



ORACLE

Learn Routing in Oracle Cloud Infrastructure Networking with Examples

September 2023, version 1.0
Copyright © 2023, Oracle and/or its affiliates
Public

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced, or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement, nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. 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, and timing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Revision History

The following revisions have been made to this document.

DATE	REVISION
September 2023	Initial publication

Table of Contents

Introduction	4
OCI VCN Routing	4
VCN Route Table Lookup Operation	5
Intra-VCN Routing	6
VCN Egress Routing	8
VCN Gateway Ingress Routing	9
DRG Routing Basics	17
Route Tables in a DRG	18
Route Preference in DRG Route Table	18
Route Propagation and Import Route Distribution Control on a DRG	20
DRG Routing Operation	21
Intra-Regional VCN Routing	24
Routing Between VCNs by Using LPGs	24
Routing Between VCNs by Using DRGs	25
Inter-Regional Routing	26
On-Premises Site to VCN Routing Examples	27
Remote On-Ramp Routing Example	30
Routing for Network Virtual Appliance Insertion Examples	32
Per-VCN Network Virtual Appliances	32
Centrally Shared Network Virtual Appliances	32
Private Access to OCI Services	34
OCI Instances Access OCI Services Through a Service Gateway	34
On-Premises Instances Access OCI Services Through a Service Gateway	36
Conclusion	38

Introduction

Oracle Cloud Infrastructure (OCI) offers a software-defined virtual network solution to our customers. An OCI network consists of virtual cloud networks (VCNs), subnets, network gateways, OCI native or third-party L4-7 network service virtual appliances, and more. *Routing* is the core function that establishes connectivity among the elements in an OCI network, or between an OCI network and on-premises networks or other cloud networks.

This technical paper explains routing in OCI cloud networks. It introduces and discusses basic OCI routing functions, and then goes in depth with typical use cases in different deployment scenarios.

Full network reachability in a network requires both network connectivity that is achieved by proper routing and network security policies that are managed through security lists or network security groups, or policies on network firewall appliances. This paper focuses solely on routing functions and designs; the management of network security policies is not in the scope of this paper.

The discussions and examples in this paper focus on IPv4 routing, but the same theories also apply to IPv6 routing. OCI uses the same routing mechanisms for IPv4 and IPv6. The network design for IPv6 requires unique considerations, however, because of the differences from IPv4, such as the different scopes of IPv6 addresses and the fact that IPv6 internet routing doesn't go through NAT. IPv6 addressing and networking are not in the scope of this paper. For more information about IPv6 support on OCI, see [IPv6 Addresses](#) in the OCI Networking documentation.

OCI VCN Routing

Routing in a cloud network serves the same purpose as routing in a traditional network: selecting the explicit path through the network to forward traffic between the source and the destination of a communication traffic flow. The source and the destination can be in different subnets, different virtual cloud networks (VCNs), different regions, or even different cloud networks. They can also be between the cloud network and on-premises networks or the internet.

OCI VCNs use *route tables* to manage the route rules for traffic forwarding. Each VCN has a system-created, default route table, but users can create more route tables for a VCN. Each VCN route table has two types of route rules: *system-defined local route rules* for the VCN CIDR blocks and *user-defined route rules*. The system-defined local route rules are created automatically by OCI for each of the VCN CIDR blocks. These local route rules aren't visible to users, but they do participate in the routing lookup process for the longest prefix match (LPM)-based best-path selection, which is explained in detail in a later part of this section. If the local route rule is resolved as the best route for a destination inside the VCN, the traffic is sent to the destination directly. Users can't modify or delete these system-defined local route rules in a VCN route table, but they can define their own route rules in the same route table for more route control based on their needs.

You can associate a VCN route table with VCN subnets or gateways for the following purposes:

- **Intra-VCN routing:** Route traffic between subnets in the same VCN.
- **VCN subnet egress routing:** Route traffic from sources in a VCN subnet to destinations outside of the VCN.
- **VCN gateway ingress routing:** Route traffic sourced from outside a VCN to destinations inside the VCN, or route transit traffic through a VCN to its destination. Many OCI gateways—such as service gateways (SGW), local peering gateways (LPG), and dynamic routing gateways (DRG)—support transit routing.

The first two types of routing, intra-VCN routing and VCN subnet egress routing, use the *subnet route table*, which is the VCN route table associated with a subnet. The third type, VCN gateway ingress routing and transit routing, requires a *VCN route table* to be associated with the gateway for route rules when it routes the traffic into the VCN or through the VCN.

The remainder of this section discusses the VCN routing lookup operation and the three basic VCN routing scenarios.

VCN Route Table Lookup Operation

Like the routing lookup operation in a traditional on-premises network, OCI uses the longest prefix match (LPM) algorithm for the best route selection. Each entry in a route table is a route rule that defines the next hop in the routing path to the destination prefix. One destination address might match more than one route rule. The most specific of the matching rules—the one with the longest subnet mask—is called the LPM. It is selected as the best route to reach the destination address.

A VCN route table contains user-defined route rules and the implicit local route rules for each of the VCN CIDRs. These local route rules aren't visible to users, but they're taken into the LPM route-selection algorithm in the same way as the user-defined route rules in the route table. To illustrate the operation, let's look at an example (Figure 1) with a simple VCN route table.

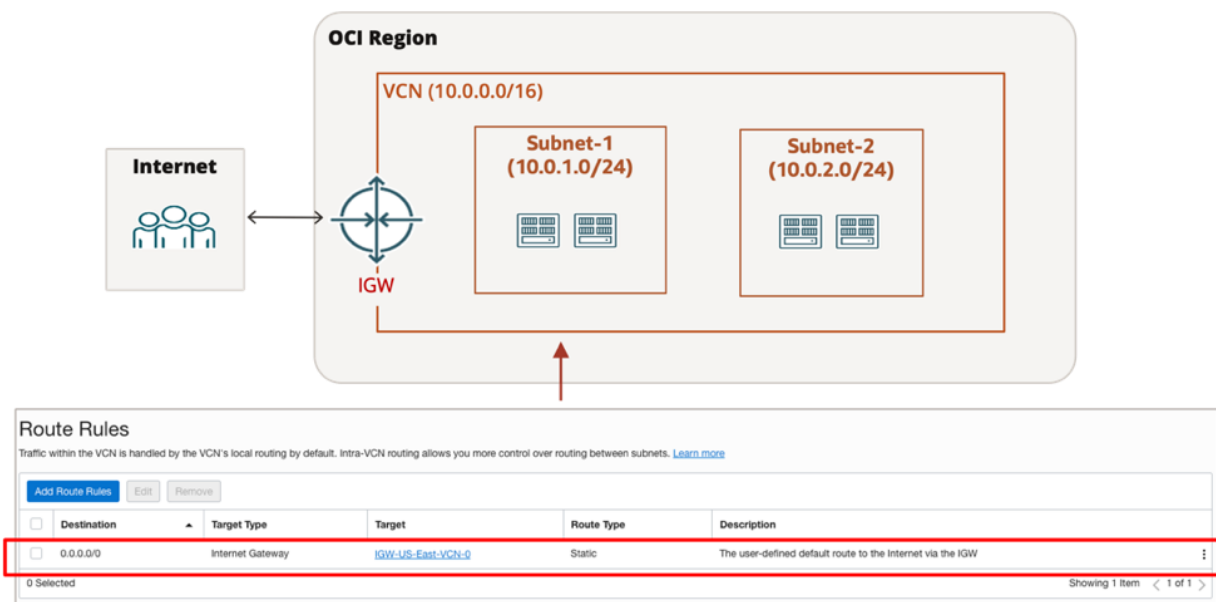


Figure 1. Example VCN Network

The example VCN has a CIDR of 10.0.0.0/16 with two subnets, Subnet-1 and Subnet-2. Subnet-1 has a route table with only a 0.0.0.0/0 route with the internet gateway (IGW) of the VCN as the target.

When a resource in Subnet-1 sends traffic to any destination, the route rules in this route table and the implicit local route rule for the VCN CIDR 10.0.0.0/16 are evaluated against the LPM algorithm for the best route selection. One of the following routing decisions is applied based on the location of the destination:

- If the destination is inside the VCN, the local route for the VCN CIDR 10.0.0.0/16 is the LPM route. OCI uses this local route to directly send the traffic to the destination inside the VCN.
- If the destination is outside the VCN, the 0.0.0.0/0 route is the LPM route and is used to route the traffic to the IGW because the IGW is the next-hop target in the configuration of the 0.0.0.0/0 route rule.

When you're evaluating a route table by yourself, it's important to consider the implicit local routes for the VCN CIDRs because they're evaluated as equally as the user-defined routes in the LPM route calculation.

Another rule in the OCI route table operation that you need to know is that users can define their own route rules for the VCN CIDR, and a user-defined route rule for the VCN CIDR is preferred over the implicit local route rule for the same VCN CIDR. This rule applies to all the routing scenarios.

To continue with the same example, let's add a firewall with a private IP address of 10.0.99.104 in a subnet with the CIDR 10.0.99.0/24, and define a route rule for the VCN CIDR 10.0.0.0/16 with a firewall private IP address of 10.0.99.104 as the target. As depicted in Figure 2, the route table shows both the default route 0.0.0.0/0 and the route rule for 10.0.0.0/16.

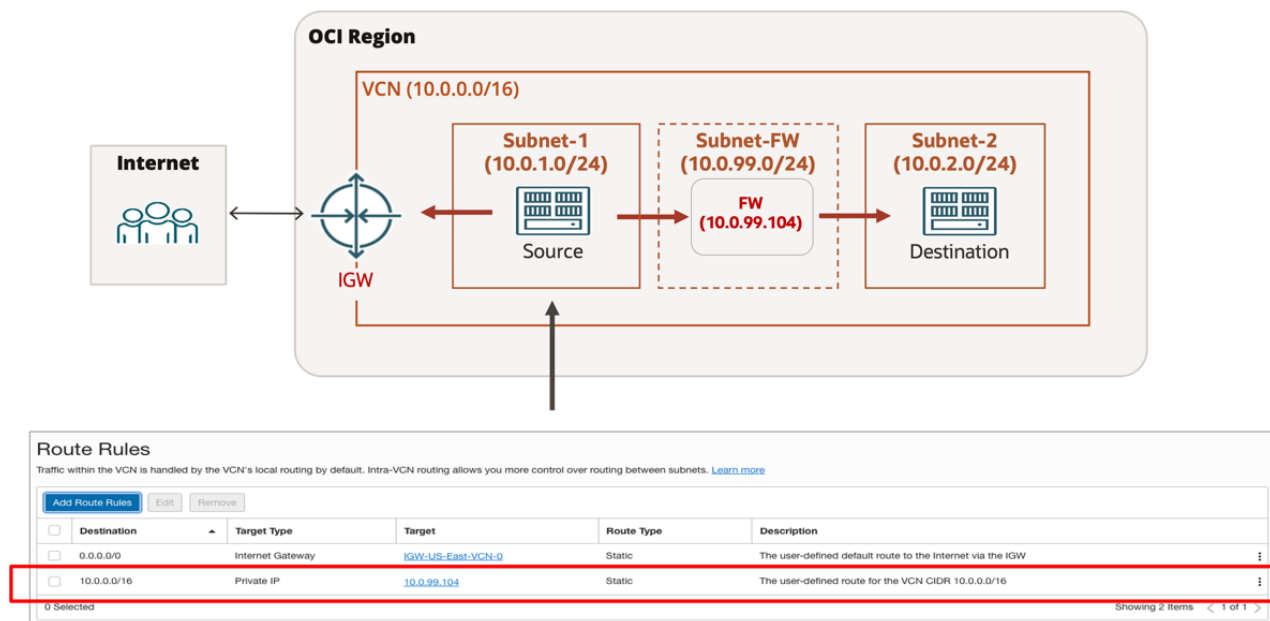


Figure 2. Example VCN Network with a Firewall

Now when OCI exercises its best route selection process, it sees both the implicit local route rule and the user-defined route rule as the LPM routes for the VCN CIDR 10.0.0.0/16 because they have the same prefix and prefix length. OCI prefers the user-defined route rule over the local route rule. Traffic sourced from Subnet-1 is subject to the following routing logic:

1. Traffic going to destinations in any other subnets in the same VCN is routed to the firewall first based on the user-defined route rule for the VCN CIDR 10.0.0.0/16 with the firewall private IP address as the target.

Note: This scenario illustrates intra-VCN routing, which is discussed in more detail in the next section.

2. Traffic going outside the example VCN is routed to the IGW based on the route rule for 0.0.0.0/0.

Intra-VCN Routing

To simplify routing in VCNs, OCI implemented direct local routing as the default within a VCN. That is, traffic between endpoints within the same VCN is directly routed from the source to the destination by default, using the OCI-created implicit local route rule for the VCN CIDRs. With this default routing behavior, users don't need to configure routing for their endpoints within a VCN.

But if you want to insert an intermediate hop (usually a network virtual appliance, for example, a virtual firewall) between the source and the destination in the same VCN, you can put the source and the destination into two separate subnets in the VCN and then define your own route rules to route the subnet-to-subnet traffic through the intermediate appliance.

Intra-VCN route rules are defined in the VCN route table associated with a subnet. That route table controls how traffic sourced from the subnet is routed towards the destination. The intra-VCN route rules in a subnet VCN route table can have the following options for its target:

- A private IP address in the VCN
- A dynamic routing gateway (DRG) that the VCN is attached to
- A local gateway (LGW) of the VCN

Figure 3 shows an example of using intra-VCN subnet-to-subnet routing to insert a firewall between the web-tier and app-tier subnets of an application.

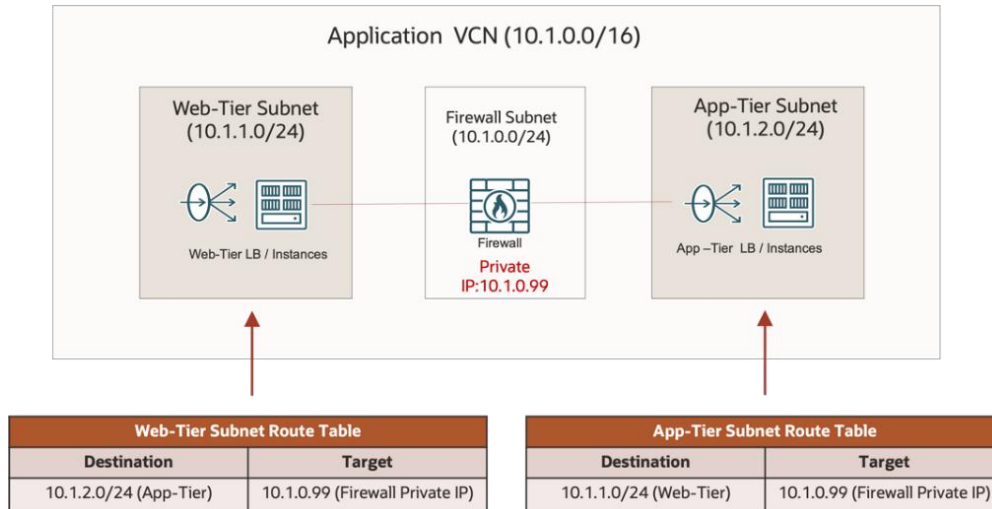


Figure 3. Firewall Insertion Using Intra-VCN Routing

In this example, the web-tier subnet CIDR is 10.1.1.0/24 and the app-tier subnet CIDR is 10.1.2.0/24. The subnets are in the same VCN CIDR 10.1.0.0/16. The subnets have a user-defined route rule for each other in their route tables, with the firewall private IP address 10.1.0.99 as the target. OCI routes traffic in both directions between these two subnets to the firewall, which is in a separate subnet. The firewall subnet uses the system-defined local route rule for the VCN CIDR 10.1.0.0/16 to route the traffic to the web-tier and the app-tier subnets.

