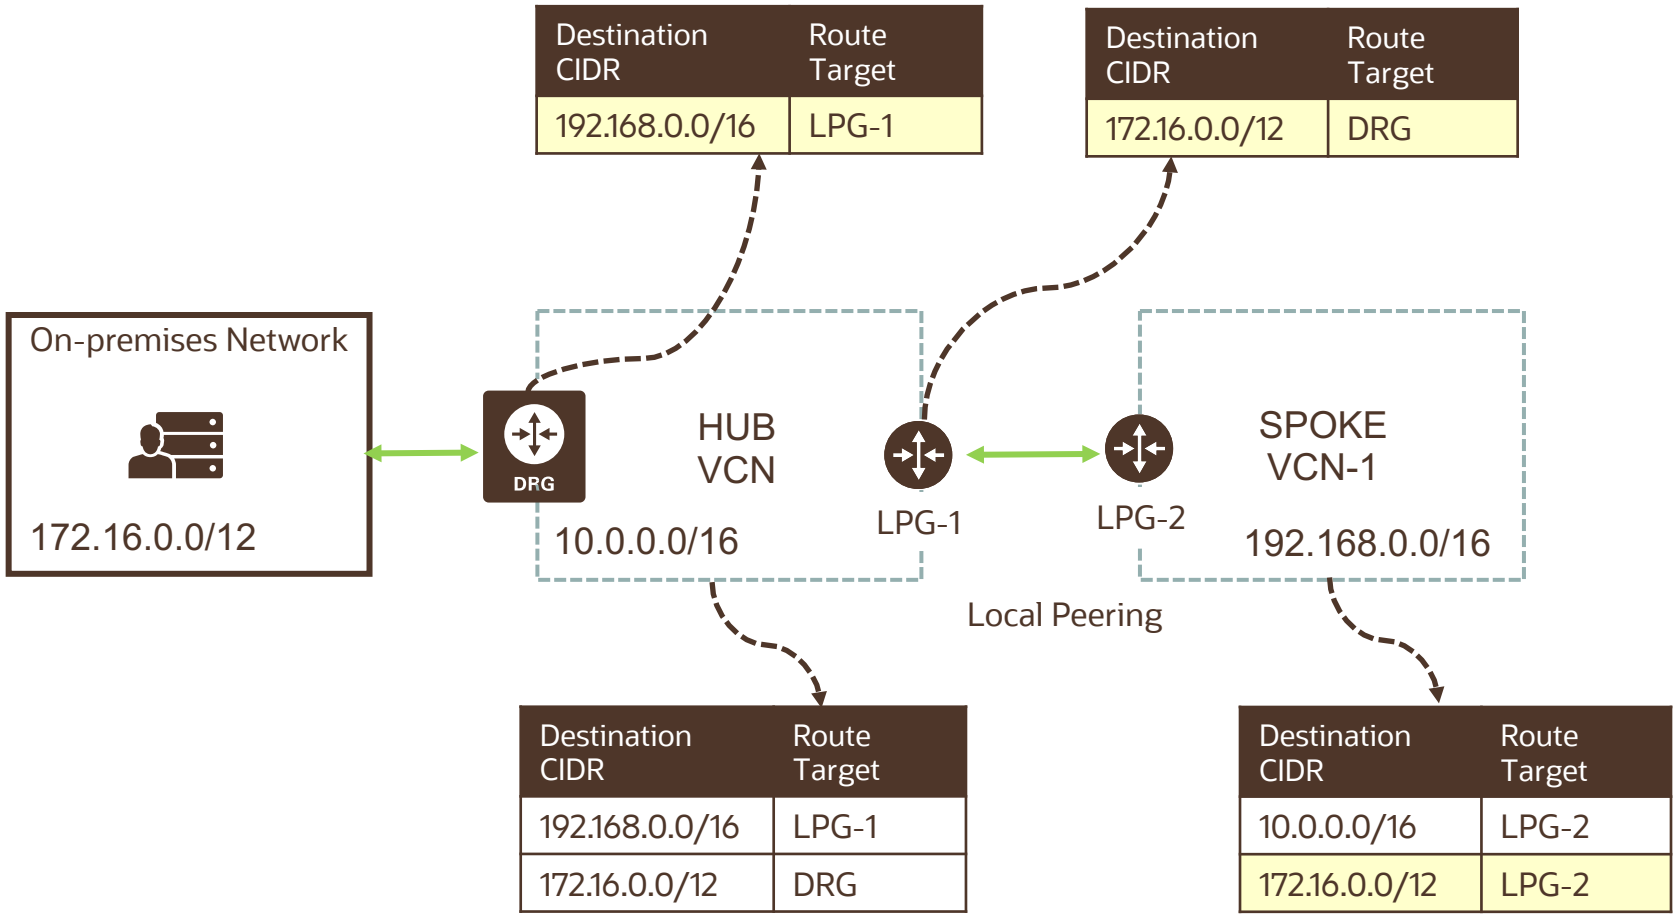

- A route table that is associated with a DRG can have only rules that target an LPG or a private IP
- A route table that is associated with an LPG can have only rules that target a DRG or a private IP
- DRG or LPG can exist without route table associated with it

Transit Routing: private access to Oracle services

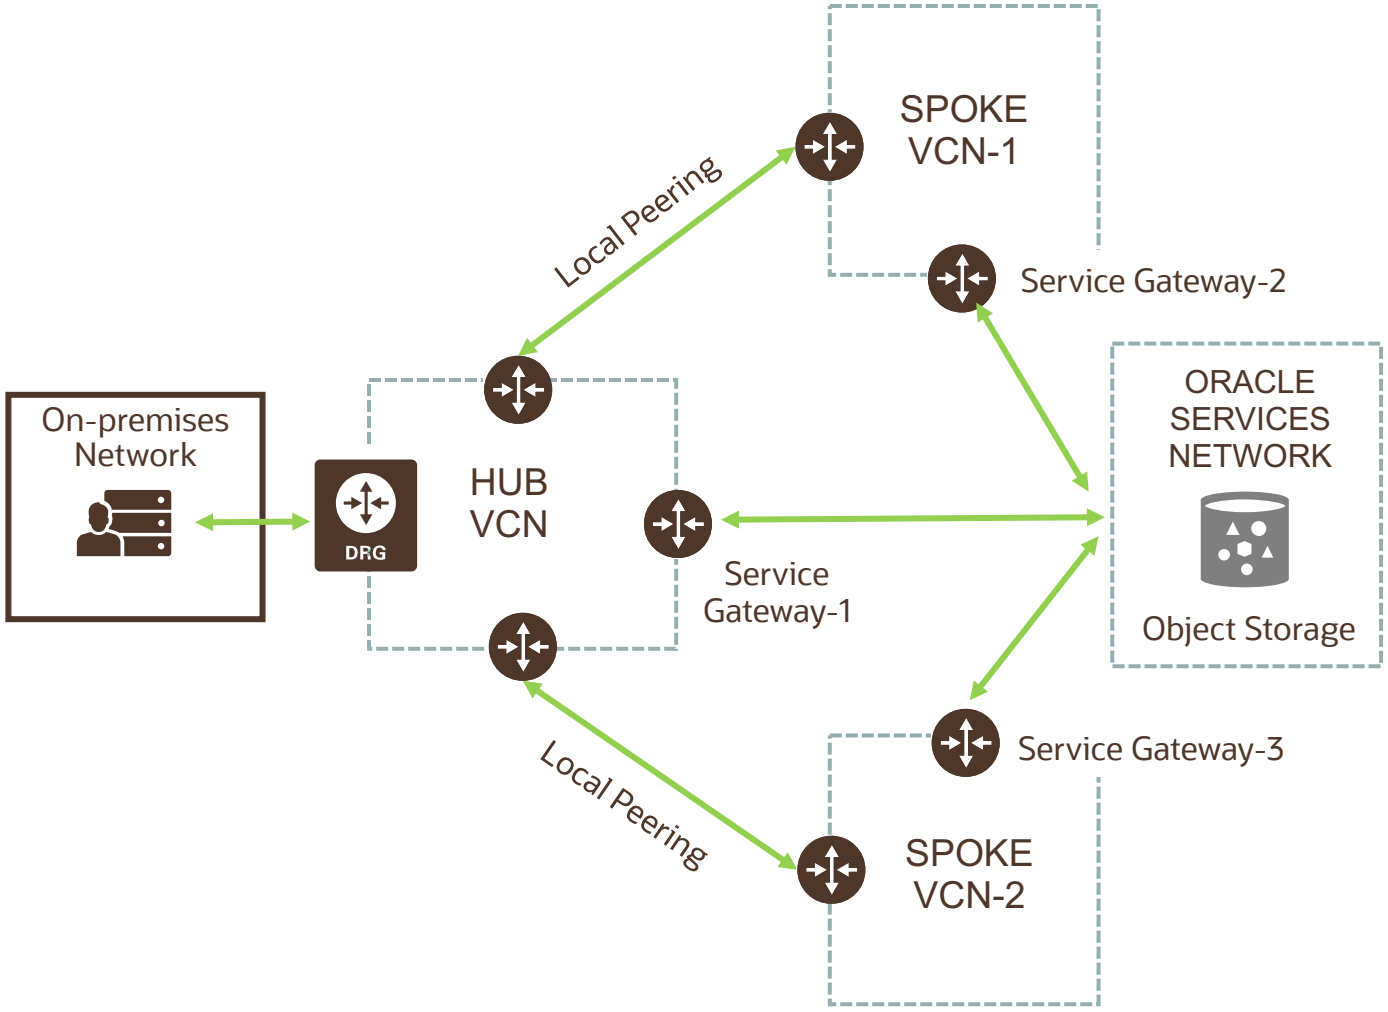

On-premises network has private access to Oracle services in the Oracle Services Network. The hosts in the on-premises network communicate with their private IP addresses

The on-premises network can reach the Oracle services only through a single VCN's service gateway (the one dedicated for this purpose, SG-1) and not through the service gateways of the other VCNs (SG-2,3).