In the routing lookup process for the subnet-to-subnet traffic, the intra-VCN subnet route and the implicit local route for the VCN CIDR are both evaluated, and the route for the subnet CIDRs is used because it's the LPM for the destination IP address. The traffic is sent to the target (the firewall) based on the route rules for the destination subnets. If there are no user-defined intra-VCN route rules for a destination subnet, the local route is used as the LPM route to route the traffic directly from the source subnet to the destination subnet.

Note: The user-defined intra-VCN route rules are for routing traffic *between* subnets within a VCN. If the source and the destination of the traffic are in the same subnet, the traffic is always directly routed. You don't need to define routes, and you can't alter this direct routing behavior for the traffic within a subnet. This is why the web tier and the app tier in the example need to be in different subnets in order to insert a firewall between them.

VCN Egress Routing

VCN egress routing refers to the routing process for traffic from a source in a VCN to a destination outside of the source VCN. The destination can be in another VCN, in an on-premises network, in an OCI service, or on the internet. This routing process is performed by using the *source subnet route table*, which is the VCN route table associated with the source subnet.

A VCN subnet always has one route table associated with it. When it's created, a subnet is associated with the system-created VCN default route table, but you can change it to use a user-defined route table. Traffic originated by resources in a subnet to destinations outside of the subnet is routed based on the route rules in the subnet route table.

Depending on the location of the destination, a route rule in a subnet route table can have different types of targets. The following table displays the target types for each destination location.

Table 1. VCN Subnet Route Table: Routing Traffic from a Subnet to a Destination Outside the Subnet

DESTINATION	TARGETS IN THE ROUTE RULE
Another subnet in the same VCN	LPG, DRG, or private IP address
Remote VCN or subnet CIDR	LPG, DRG, or private IP address
A public IP address or an IP address prefix in the internet	NAT gateway (NATGW), IGW, or a private IP address
OCI services using the service labels in the same region	SGW, DRG, or a private IP address

Note: You can add route rules to access the OCI services in the same region as the VCN by using *service labels*. Each region currently has two service labels: one for the OCI Object Storage service in the region and one for all OCI services in the region. When you use a service label as the destination of a route rule, the target can be the SGW of the VCN or a private IP address in the VCN. On the backend, a route rule with a service label as the destination is mapped to route rules for the public IP address prefixes of the corresponding services in the region (either OCI Object Storage or all OCI services). To avoid asymmetric routing for OCI services, a VCN route table isn't allowed to have both a route rule to all OCI services using the all-service label with an SGW as the target *and* a route rule for CIDR 0.0.0.0/0 with an IGW as the target at the same time.

OCI uses internet gateways (IGWs) and NAT gateways (NATGWs) to route traffic from a VCN subnet to the internet.

- An IGW is usually used to route traffic from public subnets to the internet. When routing the outbound traffic that's initiated by a resource in the VCN, the IGW modifies the source IP address in the IP packets from the resource's private IP address to its public IP address.
- A NATGW can route traffic from a private subnet to the internet for communication sessions that are initiated within the subnet. Resources in a private subnet don't have public IP addresses. A NATGW performs stateful Port Address Translation (PAT) when routing the outbound traffic. It translates the source IP address and source port combination in the original outbound packets to a unique combination of its own public IP address and an unused source port.

VCN Gateway Ingress Routing

The inbound traffic to a VCN is traffic that is sourced from outside the VCN and targets a resource or IP address inside the VCN. With the exception of traffic that is routed through a private endpoint, traffic is routed into a VCN by a *gateway*. Depending on the location of the source, the gateway could be an internet gateway (IGW), a NAT gateway (NATGW), a local peering gateway (LPG), a dynamic routing gateway (DRG), or a service gateway (SGW). The following table shows the gateway types for ingress routing from different source locations outside of the destination VCN.

Table 2. Gateways for Ingress Routing

SOURCE OF THE INGRESS TRAFFIC	GATEWAY TO ROUTE THE TRAFFIC INTO THE DESTINATION VCN
On the internet	IGW or NATGW
Another VCN	LPG or DRG
OCI services	SGW (through SGW private access) or DRG
On-premises networks	DRG

By default, a gateway routes the ingress traffic to the destinations in a VCN directly. But if you associate a VCN route table with a gateway, the gateway then uses the route rules in this route table to route ingress traffic. The route table can be the VCN default route table or a route table that you create. Note that this route table operation on the gateway for ingress traffic routing is independent of the subnet route table. The subnet route table is for routing only traffic that's going outside of the subnet. It's in the opposite direction of the gateway ingress routing, and the routing next hop can be one of the gateways if the traffic is going outside of the VCN.

Because different gateways route traffic into a VCN in different contexts (which is discussed in detail for each gateway type in a later part of this section), you may have different intentions for how the traffic should be routed. We recommend that you use separate route tables for the different gateways to capture your intention in the various routing contexts. For example, you might want the IGW to route the public traffic to destinations inside your VCN through a firewall for security inspection first because sources on the internet are considered untrusted, but you might want the DRG to route the ingress traffic from another VCN to the same destinations directly because sources in another VCN can be considered internal and trusted. In this scenario, the IGW and the DRG use different route tables, which lets you define different route rules for the same destinations.

Each gateway type is for routing the ingress traffic in certain scenarios. Therefore, they each have certain allowed next hops. The following table shows the supported target types (the next hop for routing) for each gateway.

Table 3. Gateway Ingress Route Target Types

GATEWAY TYPE	ROUTE TARGET TYPES
IGW	A private IP address in a public subnet of the VCN
NATGW	A private IP address in the VCN
LPG	A private IP address in the VCN, or a DRG
DRG	A private IP address in the VCN, or an LPG or SGW
SGW	A private IP address in the VCN, or a DRG

Internet Gateway Ingress Routing

An internet gateway (IGW) is used to route traffic from the internet to destinations in VCN public subnets. The destinations have a public IP address that is allocated from Oracle-owned public IP address prefixes or a customer's BYOIP (Bring Your Own IP) address block, and a private IP address that's allocated from the public subnet CIDR. The sources on the internet use the public IP address as the destination. When the traffic is routed to the IGW of the VCN, the IGW first translates the public IP address to the private IP address of the destination endpoint, and then routes the traffic toward the destination based on its routing lookup result for the private IP address.

By default, an IGW doesn't have a route table associated with it; it directly routes the traffic to the destination in the VCN. However, you can associate a VCN route table with an IGW and define custom routes for the IGW to route ingress traffic into the VCN. When a route table is associated with it, an IGW uses the LPM evaluation process to go through the route rules in the route table (all user-defined routes and the implicit VCN CIDR local route) to select the best route and then routes the traffic to the target of the best route. The target for a route rule in the IGW route table can be a private IP address in a public subnet in the VCN. Usually, it's a network virtual appliance (NVA).

Following are typical IGW ingress routing scenarios:

- IGW sends the traffic directly to the destination** (as shown in Figure 4): In this scenario, the IGW doesn't have an associated route table, or it has an associated route table but the table has no route rules defined for the VCN CIDR or the destination subnet CIDR. The IGW uses the implicit local route for the VCN CIDR to route the traffic directly to the destination.

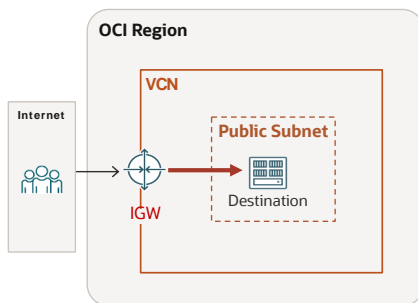
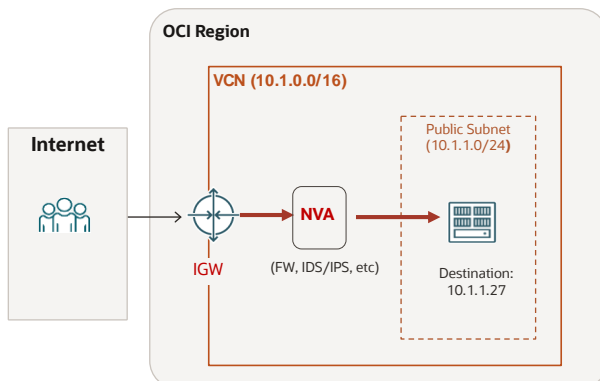


Figure 4. IGW Routes Ingress Traffic Directly to the Destination by Using the Local Route for the VCN CIDR

- IGW sends the traffic to an NVA instead of the destination** (as shown in Figure 5): In this scenario, the IGW has a route table that contains a route rule for the destination subnet CIDR or the VCN CIDR with a private IP address of a network virtual appliance (NVA) as the target. The NVA is usually a firewall, an intrusion detection system (IDS) or intrusion prevention system (IPS), or a router. After processing the traffic based on its configuration, the NVA forwards the traffic toward the destination. Usually, it's configured to route the traffic back to the virtual gateway of the VCN subnet that it resides in. The VCN subnet then uses its route table to route the traffic toward the destination.

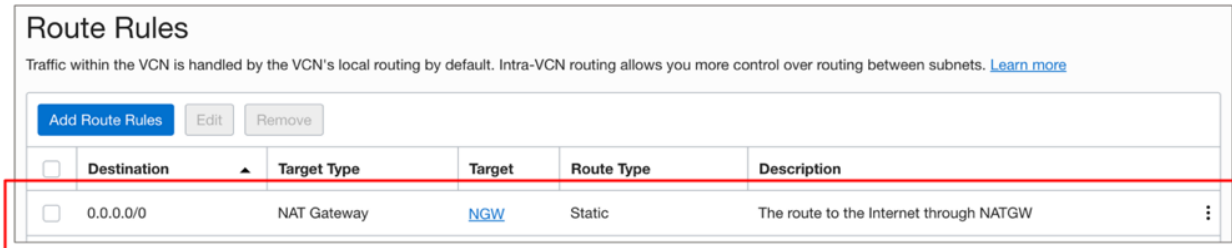


The IGW Route Table	
Destination	Target
10.1.1.0/24 (The destination subnet CIDR) or 10.1.0.0/16 (The VCN CIDR)	10.1.0.99 (The private IP of the NVA)

Figure 5. IGW Routes Ingress Traffic to an NVA by Using a User-Defined Route

NAT Gateway Ingress Routing

A NAT gateway (NATGW) provides a secure way to route traffic between a private subnet in the VCN and the internet. The communication session must be initiated by a resource in a VCN subnet. The traffic from the initiating resource in the VCN is routed to the NATGW by using route rules in the source subnet route table. For example, as shown Figure 6, the route rule can be a default rule with the NATGW as the target. It is a route to reach all IP resources (0.0.0.0/0) on the internet through this NATGW.



<input type="checkbox"/>	Destination	Target Type	Target	Route Type	Description	
<input type="checkbox"/>	0.0.0.0/0	NAT Gateway	NGW	Static	The route to the Internet through NATGW	⋮

Figure 6. An Example VCN Route Table with a Default Route Using a NATGW as the Target

Each NATGW has a public IP address. For the outbound traffic, the NATGW first uses NAT to translate the source private IP address and the source port to its own public IP address and an unused port. Then, it routes the traffic out to the destination on the internet. When the session is initiated from within a VCN, NATGW then allows and routes the return traffic in this session from the internet back to the resource.

After receiving the return traffic from the internet, the NATGW first changes the destination IP address and the destination port back to the original private IP address and the original port, and then routes it toward the destination (the resource in the VCN that initiated the session).

By default, a NATGW doesn't have a route table associated with it; it uses the local route for the VCN CIDR to route traffic directly to the destination in the VCN. However, you can associate a route table with a NATGW. The NATGW then uses the route rules in the associated route table to route the ingress traffic into the VCN. The targets of route rules in a NATGW route table can be private IP addresses in the VCN. If the associated route table has a user-defined route for the destination subnet or the VCN CIDR, the NATGW routes the traffic to the target of this route. If there's no user-defined route, the NATGW uses the implicit VCN CIDR local route in the route table to route the traffic directly to the destination.

Following are typical NATGW ingress routing scenarios:

- **NATGW sends the traffic directly to the destination** (as shown in Figure 7): In this scenario, the NATGW doesn't have an associated route table, or it has an associated route table but the table has no route rules defined for the VCN CIDR or the destination subnet CIDR. The NATGW uses the implicit local route for the VCN CIDR to route the traffic directly to the destination.

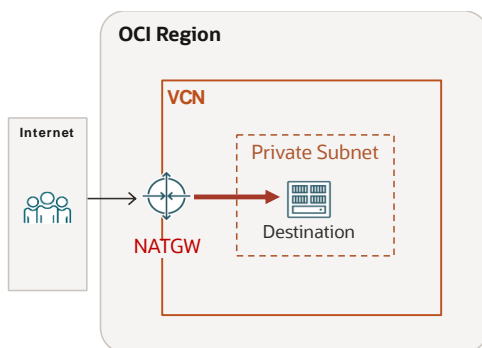


Figure 7. NATGW Routes Ingress Traffic Directly to The Destination by Using the Local Route for the VCN CIDR

- **NATGW sends the traffic to an NVA instead of the destination** (as shown in Figure 8): In this scenario, the NATGW has a route table that contains a route rule for the destination subnet CIDR or the VCN CIDR with a private IP address of a network virtual appliance (NVA) as the target. The NVA is usually a firewall, an IDS or IPS, or a virtual switch or router. After processing the traffic packets based on its configuration, the NVA forwards the traffic toward the destination. Usually, it's configured to route the traffic back to the gateway of the VCN subnet that it resides in. The VCN subnet then uses its route table to route the traffic toward the destination.

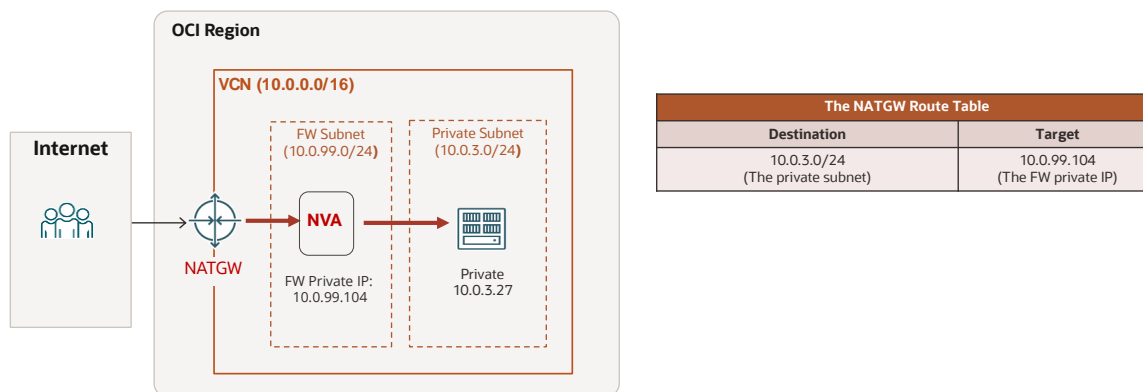


Figure 8. NATGW Routes Ingress Traffic to an NVA by Using a User-Defined Route

Service Gateway Ingress Routing and Transit Routing

You can use a service gateway (SGW) for both ingress routing and transit routing.

SGW Ingress Routing

An SGW routes traffic between a customer VCN and OCI services. It provides customer resources in customer VCNs with private access to OCI services through the OCI internal network (called the Oracle Services Network). SGW ingress routing occurs when an SGW routes traffic from an OCI service to destinations in its VCN.

Like IGWs and NATGWs, by default an SGW doesn't have a route table associated with it. It uses the implicit VCN CIDR local route to route ingress traffic directly to destinations in the VCN. To use SGW ingress routing with user-defined routes, you can associate it with a route table—either a route table that you create or the default VCN route table—and define the route rules that you want the SGW to use for ingress routing in this route table. The route rules in an SGW route table can have a private IP address or a DRG as the target.

With SGW ingress routing for VCN internal destinations, you can insert a network virtual appliance, such as a virtual firewall, into the forwarding path of the ingress traffic from OCI services to the VCN destinations. Figure 9 shows an example in which the SGW routes all ingress traffic from OCI services destined to the app subnet through a firewall.

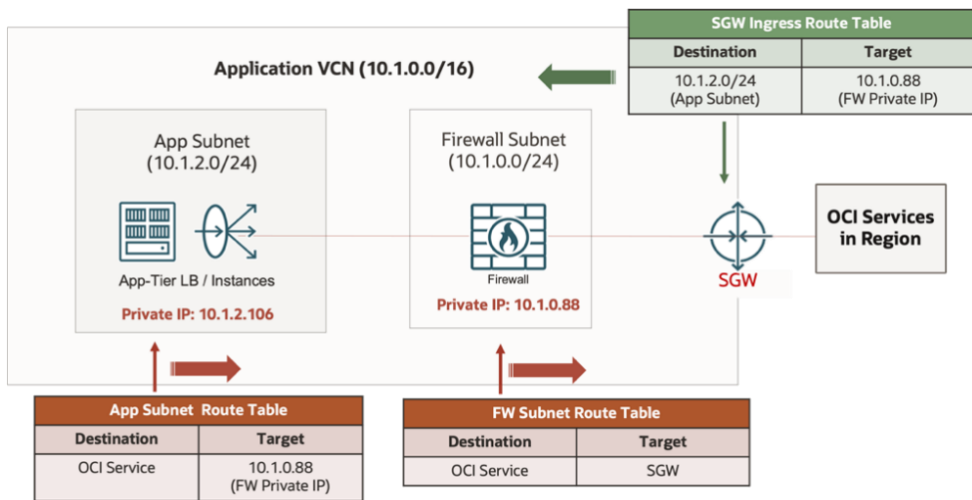


Figure 9. SGW Routes Ingress Traffic from OCI Services Through a Firewall

Transit Routing with an SGW

In addition to defining route rules to VCN internal destinations, you can define transit routes in an SGW route table. A transit route is for destinations that are outside of the VCN. The transit route supported by SGW ingress routing is used in the typical transit routing design for on-premises networks to privately access OCI services through their OCI network. Figure 10 shows a typical topology for such a transit routing design. Note that the route rule for the on-premises network 10.254.0.0/16 in the SGW ingress route table has the DRG as the target. The traffic from the OCI services traverses the Hub VCN through the SGW and DRG to reach the on-premises destination.

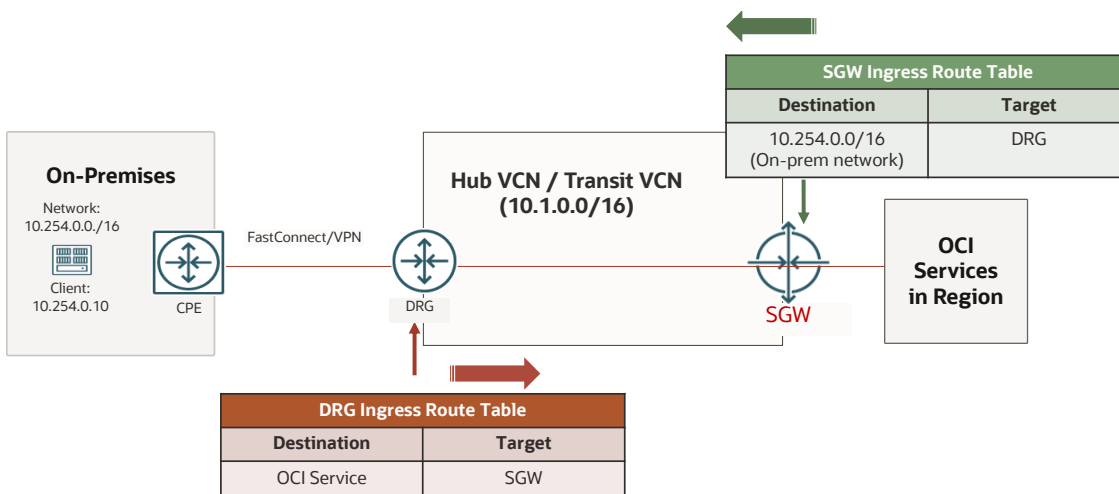


Figure 10. Transit Routing with SGW

For details about the design and configuration of transit routing to the Oracle Services Network, see the following OCI resources:

- [Transit Routing to the Oracle Services Network](#) (blog post)
- [Private Access to Oracle Services](#) (documentation)

You can also insert a firewall into this transit routing for OCI services private access design, as shown in Figure 11.

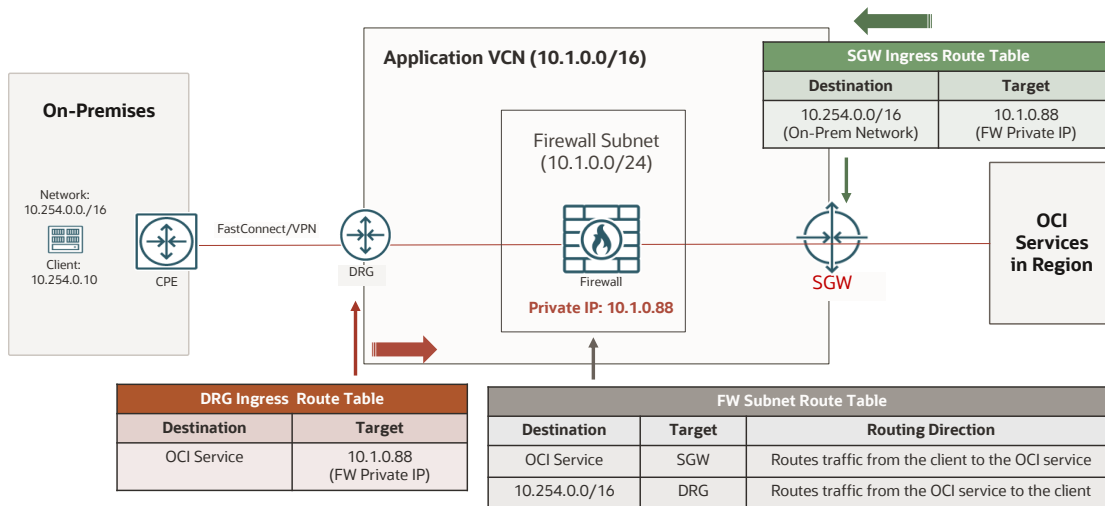


Figure 11. Transit Routing with SGW and Firewall Insertion

Local Peering Gateway Ingress Routing and Transit Routing

You can use a local peering gateway (LPG) for both ingress routing and transit routing.

LPG Ingress Routing

An LPG connects two VCNs in the same region. For VCN ingress routing, an LPG of a VCN receives traffic from the peered LPG in a remote VCN and routes the traffic to the destination in its VCN or to a DRG that its VCN is attached to. The latter is typically used in a transit routing design.

Unlike the gateways discussed previously (IGW, NATGW, and SGW), LPG ingress routing doesn't support user-defined routing for destinations in the VCN. You can't define route rules for prefixes inside the VCN CIDR in the route table for an LPG. For traffic going to destinations inside the VCN, the LPG sends it directly to the destination by using the local route for the VCN CIDR. Figure 12 shows an example of direct local routing by an LPG for ingress traffic to destinations in the VCN.

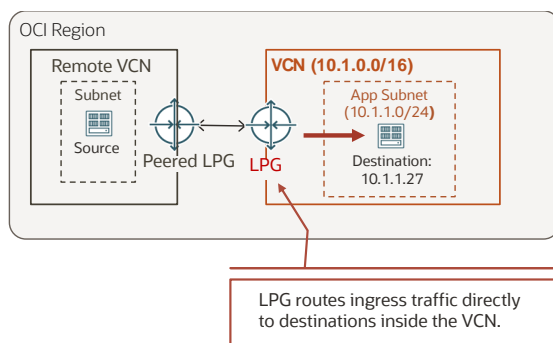


Figure 12. LPG Routes Ingress Traffic Directly to Destinations in the VCN

Transit Routing with an LPG

In the route table associated with an LPG for its ingress routing, you can define route rules for prefixes that are outside of the VCN with a private IP address in the VCN or a DRG as the target. These route rules are for transit routing, that is, for the LPG to route traffic for a destination that is not in its VCN.

Figure 13 shows a typical transit routing design between an on-premises network and a VCN network. The details of this design are provided in a later section, “Transit Routing with a DRG.” The illustration here shows that LPG-1 in the hub VCN uses its ingress route table to route traffic from a spoke VCN to the on-premises network through the hub VCN by way of the DRG that both the hub VCN and the FastConnect circuits or VPN tunnels are attached to. There are usually multiple on-premises prefix routes in the DRG route table. For demonstration purpose, we show the route for only one on-premises prefix, 10.254.0.0/16, in the illustrations.

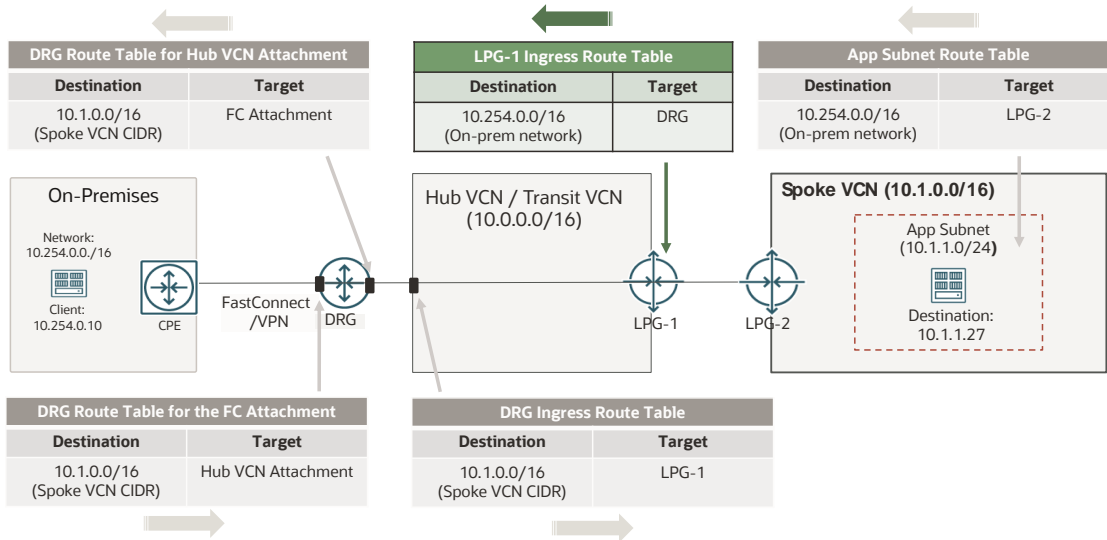


Figure 13. Transit Routing with LPG

In this design, you can also add a firewall in the hub VCN in the routing path. As depicted in Figure 14, the target of the route rules in the LPG ingress route table is the private IP address of the firewall. The firewall subnet route table has a route for the on-premises prefix, 10.254.0.0/16, with the DRG as the target.

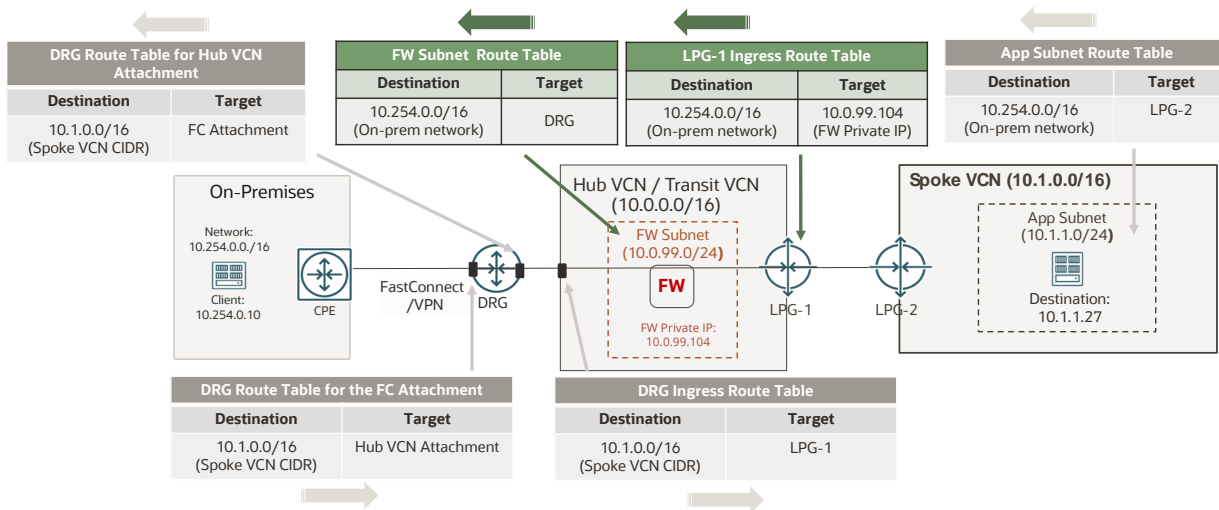


Figure 14. Transit Routing with LPG and Firewall Insertion

VCN Ingress Routing Through a Dynamic Routing Gateway

A dynamic routing gateway (DRG) provides connectivity among VCNs and connects VCNs with on-premises networks, remote regions, or other clouds. A DRG uses its route tables to route traffic between its *attachments*. It uses the DRG route table associated with the ingress attachment of the traffic to decide the egress attachment where it needs to forward the traffic. If the DRG decides that the traffic needs to be sent to a VCN attachment, the next routing decision point is when the traffic enters that VCN. How the VCN routes the traffic toward its destination is controlled by the VCN ingress routing through the DRG.

When a VCN is attached to a DRG, the DRG is represented as a DRG attachment in the VCN. Traffic is routed into the VCN through the DRG attachment. By default, a DRG attachment in a VCN doesn't have a VCN route table associated with it, and it uses the implicit local route for the VCN CIDRs to route traffic directly to the destinations in the VCN. But you can associate a VCN route table with the DRG attachment and define route rules in this VCN route table to control the ingress routing through the DRG attachment.