For those other VCNs, only the resources inside those VCNs can reach Oracle services through their VCN's service gateway.

ORACLE

Security

Security List (SL)

	Direction	CIDR	Protocol	Source Port	Dest Port
Stateful	Ingress	0.0.0.0/0	TCP	All	80
Stateful	Egress	10.0.2.0/24	TCP	All	1521

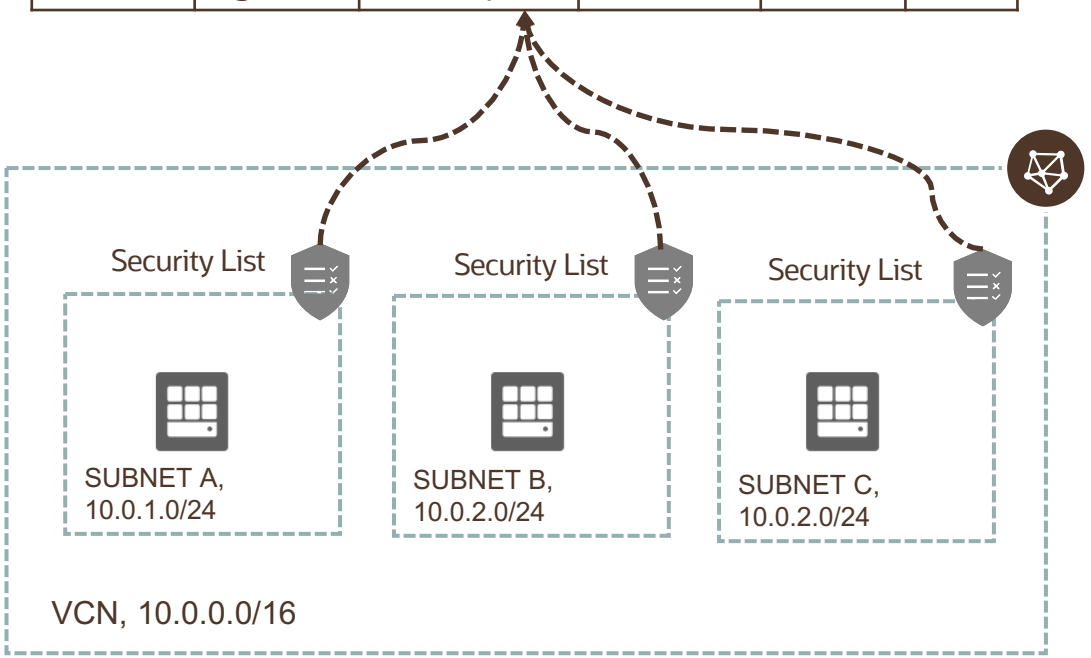

A common set of firewall rules associated with a subnet and applied to all instances launched inside the subnet

- Security list consists of rules that specify the types of traffic allowed in and out of the subnet
- To use a given security list with a particular subnet, you associate the security list with the subnet either during subnet creation or later.
- Security list apply to a given instance whether it's talking with another instance in the VCN or a host outside the VCN
- You can choose whether a given rule is stateful or stateless

Network Security Group (NSG)

		Direction	CIDR	Protocol	Source Port	Dest Port
NSG-A	Stateful	Ingress	0.0.0.0/0	TCP	All	80
NSG-B	Stateful	Ingress	0.0.0.0/0	TCP	All	22

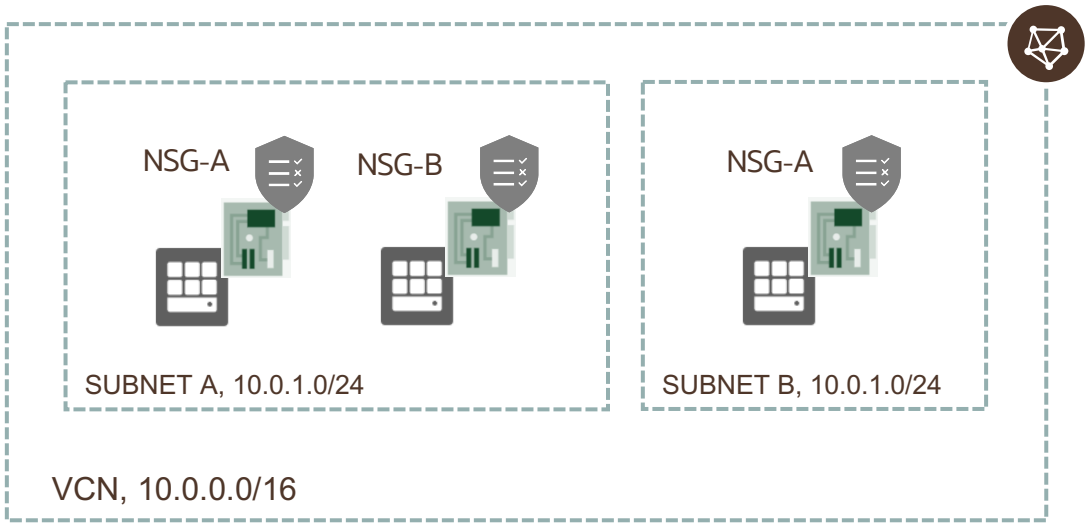

A network security group (NSG) provides a virtual firewall for a set of cloud resources that all have the same security posture

- NSG consists of set of rules that apply only to a set of VNICs of your choice in a single VCN
- Currently, compute instances, load balancers and DB instances support NSG
- When writing rules for an NSG, you can specify an NSG as the source or destination. Contrast this with SL rules, where you specify a CIDR as the source or destination
- Oracle recommends using NSGs instead of SLs because NSGs let you separate the VCN's subnet architecture from your application security requirements