Note: The VCN route table for VCN ingress routing through the DRG is separate from the DRG route table for the VCN attachment. The former controls how traffic is routed into or through the VCN through the DRG, whereas the latter controls how the DRG routes the traffic from the VCN. DRG routing is covered in detail later in “DRG Routing Basics” and accompanying sections. In this section, we focus on the VCN ingress routing through the DRG.

Routing Traffic into a VCN Through a DRG

In the VCN route table for a DRG attachment in a VCN, you can define route rules for traffic destined to resources in the VCN with a private IP address in the VCN as the target. This enables you to insert a network virtual appliance, such as a virtual firewall, in the routing path between the DRG and the destination in the VCN. If the VCN route table for the DRG attachment doesn't have a route rule defined for a destination in the VCN, or the DRG attachment doesn't have a route table associated with it, the implicit local route for the VCN CIDR is used by the DRG to route the ingress traffic directly to the destination.

Figure 15 shows an example of VCN ingress routing through a DRG in which the ingress traffic to local resources in the VCN is sent to a firewall first.

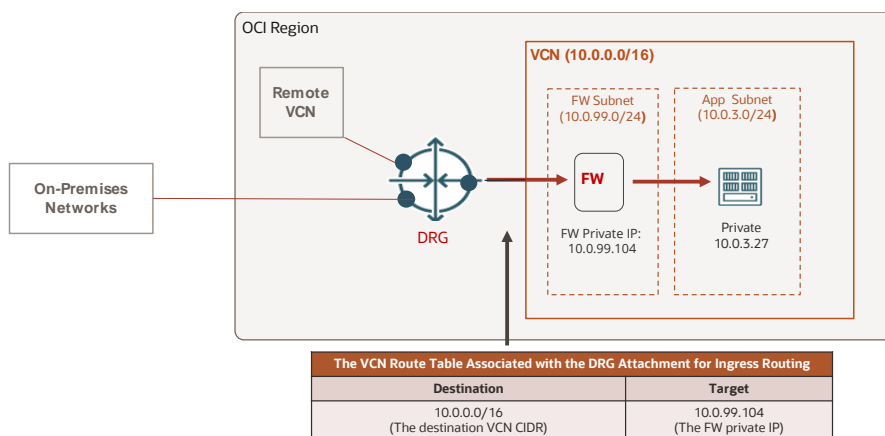


Figure 15. Routing Traffic into a VCN Through a DRG

As illustrated in the diagram, the traffic entering a VCN through a DRG is from the other attachments on the same DRG. The source of the traffic can be another VCN, or an on-premises network through FastConnect circuits or VPN tunnels. It can also be from VCNs in a remote region through a remote peering connection (RPC). These designs with the DRG are discussed in detail later in this paper. The focus in this example is how the DRG routes traffic into a VCN.

Transit Routing with a DRG

In addition to route rules to VCN internal destinations, you can define transit routes in a VCN route table associated with a DRG. A transit route is for destination prefixes that are outside of the VCN. The transit route rules in the VCN route table for a DRG can have a private IP address, an LPG, or an SGW as the target. The transit route support by DRG ingress routing is used in the typical transit routing design for on-premises networks to communicate with multiple OCI VCNs or for on-premises networks to privately access OCI services through a VCN's SGW.

Figure 16 shows a typical topology for such a transit routing design to allow an on-premises network to reach a spoke VCN through the hub VCN that's attached to the DRG. DRG ingress routing is used for the DRG to route the traffic from the on-premises network to the destination spoke VCN through the hub VCN. For the inverse routing in this scenario, see the "Local Peering Gateway Ingress Routing and Transit Routing" section.

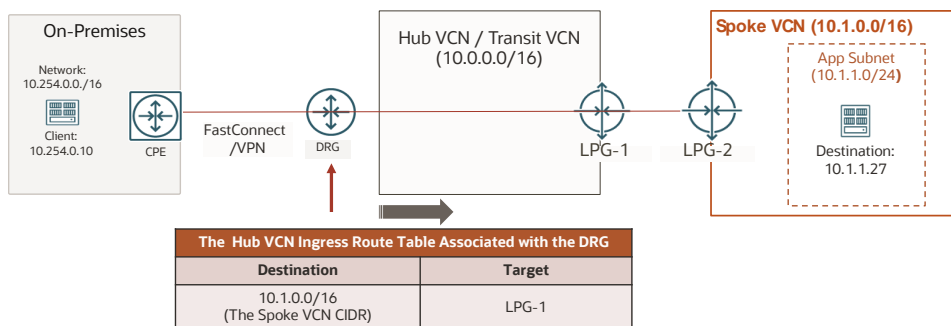


Figure 16. Transit Routing with DRG

The later part of this paper provides more discussions about transit routing or routing between on-premises networks and OCI VCN networks. The example here is to illustrate how the DRG ingress routing is used in such routing designs.

DRG Routing Basics

A DRG is a regional virtual router that interconnects VCNs in a region and connects the VCNs with on-premises networks through FastConnect virtual circuits or IPsec VPN tunnels. It also provides network connectivity between regions through remote peering connection (RPC).

A DRG acts like a central hub to connect the network resources that are attached to it. The network resources can be VCNs, site-to-site IPsec VPN tunnels, FastConnect virtual circuits, or RPCs. When a network resource is attached to a DRG, an attachment of the corresponding type is created:

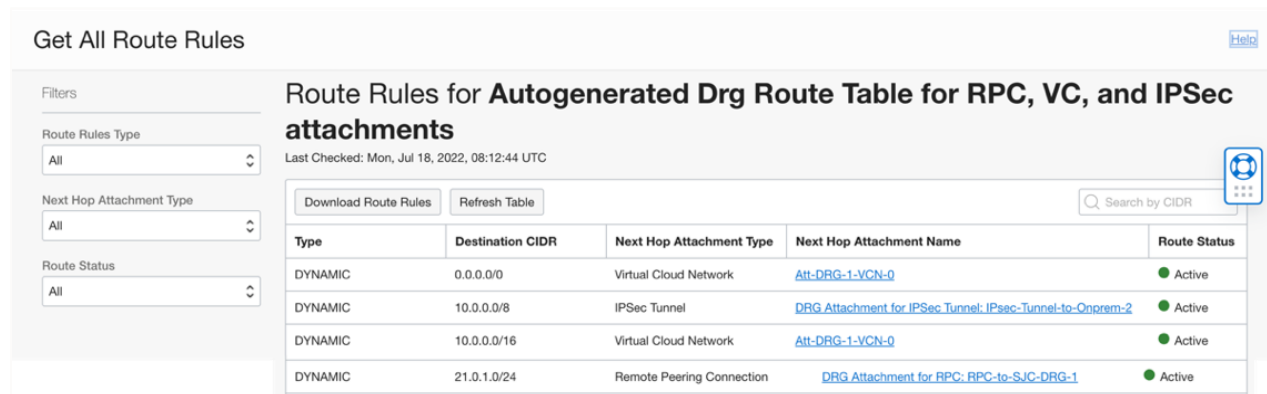
- **VCN attachment:** Created when a VCN is attached to the DRG
- **Virtual Circuit (VC) attachment:** Created when a FastConnect virtual circuit is attached to the DRG
- **IPsec tunnel attachment:** Created when an IPsec tunnel is attached to the DRG
- **RPC attachment:** Created when an RPC is attached to the DRG

A DRG routes traffic between the attachments by using DRG route tables. Each attachment is associated with a DRG route table. Traffic enters a DRG through an attachment and is routed to another attachment by the DRG based on the DRG route table associated with the ingress attachment of the traffic.

Route Tables in a DRG

A DRG uses *DRG route tables* to route traffic between its attachments. OCI automatically generates two route tables for each DRG: one for VCN attachments, and one for IPsec, VC, and RPC attachments. You can create more DRG route tables. The route rules in a DRG route table contain the route type, the destination CIDR, and the next-hop attachment type and name.

Figure 17 shows an example of a DRG route table.



Get All Route Rules Help

Filters

Route Rules Type: All

Next Hop Attachment Type: All

Route Status: All

Route Rules for Autogenerated Drg Route Table for RPC, VC, and IPsec attachments

Last Checked: Mon, Jul 18, 2022, 08:12:44 UTC

Download Route Rules Refresh Table

Search by CIDR

Type	Destination CIDR	Next Hop Attachment Type	Next Hop Attachment Name	Route Status
DYNAMIC	0.0.0.0/0	Virtual Cloud Network	Att-DRG-1-VCN-0	Active
DYNAMIC	10.0.0.0/8	IPSec Tunnel	DRG Attachment for IPSec Tunnel: IPSec-Tunnel-to-Onprem-2	Active
DYNAMIC	10.0.0.0/16	Virtual Cloud Network	Att-DRG-1-VCN-0	Active
DYNAMIC	21.0.1.0/24	Remote Peering Connection	DRG Attachment for RPC: RPC-to-SJC-DRG-1	Active

Figure 17. Example of a DRG Route Table

The route type can be dynamic or static. Dynamic routes are imported from the DRG attachments. Static routes are created by users through the API or the Oracle Cloud Console. The next hop of a route rule in a DRG route table is the DRG attachment of the network where the destination resides or in route to the destination. For a static route in a DRG route table, the next-hop attachment can be a VCN attachment, a cross-regional RPC attachment, or a cross-tenancy RPC attachment. It can't be a VPN attachment or FastConnect VC attachment.

Each DRG attachment has one DRG route table associated with it. By default, it's the autogenerated DRG route table for the attachment type. You can change it to a user-created DRG route table.

When traffic gets onto a DRG, the DRG performs ingress routing lookup based on the DRG route table associated with the ingress attachment of the traffic. The routing lookup resolves the next-hop attachment, which is the egress attachment. The DRG then sends the traffic onto the egress attachment through which the traffic gets to the next-hop network. There is no routing lookup at the egress attachment on the DRG.

Route Preference in DRG Route Table

A DRG route table could have multiple routes for the same prefix and mask. The DRG has a built-in mechanism for resolving such route conflicts. The decision is made based on the following route preference, and the evaluation is done in this exact order:

1. In a DRG route table, static routes have higher preference than dynamic routes.
2. Among the dynamic routes in a DRG route table, routes with a shorter AS path are preferred over routes with a longer AS path.

Routes with a route source of VCN or Static always have an empty AS path. Routes with a route source of IPsec VPN tunnel or FastConnect VC have the AS paths shown in the following table.

Table 4. AS Paths for IPsec VPN Tunnel or FastConnect Virtual Circuit Route Sources

ROUTE SOURCE	DETAILS OF HOW ORACLE PREFERS THE PATH	RESULTING AS PATH FOR THE ROUTE
FastConnect VC	OCI prepends no ASNs to the routes, which results in a total AS path length of 1.	Customer ASN
Site-to-site VPN with border gateway protocol (BGP) routing	OCI prepends a single private ASN on all the routes that the customer edge device advertises over site-to-site VPN with BGP, for a total AS path length of 2.	Private ASN, Customer ASN
Site-to-site VPN with static routing	OCI advertises those static routes to the DRG as BGP dynamic routes. OCI prepends 3 private ASNs on these routes, which results in a total AS path length of 3.	Private ASN, Private ASN, Private ASN

3. The attachment type that imported the route is evaluated according to the following priority based on the attachment type:
 - A. **VCN**
 - B. **VIRTUAL_CIRCUIT**: If equal-cost multipath (ECMP) is disabled for the DRG route table, the DRG makes an arbitrary but stable selection. If ECMP is enabled, all routes are added to the route table and the DRG uses ECMP to make routing choices. The maximum supported ECMP width inside a DRG is 8.
 - C. **IPSEC_TUNNEL**: If ECMP is disabled for the DRG route table, the DRG makes an arbitrary but stable selection. If ECMP is enabled, all routes are added to the route table and the DRG uses ECMP to make routing choices. The maximum supported ECMP width inside a DRG is 8.
 - D. **REMOTE_PEERING_CONNECTION (RPC)**: The DRG chooses the route with the lowest network distance. If two routes have identical network distances, the DRG selects the route with the highest priority route source (STATIC > VCN > VIRTUAL_CIRCUIT > IPSEC_TUNNEL).

If two routes have the same route source, the DRG makes an arbitrary but stable selection.
4. If conflicting routes are imported from attachments of the same type, the conflict is resolved differently depending on the attachment type:
 - **VCN attachments**: If identical CIDRs are imported from two VCN attachments, only one is selected by using an arbitrary but stable decision procedure.
 - **VIRTUAL_CIRCUIT and IPSEC_TUNNEL attachments**: If multiple routes with the same CIDR and different AS_PATH lengths are imported into a DRG route table, the route with the lowest AS_PATH length is selected. Otherwise, one route is chosen by using an arbitrary but stable decision procedure.
 - **RPC attachments**: If identical CIDRs are imported from two RPC attachments, one of them is chosen by using an arbitrary stable decision procedure.

Note: The preceding content is from the Oracle document about DRG route conflicts resolution at docs.oracle.com/iaas/Content/Network/Tasks/managingDRGs.htm#managingDRGs_topic_drg_routing.

Route Propagation and Import Route Distribution Control on a DRG

Network resources such as VCNs, FastConnect VCs, and IPsec VPN tunnels can be attached to a DRG. The routes associated with these network resource are propagated to the DRG. They can then be imported into a DRG route table as dynamic routes with an import route distribution policy.

Route Propagation on a DRG

This section shows what routes are associated with the network resource of each DRG attachment type and are propagated to the DRG.

VCN Attachment

The routes in the VCN route table that is associated with the DRG attachment in the VCN, and the VCN CIDRs and its subnet CIDRs. As discussed in the “VCN Ingress Routing Through a Dynamic Routing Gateway” section, after a VCN is attached to a DRG, the DRG is represented as a DRG attachment in the VCN. A route table of the VCN can be associated with the DRG attachment for the VCN ingress routing through the DRG. The routes in this VCN route table are the ones that are propagated to the DRG. If no VCN route table is associated with the DRG attachment, only the VCN CIDRs and its subnet CIDRs are propagated to the DRG.

When these routes are imported into a DRG route table, this VCN attachment is the next-hop attachment in the routes.

IPsec Tunnel Attachment

The routes advertised by the IPsec customer-premises equipment (CPE) when border gateway protocol (BGP) dynamic routing is used on the IPsec connection or the configured static routes if static routing is used on the IPsec connection.

When these routes are imported to a DRG route table as dynamic routes, the IPsec tunnel attachment is the next-hop attachment for the routes.

VC Attachment

The routes advertised by the FastConnect CPE through BGP.

When these routes are imported to a DRG route table, the VC attachment is the next-hop attachment in the routes.

RPC Attachment

All routes in the DRG route table associated with the remote DRG's RPC attachment are propagated to the local DRG.

When these routes are imported to a local DRG route table, the RPC attachment is the next-hop for the routes.

Import Route Distribution Control on a DRG

For a given DRG route table, you can create and apply an import route distribution policy to control which routes get imported to the route table. You can match by attachment type (for example, match all VCN attachments), a specific attachment, or match all.

The automatically generated DRG route table for VCN attachments has a default import route distribution that has a “match all” statement to import routes from all DRG attachments

The automatically generated DRG route table for IPsec, VC, and RPC attachments has a default import route distribution that has a “match type VCN” statement to import routes only from all VCN attachments. Note that this default import distribution policy doesn't import routes propagated by other attachment types into this DRG route table.

If you use the automatically generated DRG route table for all your VCN attachments, you achieve fully meshed routing connectivity among your VCNs, and all your VCNs have routes to reach all your on-premises networks and VCNs in the remote region.