SL + NSG

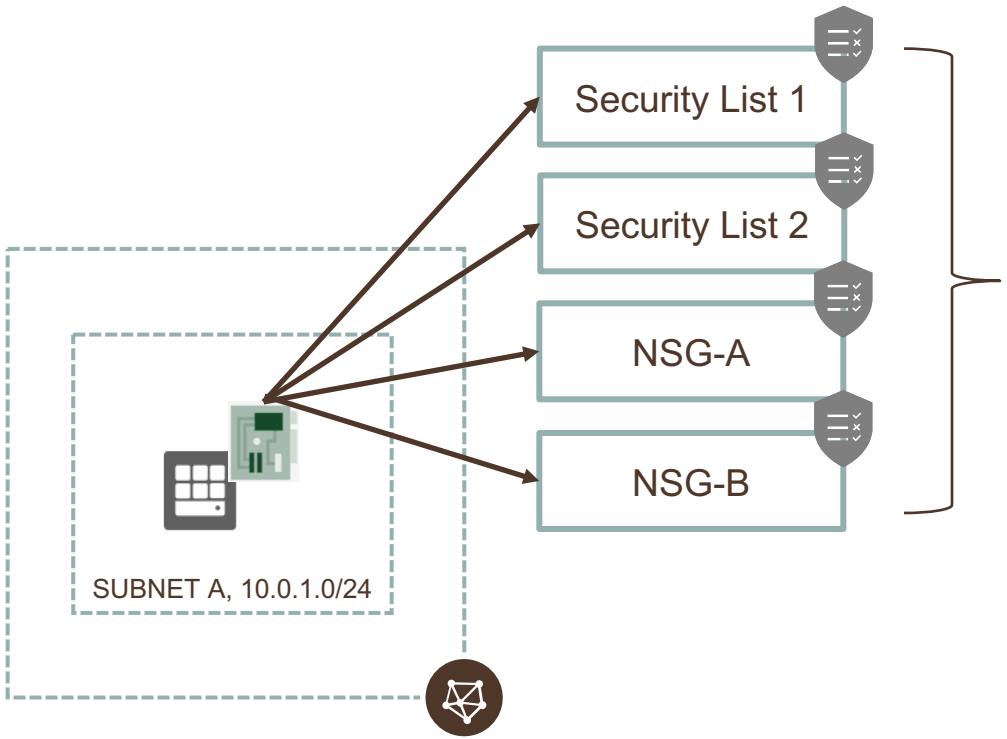

- You can use security lists alone, network security groups alone, or both together
- If you have security rules that you want to enforce for all VNICs in a VCN: the easiest solution is to put the rules in one security list, and then associate that security list with all subnets in the VCN
- If you choose to use both SLs and NSGs, the set of rules that applies to a given VNIC is the union of these items:
  - The security rules in the SLs associated with the VNIC's subnet
  - The security rules in all NSGs that the VNIC is in
  - A packet in question is allowed if any rule in any of the relevant lists and groups allows the traffic

Stateful Security Rules

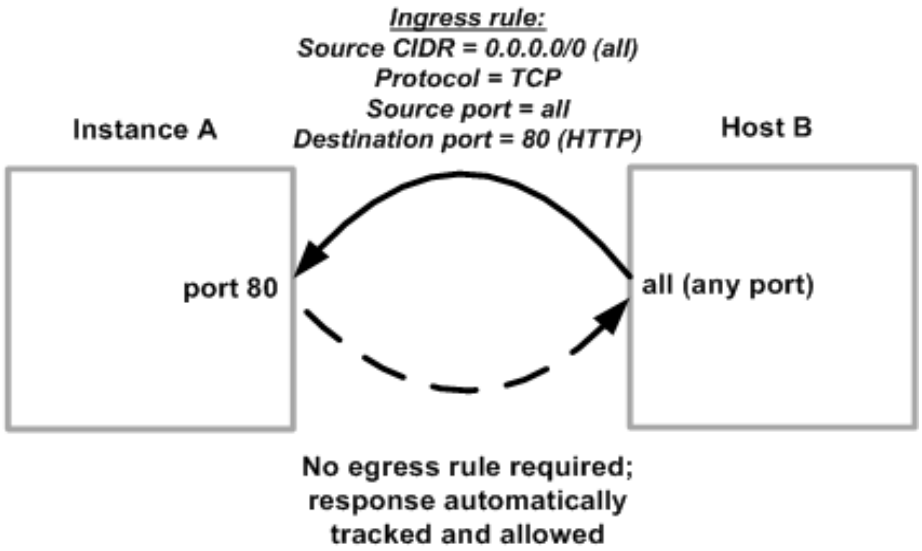

- Connection Tracking: when an instance receives traffic matching the stateful ingress rule, the response is tracked and automatically allowed regardless of any egress rules; similarly for sending traffic from the host
- Default Security List rules are stateful

SOURCE TYPE	SOURCE CIDR	IP PROTOCOL	SOURCE PORT RANGE *(OPTIONAL)*	DESTINATION PORT RANGE *(OPTIONAL)*
CIDR	0.0.0.0/0	TCP	All	80
	Specified IP addresses: 0.0.0.0-255.255.255.255 (4,294,967,296 IP addresses)	*(more information)*	Examples: 80, 20-22 or All *(more information)*	Examples: 80, 20-22 or All *(more information)*

Hosts in this group are reachable from the internet on Port 80

Stateless Security Rules

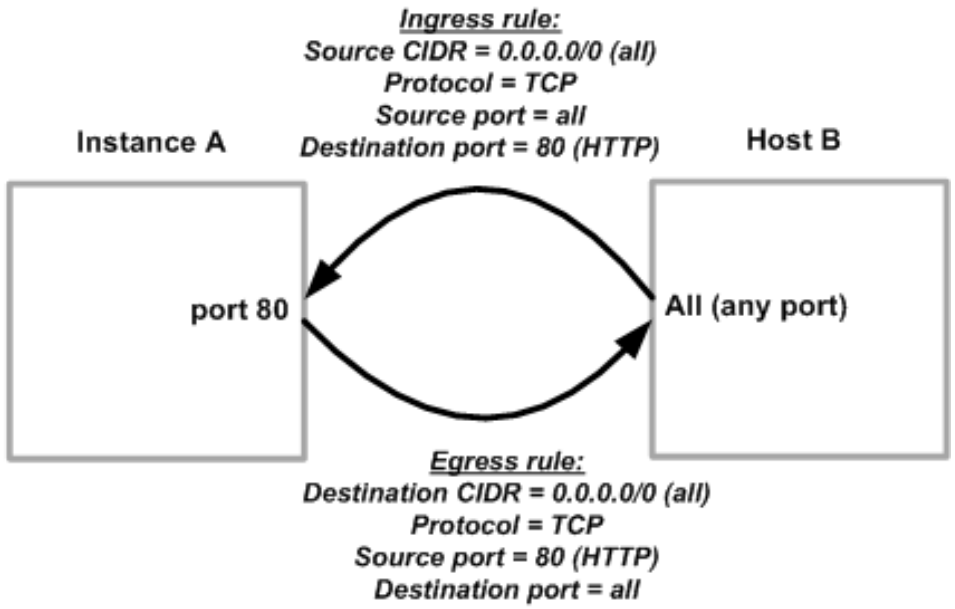

- With stateless rules, response traffic is not automatically allowed
- To allow the response traffic for a stateless ingress rule, you must create a corresponding stateless egress rule
- If you add a stateless rule to a security list, that indicates that you do NOT want to use connection tracking for any traffic that matches that rule
- Stateless rules are better for scenarios with large numbers of connections (Load Balancing, Big Data)