If you want to establish more restricted routing connectivity or create routing segmentation in your network, you can use separate DRG route tables with different import route distribution policies for different DRG attachments. In the example shown in Figure 18, we created three routing segments on the same DRG by using different DRG route tables and different import route distributions for the VCNs:

- A fully meshed connectivity between VCN-1, VCN-2, and VCN-3
- Connectivity between VCN-4 and VCN-5
- Connectivity between VCN-6 and the on-premises networks

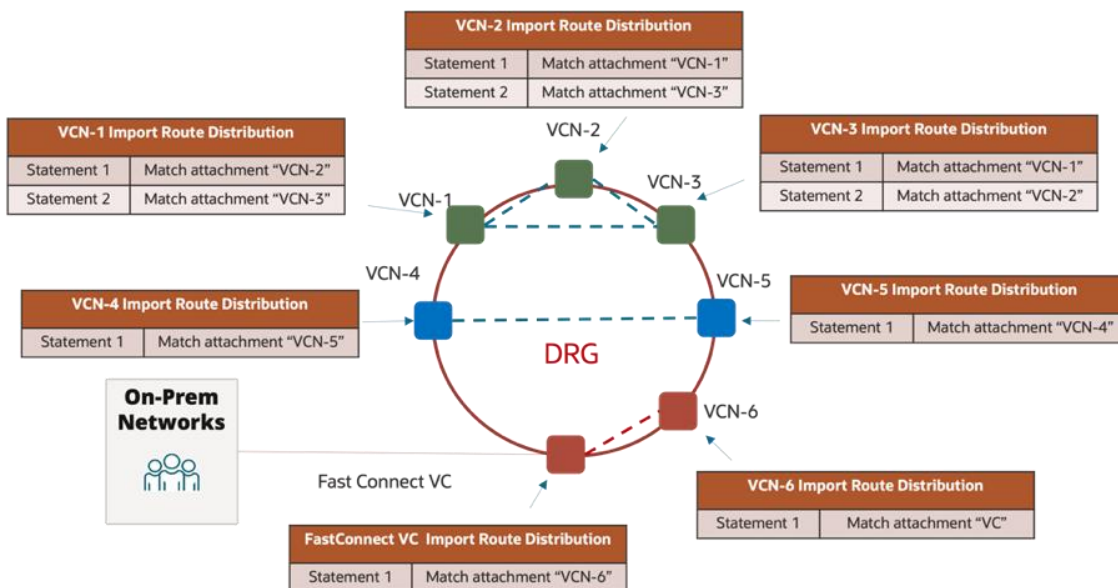


Figure 18. Example DRG Import Route Distribution Control

Although by default all the attachments use the automatically generated DRG route table for its type, real network designs often require some attachments of the same type to have different route rules or different route import distribution policies. It's a good practice to create separate DRG route tables for these attachments.

DRG Routing Operation

A DRG routes traffic between its attachments. For a given traffic flow, the DRG has an ingress attachment and an egress attachment. The DRG uses an ingress routing model: when traffic enters the DRG through the ingress attachment, the DRG uses the DRG route table associated with the ingress attachment to decide where the traffic goes.

If a route for the destination exists in the ingress attachment's DRG route table, the next hop of the route must be another attachment on the DRG. It could be a VCN attachment (if the destination is in a VCN), an IPsec or VC attachment (if the destination is in an on-premises network connected to the DRG through IPsec tunnels or FastConnect), or an RPC attachment (if the destination is in a remote region). If no matching route exists for the destination, the traffic is dropped.

Figure 19 shows an example for DRG route operation. It shows the DRG routing lookup for traffic that comes from Attachment-1 and goes to a destination network that is on Attachment-2. The routing lookup takes place in the DRG route table of the ingress attachment (Attachment-1). The route table has a route rule for the destination with the egress attachment (Attachment-2) as the next-hop attachment.

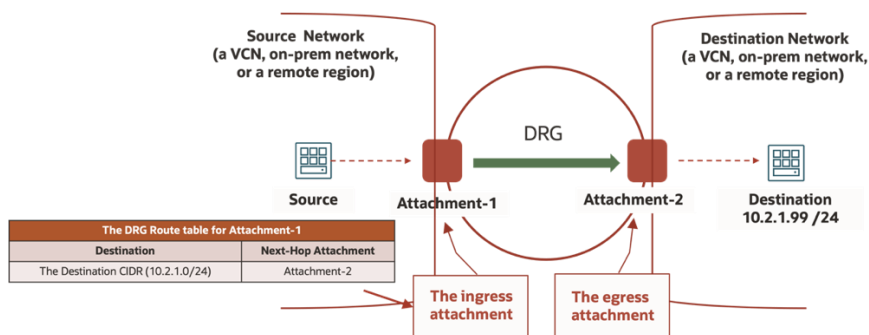


Figure 19. Example DRG Routing Operation

Because DRG attachment is a logical point-to-point connection between the DRG and the network resource behind the attachment, DRG doesn't need to perform another routing lookup on the egress attachment—it just forwards the traffic to the next network resource through the attachment. The next network resource could be a VCN, the routing device on the other side of an IPsec tunnel or a FastConnect VC, or a DRG in a remote region. This next resource must perform its own routing lookup to decide where to forward the traffic.

For example, if the next-hop attachment is a VC attachment, DRG routes the traffic through the FastConnect VC. The routing device on the other end of the VC performs its own routing upon receiving the traffic. If the next hop is a VCN attachment, the VCN ingress routing through the DRG occurs. For details about the routing operation, see the previous “VCN Ingress Routing Through a Dynamic Routing Gateway” section.

Let's use the example in Figure 20 to illustrate the routing lookup process along a multiple-hop network path.

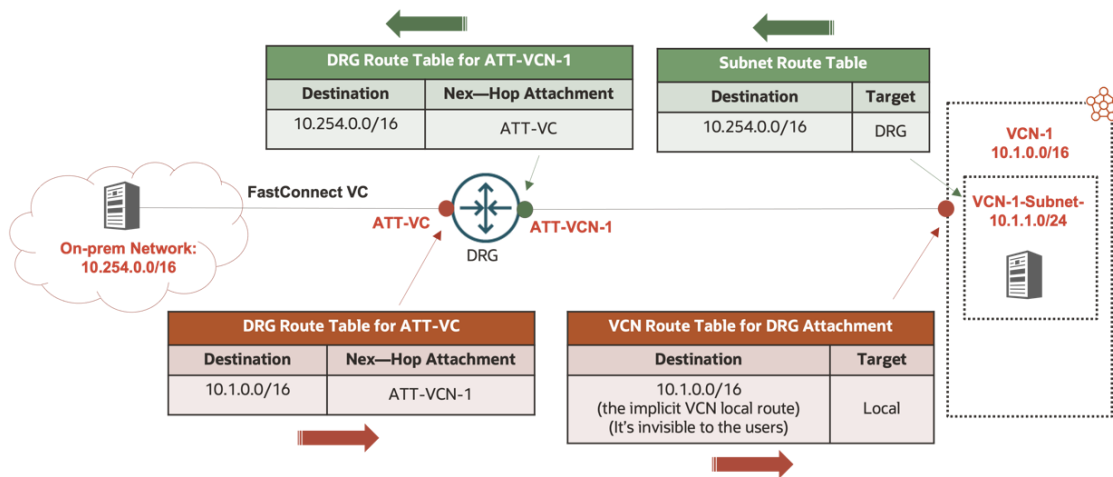


Figure 20. Routing Lookup Along a Multiple-Hop Network Path

The diagram shows the routing path between an on-premises network, 10.254.0.0/16, and a VCN subnet, 10.1.1.0/24. The default local routing in the VCN is used for simplicity. To achieve the end-to-end routing connectivity, multiple DRG route tables and VCN route tables are used at different points for the routing lookup for each direction:

- The DRG route table for the FastConnect VC attachment routes traffic from the on-premises network to VCN-1.
- The DRG route table for the VCN-1 attachment routes traffic from VCN-1 to the on-premises network.
- The VCN route table for the DRG attachment in the VCN uses VCN ingress routing through the DRG into VCN-1. If user-specified route table is associated with the DRG attachment in the VCN, the default local route for the VCN CIDRs is used, that is, the DRG routes the traffic directly to the destinations in the VCN.
- The VCN route table for the subnet 10.1.1.0/24 routes traffic from the VCN-1 subnet to the DRG.

As shown in Figure 21, the VCN subnet route table and the DRG route table for the VCN-1 attachment are used to route the traffic from a resource in the VCN-1 subnet to a resource in the on-premises network.

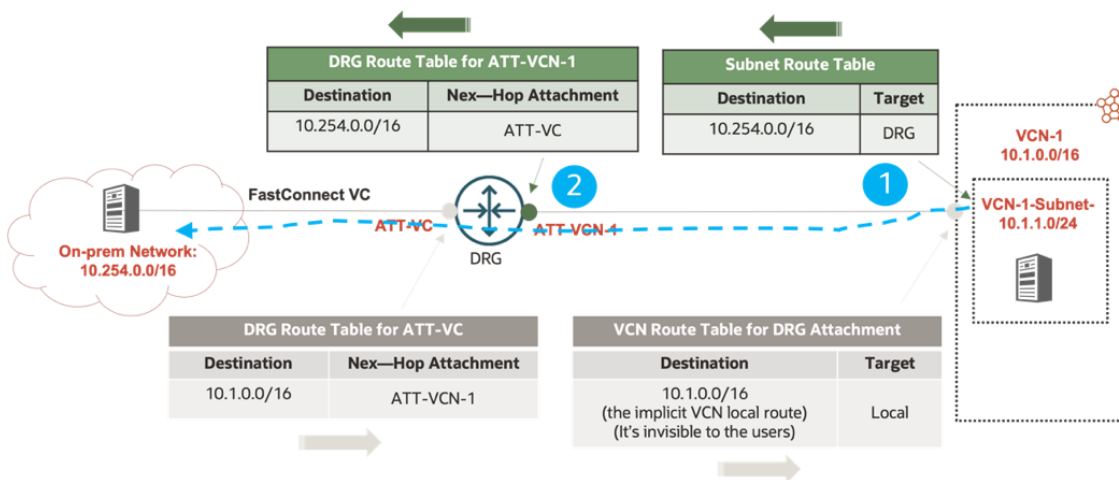


Figure 21. Routing the Traffic from VCN-1 to the On-Premises Network

As shown in Figure 22, the DRG route table for the FastConnect VC attachment and the VCN route table associated with the DRG attachment in the VCN for DRG ingress routing are used to route the traffic from the on-premises resource to a resource in the VCN-1 subnet.

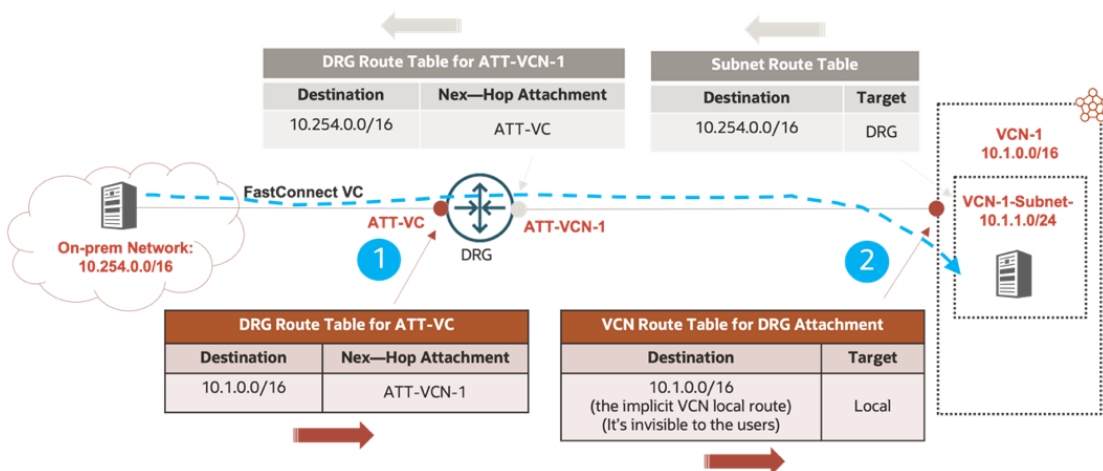


Figure 22. Routing the Traffic from the On-Premises Network to VCN-1

Intra-Regional VCN Routing

Intra-regional VCN routing refers to routing among network resources in different VCNs that are in the same region. In other words, it's inter-VCN routing within a region. Because VCNs in a region can be connected to one another by LPG peering or through a DRG, this section discusses inter-VCN routing using those the two options.

LPG peering is a traditional way of connecting VCNs. It establishes a point-to-point connectivity between two VCNs by peering their LPGs. When multiple VCNs are connected for full-mesh or partial-mesh communication, an LPG peering needs to be established between each pair of VCNs for the necessary communication. This setup presents complexity and scale challenge when the number of VCNs in the network increases.

The latest version of DRG has the capability to connect multiple VCNs. To connect multiple VCNs for full-mesh or partial-mesh communication, you just need to attach the VCNs to the same DRG and manage the route rules and route import distribution policies in the DRG route tables for the VCN attachments to achieve the necessary communication. As a result, using DRG for inter-VCN connectivity is now the recommended method for the simplicity and scalability that it provides.

Routing Between VCNs by Using LPGs

In this scenario, VCN-1 Subnet-1 and VCN-2 Subnet-2 need to send traffic to each other. Each VCN has an LPG deployed, and local peering is established between the two LPGs. The routing process is as follows:

1. In the source subnet (Subnet-1 in VCN-1), subnet egress routing occurs based on the subnet route table.
2. Traffic is routed to the local LPG (LPG-1) and then crosses the local peering to reach the remote LPG (LPG-2).
3. The remote LPG (LPG-2) routes the traffic directly to the destination in the remote VCN (Subnet-2 in VCN-2) by using the implicit local route for the VCN CIDR.

Figure 23 depicts this routing lookup process.