ORACLE

Default VCN, Internal DNS

Default VCN components

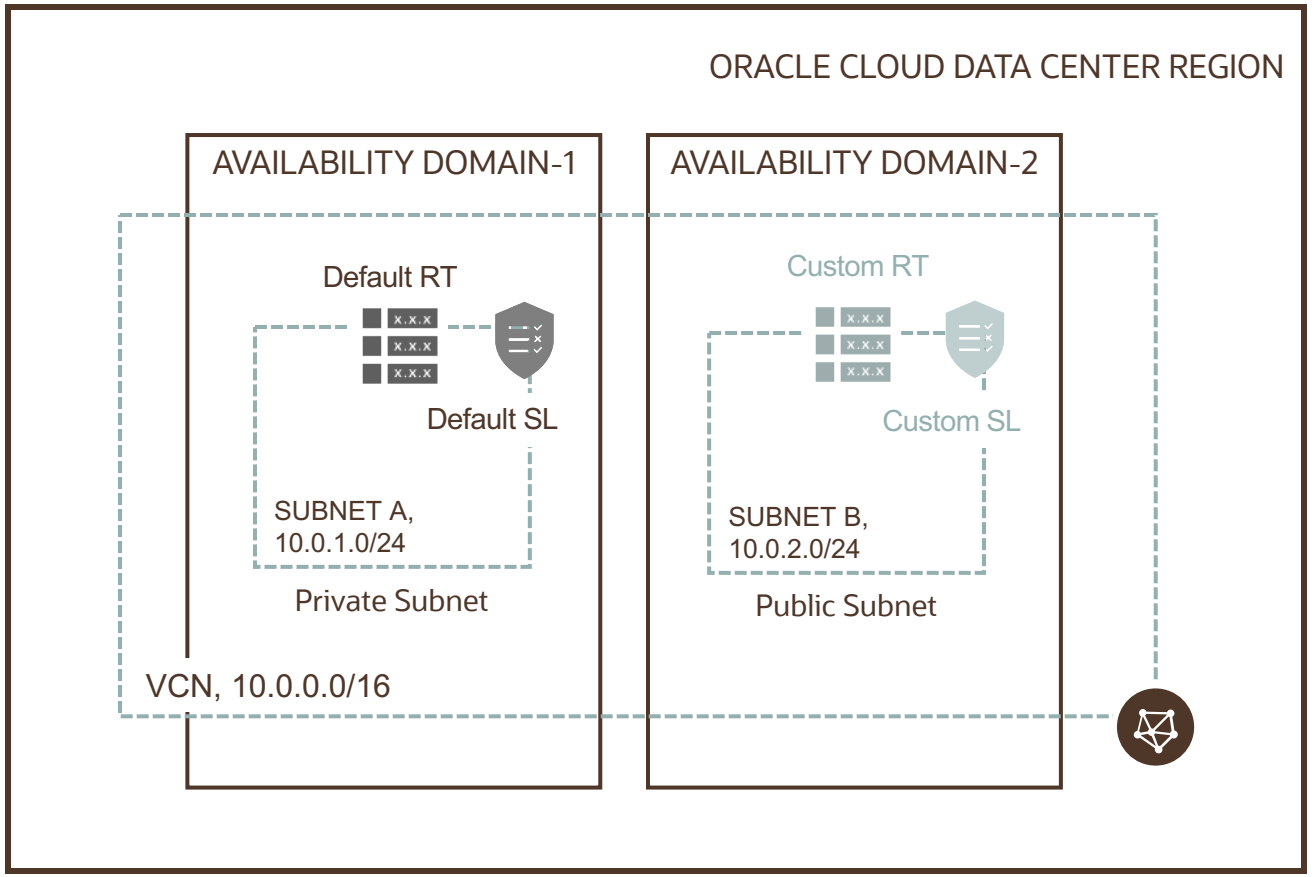

Your VCN automatically comes with some default components

- Default Route Table
- Default Security List
- Default set of DHCP options

You can't delete these default components; however, you can change their contents (e.g. individual route rules). And you can create more of each kind of component in your cloud network (e.g. additional route tables).

Internal DNS

- The VCN Private Domain Name System (DNS) enables instances to use hostnames instead of IP addresses to talk to each other
- Options:
  - Internet and VCN Resolver: default choice for new VCNs
  - Custom Resolver: lets instances resolve the hostnames of hosts in your on-premises network through IPsec VPN/FastConnect
- Optionally specify a DNS label when creating VCN/subnets/instances
  - VCN: <VCN DNS label>.oraclevcn.com
  - Subnet: <subnet DNS label>.<VCN DNS label>.oraclevcn.com
  - Instance FQDN: <hostname>.<subnet DNS label>.<VCN DNS label>.oraclevcn.com

- Instance FQDN resolves to the instance's Private IP address
- No automatic creation of FQDN for Public IP addresses (e.g. cannot SSH using <hostname>.<subnet DNS label>.<VCN DNS label>.oraclevcn.com)

ORACLE

Putting it all together

VCN Review

- Subnets can have one Route Table and multiple (5*) Security Lists associated to it
- Route table defines what can be routed out of VCN
- Private subnets are recommended to have individual route tables to control the flow of traffic outside of VCN
- All hosts within a VCN can route to all other hosts in a VCN (no local route rule required)
- Security Lists manage connectivity north-south (incoming/outgoing VCN traffic) and east-west (internal VCN traffic between multiple subnets)
- OCI follows a white-list model (you must manually specify white listed traffic flows); By default, things are locked down
- Instances cannot communicate with other instances in the same subnet, until you permit them to!
- Oracle recommends using NSGs instead of SLs because NSGs let you separate the VCN's subnet architecture from your application security requirements