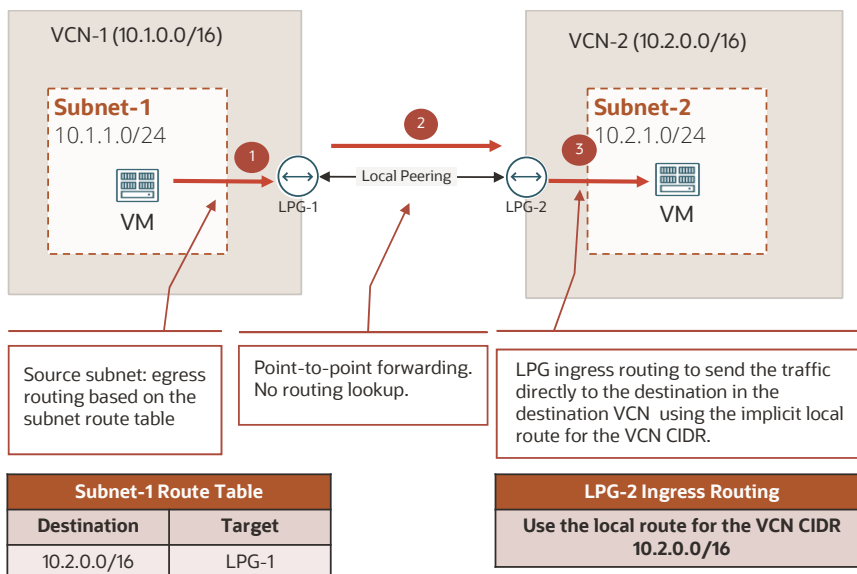


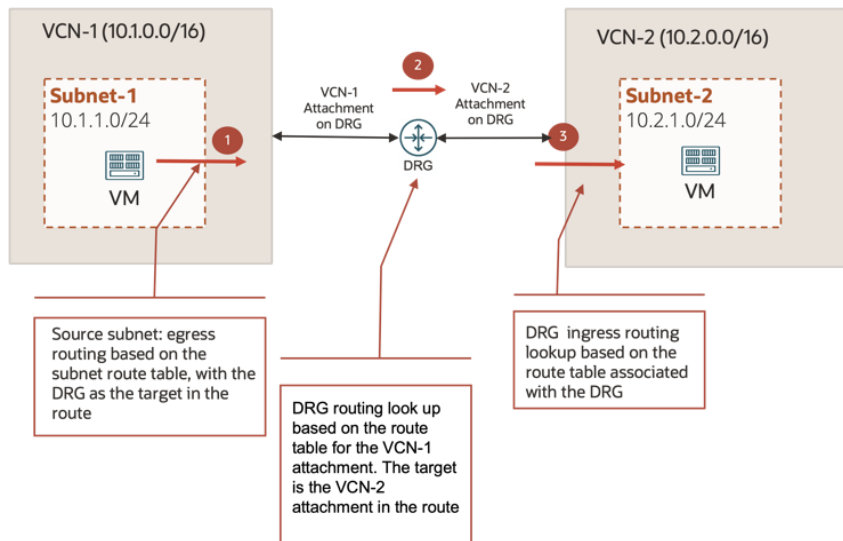
Figure 23. Inter-VCN Routing by Using LPGs

Routing Between VCNs by Using DRGs

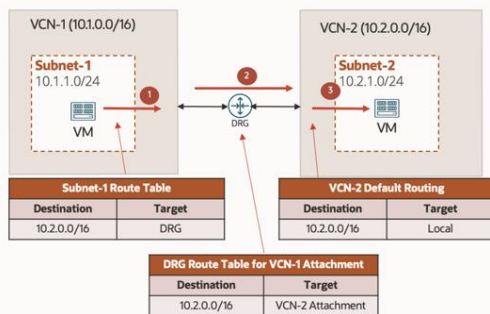
This section describes the scenario in which two VCNs (VCN-1 and VCN-2) are attached to the same DRG and need to send traffic to each other. In this scenario, the DRG routes the traffic between them. The routing process between the source and the destination in the two VCNs includes the following steps:

1. In the source subnet (Subnet-1 in VCN-1), subnet egress routing occurs based on the subnet route table, which resolves the route to the destination with the DRG as the target. The traffic is routed to the DRG.
2. The DRG performs a routing lookup in the DRG route table associated with the source VCN (VCN-1) attachment. It resolves the route to the destination with the destination VCN (VCN-2) attachment as the next-hop attachment. The traffic is routed onto the destination VCN (VCN-2).
3. The VCN ingress routing occurs through the DRG into the destination VCN. It uses the VCN route table associated with the DRG. If no VCN route table is associated with the DRG attachment, the implicit local route for the VCN CIDR is used to send the traffic directly to the destination (Subnet-2 in VCN-2).

Figure 24 depicts this routing lookup process:



Example 1: No route table is associated with the DRG. DRG uses the VCN default routing to route traffic directly to the destination.



Example 2: A user-defined route table is associated with the DRG. DRG performs ingress routing based on this route table. The target can be an NLB, LBaaS, private IP, etc. In this example, the DRG routes the traffic to a firewall.

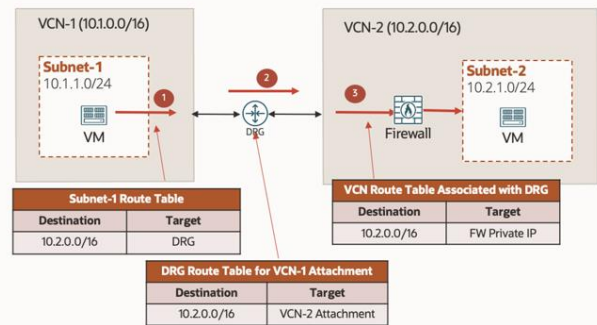


Figure 24. Routing Between Two VCNs by Using a DRG

Inter-Regional Routing

Customers often need to send traffic between resources in VCNs that are in different OCI regions. This task is done by using remote peering connection (RPC) between the DRGs in different regions. The traffic over RPC traverses the OCI backbone network for secure and high-performance data transfer.

Let's use the example depicted in Figure 25 to study the inter-regional routing process. In this example, resources in Subnet-1 of VCN-1 in Region-1 need to communicate with resources in Subnet-2 of VCN-2 in Region-2. An RPC connection is established between DRG-1 in Region-1 and DRG-2 in Region-2 to transfer the traffic between the two regions.

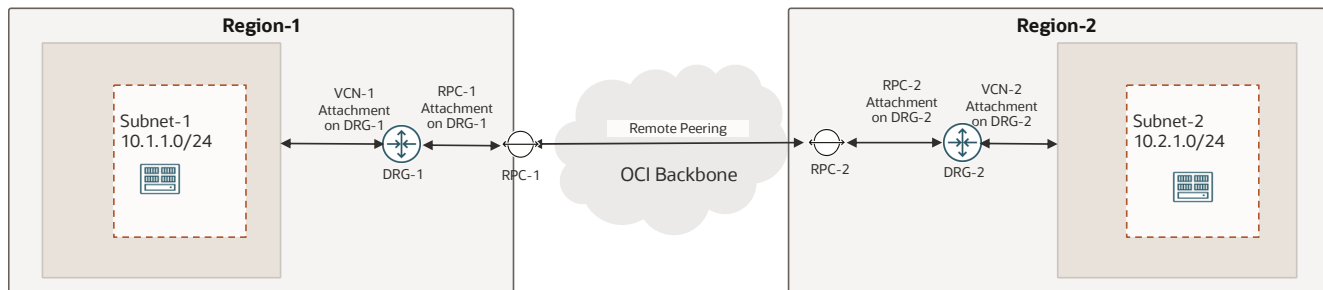


Figure 25. Example of Inter-Regional Routing

The routing process between the source and the destination in the two subnets of the two regions includes the following steps:

1. In the source subnet (Subnet-1 on VCN-1), subnet egress routing occurs based on the subnet route table, which resolves the route to the destination with the local DRG as the target. The traffic is routed to the local DRG.
2. The local DRG performs a routing lookup operation in the DRG route table associated with the source VCN attachment. It resolves the route to the destination with the RPC-1 attachment as the next-hop attachment.
3. The traffic is routed to the DRG in the remote region over the RPC connection.
4. The remote DRG performs a routing lookup operation in the DRG route table associated with the RPC-2 attachment. It resolves the route to the destination with the destination VCN attachment as the next-hop attachment. The remote DRG routes the traffic to the destination VCN (VCN-2).
5. The traffic is routed to the destination (Subnet-2 on VCN-2) through the destination VCN ingress routing. If a VCN route table is associated with the DRG, it uses this route table for the ingress routing lookup. Otherwise, the implicit local route for the VCN CIDR is used to send the traffic directly to the destination.

Figures 26 and 27 on the following page depict the routing lookup process in each direction, respectively.

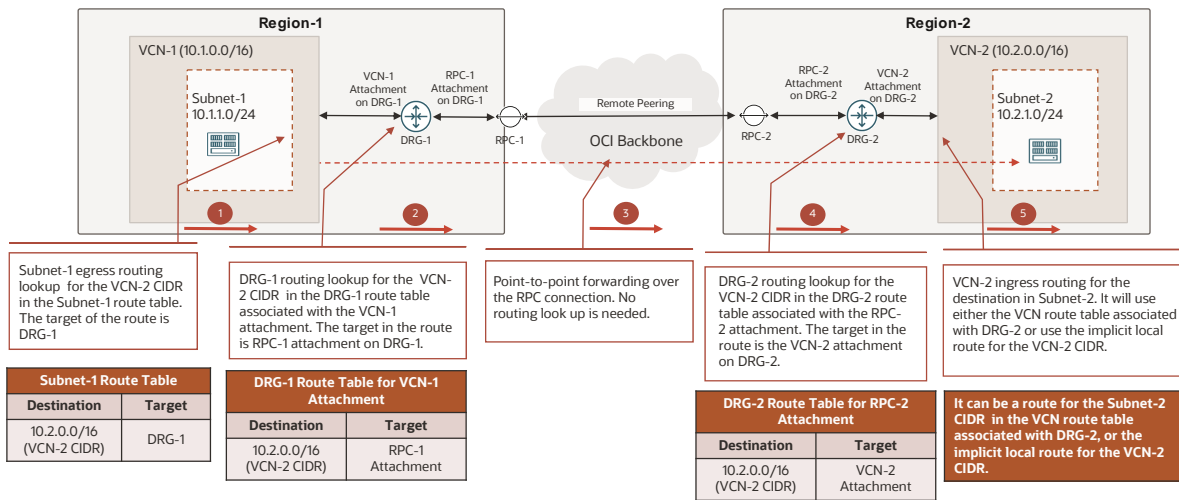


Figure 26. Routing Lookup Process from VCN-1 in Region-1 to VCN-2 in Region-2

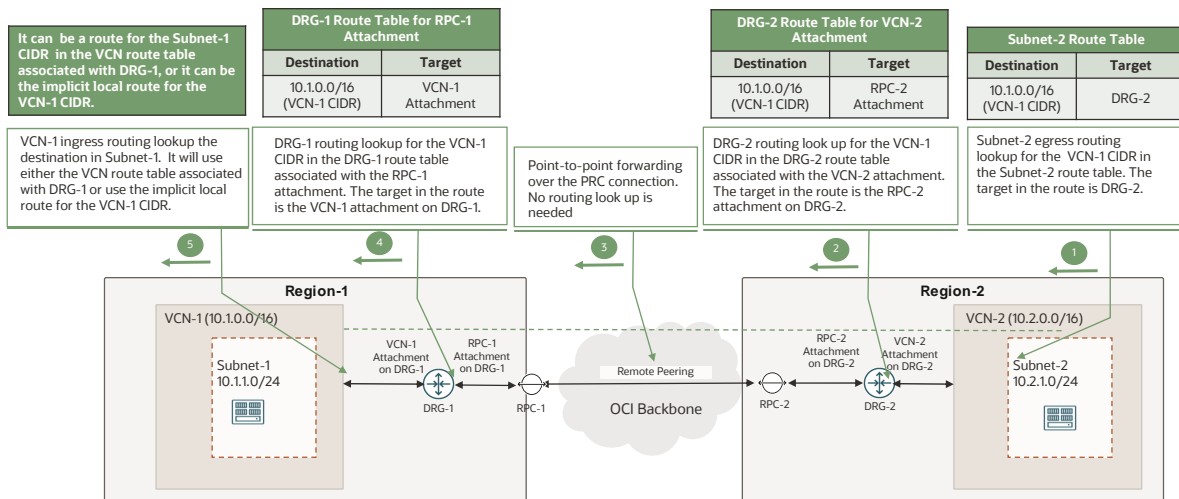


Figure 27. Routing Lookup Process from VCN-2 in Region-2 to VCN-1 in Region-1

On-Premises Site to VCN Routing Examples

Hybrid cloud deployments are common. In a hybrid network, hosts in the on-premises networks and hosts in the cloud networks need to communicate with each other.

With enhanced DRG, customers can connect multiple VCNs with their on-premises networks through the same DRG. The on-premises networks can be attached to the DRG through either VPN tunnels or FastConnect virtual circuits (VCs). If BGP is running between the DRG and the on-premises CPEs, the VCN CIDR and subnet CIDR routes can be advertised to the CPEs by the DRG, and the CPEs can advertise the on-premises network routes to the DRG as well.

As depicted in Figure 28, the traffic coming from the on-premises networks to destinations in a OCI VCN go through the following routing process:

1. The on-premises network routes the traffic to the DRG through the VPN tunnel or the FastConnect VC.
2. The DRG performs a routing lookup operation in the DRG route table associated with the VPN or FastConnect attachment for the on-premises networks. It resolves a route to the destination in the VCN with the VCN attachment as the next hop.

- The traffic gets to the VCN and is subject to the VCN ingress routing based on the VCN route table associated with the DRG attachment. If no VCN route table is associated with the DRG attachment, the VCN uses the implicit local route for the VCN CIDR to route the traffic to the destination directly.

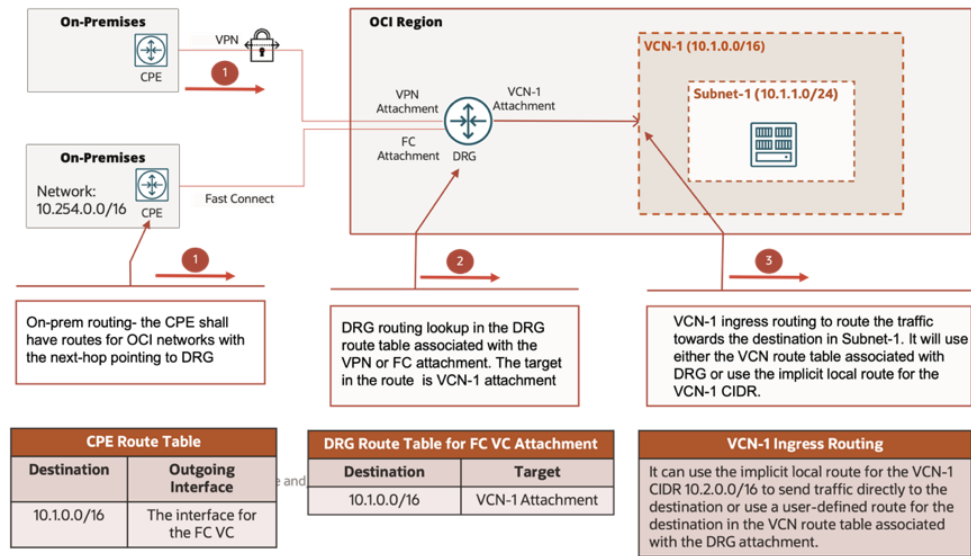


Figure 28. Routing Lookup Process from On-Premises Network to VCN

As depicted in Figure 29, traffic in the reverse direction goes through the following routing process:

- The egress routing of the VCN source subnet uses the VCN route table associated with the subnet. It resolves a route to the destination on-premises network with the DRG as the next hop. As a result of this subnet egress routing, the traffic is routed out of the VCN into the DRG.
- The DRG performs a routing lookup operation in the DRG route table associated with the VCN attachment. It resolves a route to the destination on-premises network with the VPN or FastConnect attachment as the target. The traffic is routed towards the CPE through the VPN or the FastConnect attachment.
- The CPE performs the on-premises routing lookup operation. It resolves a route to the on-premises network with the next on-premises routing device as the next hop.

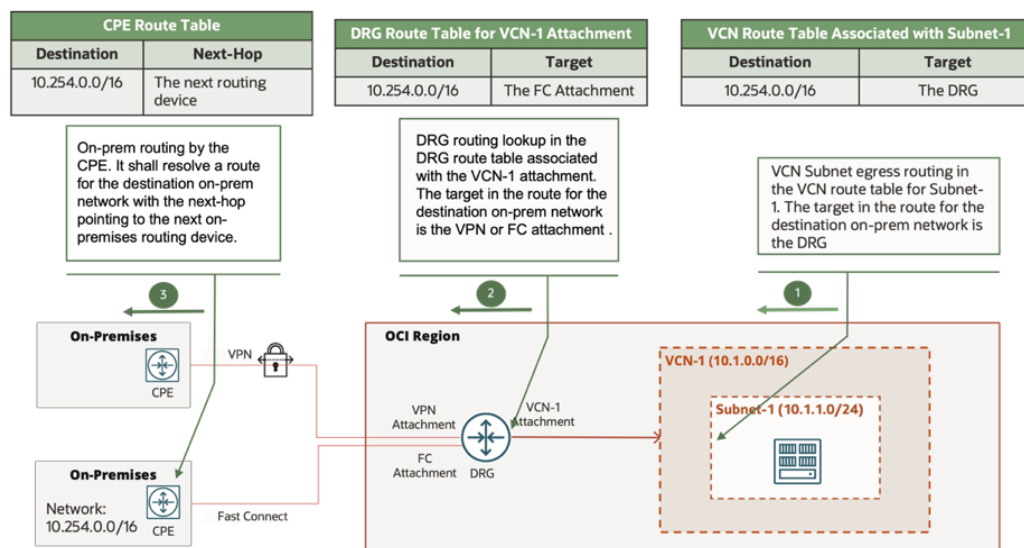


Figure 29. Routing Lookup Process from VCN to On-Premises Network

With the latest capability of DRG, we recommend using DRG as the central hub to interconnect your VCNs and to connect them to the on-premises networks directly. The diagram in Figure 30 shows such a network design.

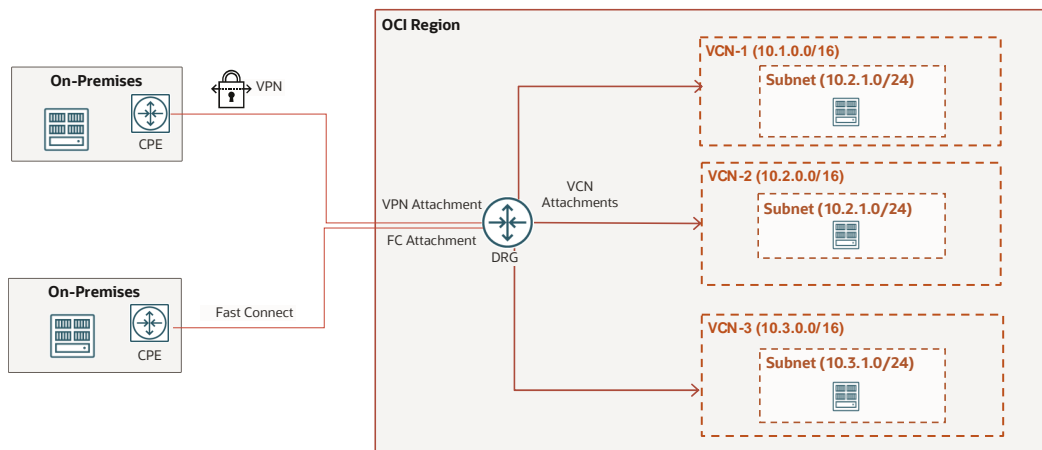


Figure 30. On-Premises Network and VCN Connectivity with DRG as the Hub

However, some customers are still using the legacy DRG to connect their on-premises network to their OCI environment. The legacy DRG supports only a single VCN attachment. To achieve the network connectivity between the on-premises networks and multiple VCNs, we recommend using the transit routing design depicted in Figure 31 as a solution in the case of the legacy DRG. This design is a typical hub-spoke topology in which the hub VCN is connected to the on-premises networks through a legacy DRG, and the spoke VCNs are connected to the hub VCN through point-to-point LPG peering.

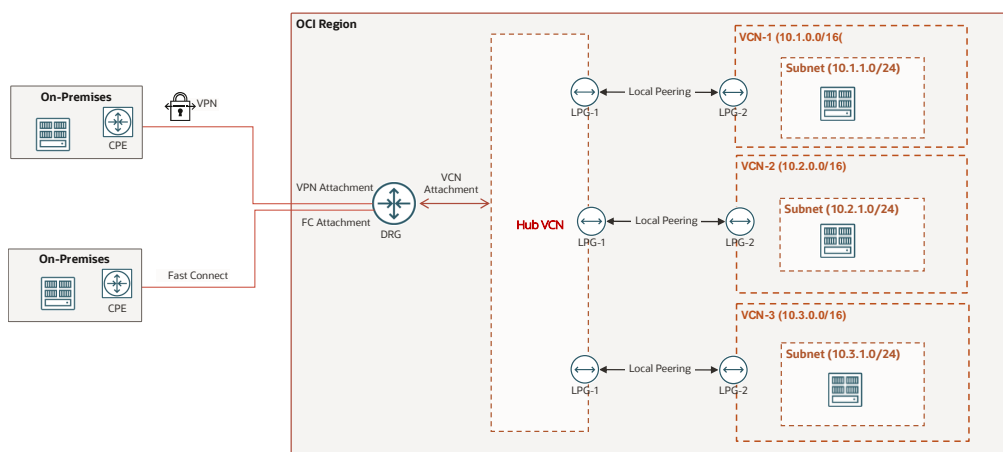


Figure 31. On-Premises and VCN Connectivity with a Legacy DRG and Hub-Spoke VCN Design

Remote On-Ramp Routing Example

The remote on-ramp solution allows your on-premises network to access resources in a remote OCI region by connecting to its local OCI region and traversing the local region through the OCI backbone to reach the remote region. Figure 32 depicts a typical network topology for a remote on-ramp design.

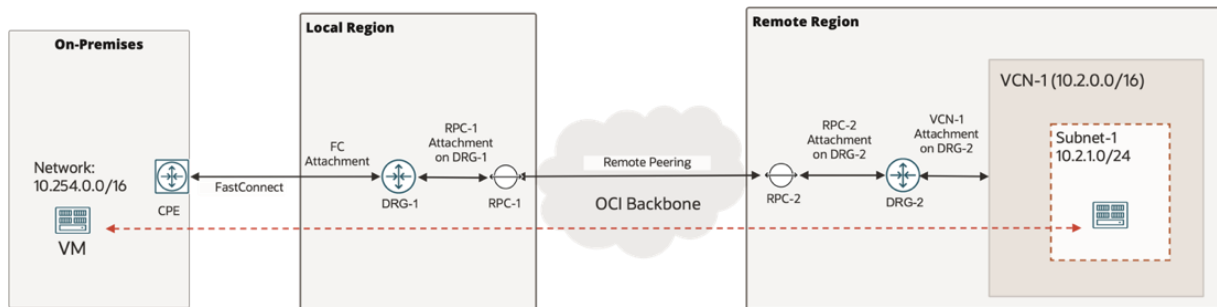


Figure 32. A Typical Remote On-Ramp Design

In this design, the on-premises network needs to access some OCI resources in VCN-1 in the remote region. The on-premises network is connected to the local OCI region through FastConnect virtual circuits (VCs). The local region DRG (DRG-1) has a remote peering connection (RPC) established with the DRG in the remote region (DRG-2). DRG-2 then connects to VCN-1 where the OCI resources reside.

From routing perspective, the key things necessary for this design to work for bi-directional communication include the following items:

1. The route for the on-premises network needs to be propagated to the remote DRG (DRG-2) by the local DRG (DRG-1) through the RPC connection.
2. The remote region VCN-1 or VCN-1 Subnet-1 CIDR routes need to be propagated to the local DRG (DRG-1) by the remote DRG (DRG-2) through the RPC connection, and then get imported into the DRG route table of the FastConnect VC attachment.
3. The local DRG (DRG-1) needs to advertise the routes for the remote region VCN-1 or VCN-1 Subnet-1 to the CPE.

In order for the on-premises network route to be propagated to the remote region through the RPC connection, the DRG-1 in the local region needs to have the route for the on-premises network in its route table associated with its RPC attachment (RPC-1) because the routes in this DRG route table are propagated on DRG-2 in the remote region. This requires DRG-1 to import this on-premises route into the DRG route table for the RPC attachment. To achieve it, you need to have an import distribution policy in the DRG route table of the RPC attachment to match the FastConnect attachment.

Similarly, in order for the VCN-1 or VCN-1 Subnet-1 routes in the remote region to be propagated to the CPE, the routes need to be imported into the local DRG-1 route table for the FastConnect VC attachment. This DRG route table needs to have an import distribution policy that matches the RPC attachment.

If you're using the automatically generated DRG route table for RPC, VC, and IPsec attachments for your FastConnect or RPC attachments on the DRGs, you need to remember that this route table by default imports only the routes propagated by VCN attachments. It doesn't automatically import the routes propagated by the RPC, FastConnect, or IPsec attachments. You need to add the necessary import policies yourself. To avoid mistakes and to allow for more policy flexibility, we recommend that you create separate route tables for these attachments.

Figure 33 depicts the routing process for the traffic from the on-premises network to the VCN subnet in the remote region. Using a traffic flow from a host in the on-premises network (10.254.0.0/16) to a network resource in VCN-1 (10.2.0.0/16) Subnet-1 (10.2.1.0/24) as an example, it shows the needed routes in the relevant route tables along the path, including the CPE route table, the DRG route tables for the FastConnect VC attachment on DRG-1 and the RPC attachment on DRG-2, and the VCN-1 route table for VCN ingress routing through DRG-2.

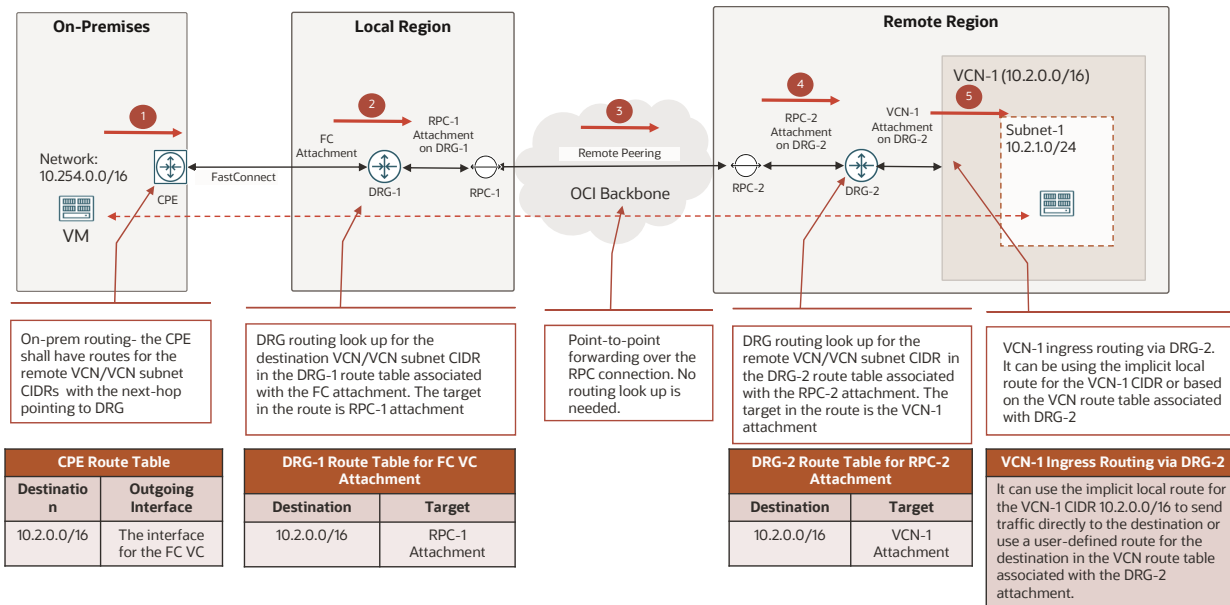


Figure 33. Routing Lookup Process from the On-Premises Network to the Remote Region

Figure 34 shows the routing lookup process for the traffic in the reverse direction, that is, the traffic from the VCN subnet in the remote region to the on-premises network. Using a traffic flow from the network resource in VCN-1 (10.2.0.0/16) Subnet-1 (10.2.1.0/24) to a host in the on-premises network (10.254.0.0/16) as an example, it shows the needed route in the relevant route tables along the path, including the VCN subnet route table, the DRG route tables for the VCN attachment and the RPC attachments, and the route table on the CPE.

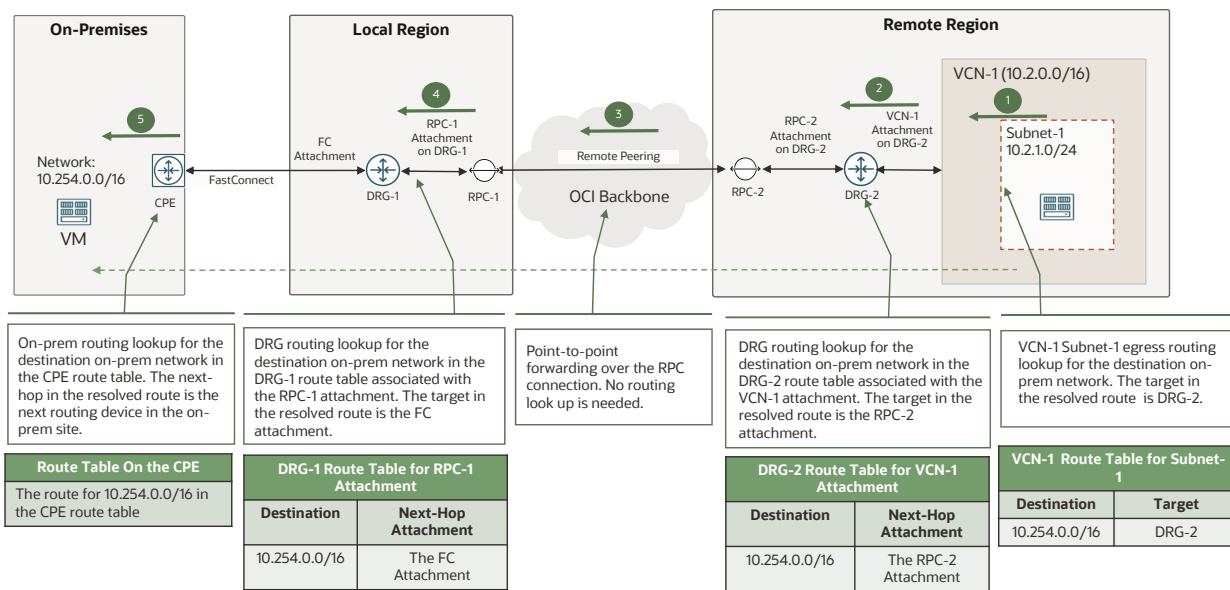


Figure 34. Routing Lookup Process from the Remote Region to the On-Premises Network

Routing for Network Virtual Appliance Insertion Examples

Network virtual appliances often need to be inserted into the traffic routing path for additional processing of the traffic before it reaches the destination. One typical example is to add a firewall for security inspection. There are different network designs for service insertion:

- **Per-VCN network virtual appliance:** Each VCN has its own network virtual appliance.
- **Centrally shared network virtual appliance:** Multiple VCNs share the same centrally deployed network virtual appliance.

The remainder of this section uses a virtual firewall as the example for routing in these two design options.