VCN Review

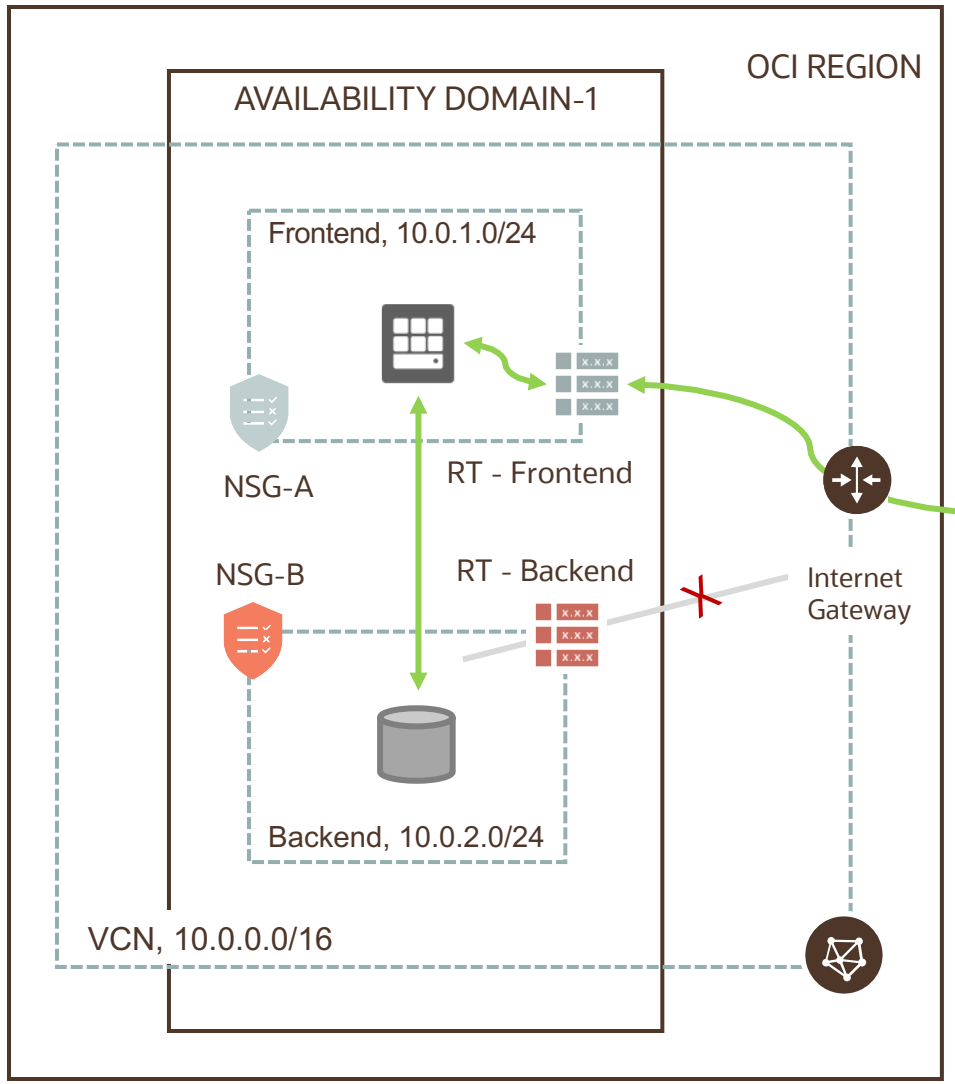

Destination CIDR	Route Target
0.0.0.0/0	Internet Gateway

	Type	CIDR	Protocol	Source Port	Dest Port
Stateful	Ingress	0.0.0.0/0	TCP	All	80
Stateful	Egress	NSG-B	TCP	All	1521

Destination CIDR	Route Target
0.0.0.0/0	NAT/ Service gateway /DRG

	Type	CIDR	Protocol	Source Port	Dest Port
Stateful	Ingress	NSG-A	TCP	All	1521
~~Stateful~~	~~Egress~~		~~All~~	~~All~~	

Summary

- Key Virtual Cloud Network (VCN) concepts

  Subnets, Route Table, Private IP, Public IP, Internal DNS
- Gateways and Routing

  Internet Gateway, NAT Gateway, Service Gateway, Local and Remote Peering

  Transit Routing

  VPN, FastConnect (next module)

VCN Security

Security List, Network Security Groups

ORACLE

Oracle Cloud always free tier:
oracle.com/cloud/free/

OCI training and certification:
https://www.oracle.com/cloud/iaas/training/
https://www.oracle.com/cloud/iaas/training/certification.html
education.oracle.com/oracle-certification-path/pFamily_647

OCI hands-on labs:
ocitraining.qloudable.com/provider/oracle

Oracle learning library videos on YouTube:
youtube.com/user/OracleLearning

Thank you