Per-VCN Network Virtual Appliances

This design uses intra-VCN routing to insert the virtual firewall between subnets in the same VCN. The diagram in Figure 35 shows an example design in which an OCI network firewall is inserted between the web-tier and the app-tier subnets of an application in the same VCN.

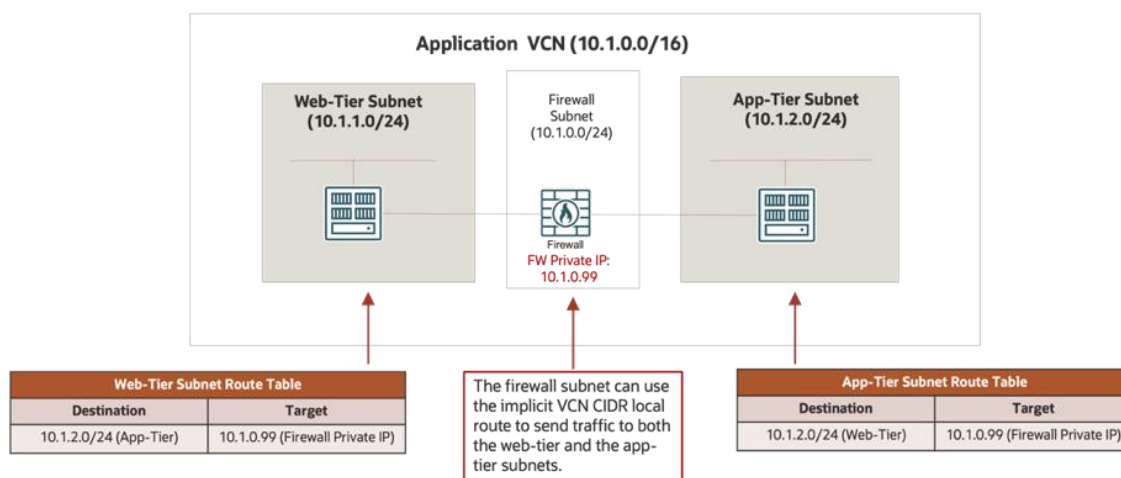


Figure 35. Per-VCN Firewall Design

The intra-VCN subnet routes in both route tables ensure that the traffic in both direction is routed to the firewall instead of reaching the destination directly.

Centrally Shared Network Virtual Appliances

In this design, a firewall is deployed in a central service-hub VCN, and applications deployed in multiple spoke VCNs share the firewall. The service-hub VCN and the application spoke VCNs must be attached to the same DRG. DRG routing and intra-VCN routing are used for the routing in this design.

The diagram in Figure 36 shows an example of such a design. In this example, multiple spoke VCNs for different applications are connected to the service-hub VCN through the attachments to the common DRG. A central firewall is deployed in the service-hub VCN that is shared by the applications. Traffic between the web tier and app tier in each application is sent to this central firewall for inspection. The dotted lines in the diagram show the logical path between the application tiers.

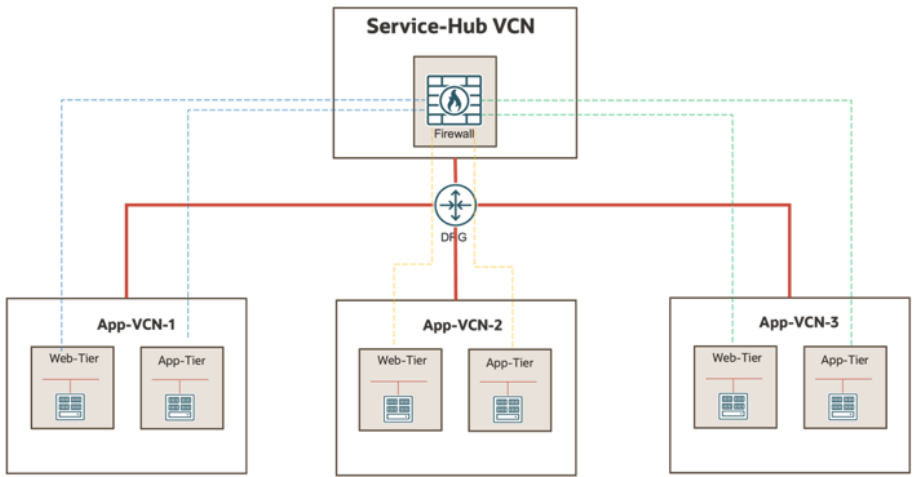


Figure 36. Central Firewall Design

To study the routing process along the traffic path between the two tiers in each application, let's focus on one application spoke VCN in which a web tier and an app tier are deployed in two subnets. Traffic between the two tiers must go through the central firewall in the service-hub VCN for inspection.

Figure 37 shows the hop-by-hop routing process for traffic from the web-tier subnet to the app-tier subnet.

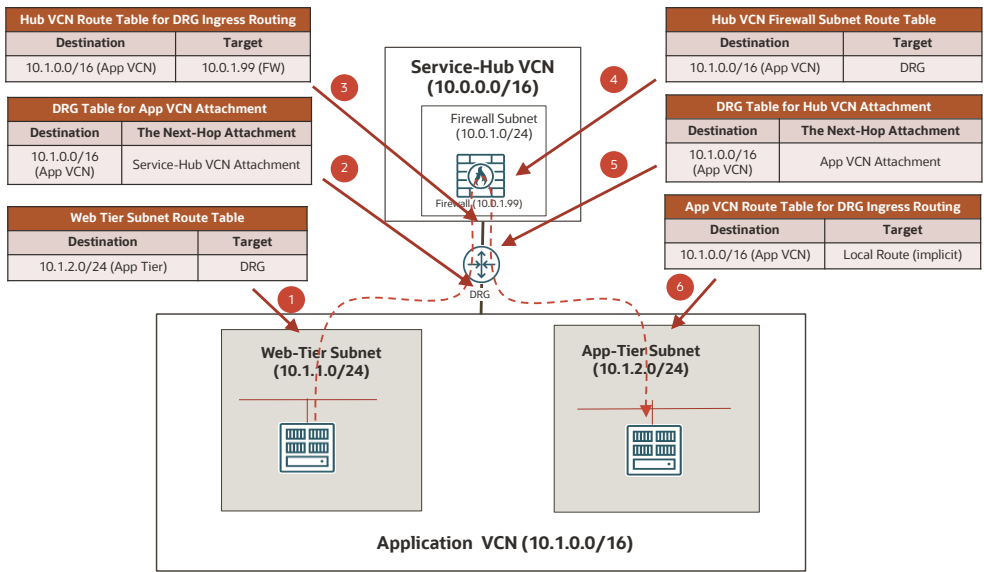


Figure 37. Routing Lookup Process from the Web Tier to the App Tier

Figure 38 shows the hop-by-hop routing process for the traffic from the app-tier to the web-tier.

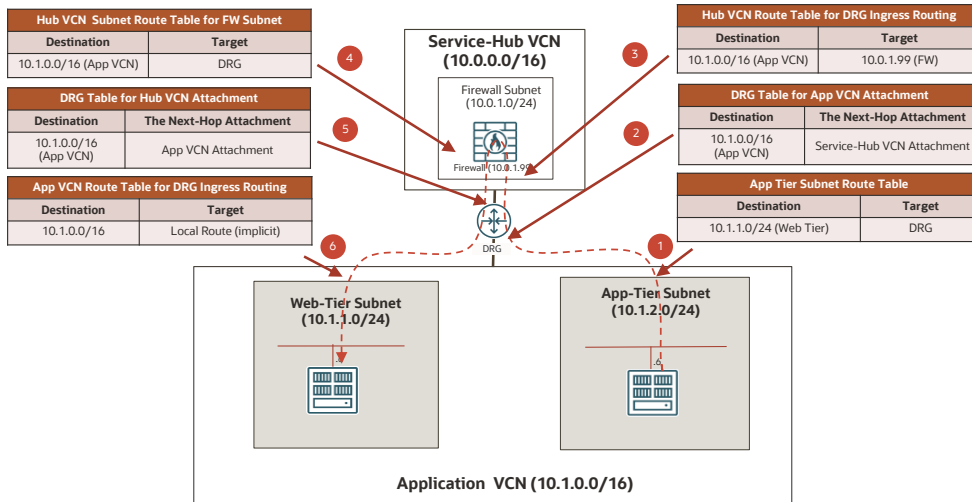


Figure 38. Routing Lookup Process from the App Tier to the Web Tier

This design with a shared firewall is often considered to be more cost-effective and simpler to manage compared to the design with separate firewalls per applications.

Private Access to OCI Services

OCI offers different ways for you to securely access OCI services from your private cloud network by traversing the OCI internal network (called the Oracle Services Network). The options include secure private access through a VCN service gateway (SGW) and secure private access through private endpoints.

Because a private endpoint is represented as an endpoint in the user's VCN and has a private IP address allocated from the VCN CIDR range, routing to and from a private endpoint in the user's networks is just like routing to a user's deployed instance in the VCN. All the routing design and operations discussed in previous sections also apply to private endpoints. Therefore, we don't separately discuss the routing for private endpoints in this section. Instead, we discuss routing in the private access design with a VCN SGW.

OCI Instances Access OCI Services Through a Service Gateway

Instances deployed in OCI VCNs can access OCI services in the same region through the private VCN network and the OCI internal network by using a VCN SGW.

When accessing OCI services from a VCN through an SGW, you need to add a route rule for the service label with the SGW as the target in the VCN route table for the subnet from which you're trying to access the service. Currently, OCI supports two service labels in each region: one for the Object Storage service, and one for all services. Figure 39 shows an example route table that has a user-defined route rule to the OCI Object Storage service in the US East (Ashburn) region.

Destination	Target Type	Target	Route Type	Description
0.0.0.0/0	Internet Gateway	igw-us-east-vcn-0	Static	
10.0.0.0/16	Private IP	10.0.1.3	Static	
OCI:AD.ObjectStorage	Service Gateway	sgw-lad-vcn-0	Static	

Figure 39. Example of Route Rules to OCI Services in a VCN Route Table

The diagram in Figure 40 shows a simple example in which instances in the App subnet of the Application VCN access the regional Object Storage service through the SGW in the VCN.

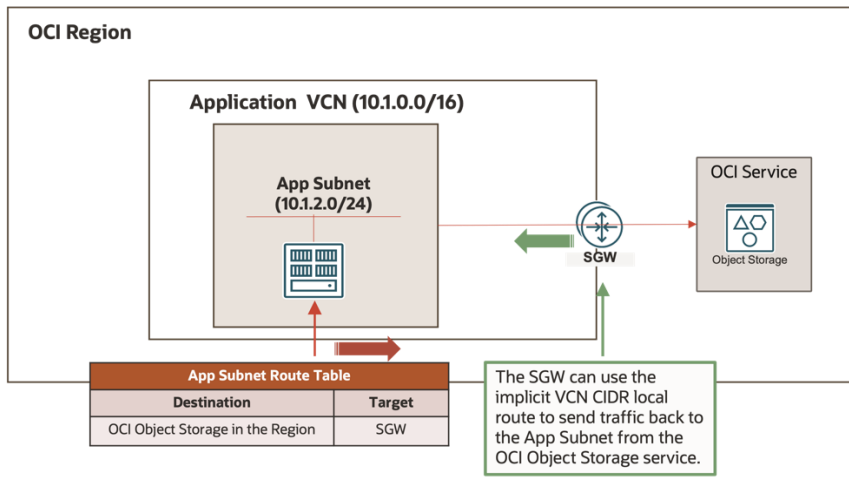


Figure 40. OCI Instance Accesses OCI Services Through an SGW

For traffic from instances inside the App subnet to the Object Storage service, the App subnet route table is used for routing. It has a route to the Object Storage label with the SGW as the target. It routes the traffic to the SGW, and the SGW performs the rest of the forwarding and sends the traffic to the service.

For traffic from the service back to the instances in the App subnet, OCI forwards the traffic to the SGW. In this simple example, the SGW uses the default implicit VCN CIDR route to route the traffic back to the instance directly.

User often want to insert a firewall in front of the SWG to filter on Layer-7 URLs to limit what services are accessible, and to inspect the traffic based on their security policies. The firewall insertion along the forwarding path through the SGW can be done by using SGW ingress routing. Figure 41 shows such a design for the same App subnet to access the Object Storage service through the SGW with an OCI network firewall inspecting traffic from the service.

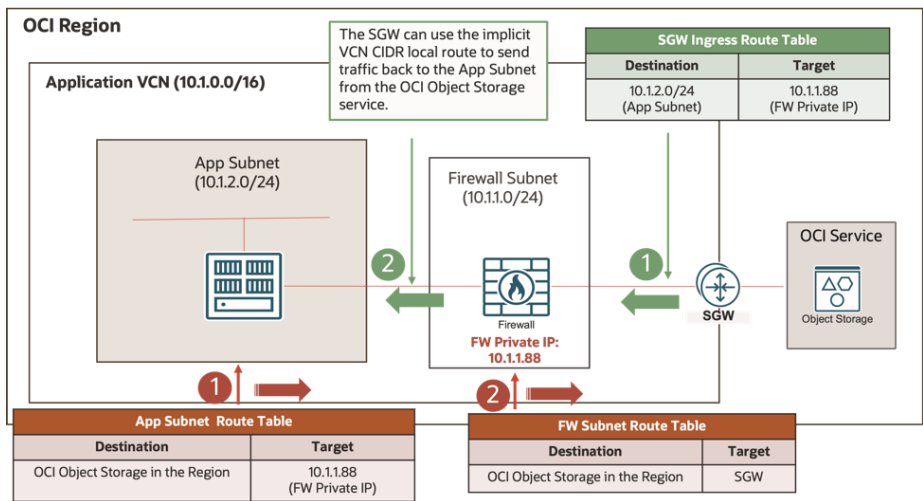


Figure 41. OCI Instances Access OCI Services Through a Firewall and SGW

This design requires the following OCI routing configuration:

- The App subnet route table now has a route rule for the OCI Object Storage label with the OCI network firewall private IP address as the target. This route rule ensures that the traffic from the instances is routed to the firewall.

- The firewall subnet route table has a route rule for the OCI service with the SGW as the target. This rule further routes the service-bound traffic to the SGW after it goes through the firewall.
- The SGW is associated with a route table that has a route rule for the App subnet CIDR with the firewall private IP address as the target. This rule routes the ingress traffic from the service to the VCN onto the firewall.
- The firewall subnet can use the implicit VCN CIDR local route to route traffic from the service directly to the destinations in the App subnet after it goes through the firewall inspection.

The traffic from instances in the App subnet to the OCI service goes through the following routing process:

1. The App subnet routes the traffic to the firewall by using the route for the OCI service in the subnet route table. The firewall private IP address is the target of the route.
2. The firewall processes the traffic and then routes it to the firewall subnet gateway. The firewall subnet routes the traffic to the SGW by using the route rule for the OCI service in the subnet route table. The SGW is the target of the route rule.
3. After the traffic is sent to the SGW, OCI forwards it to the service.

The traffic from the OCI service to instances in the App subnet goes through the following routing process:

1. OCI performs the routing from the service to the SGW.
2. The SGW routes the traffic onto the firewall in the VCN by using the route rule for the App subnet in its ingress routing table. The firewall private IP address is the target of the route rule.
3. The firewall processes the traffic and then routes it to the firewall subnet gateway. The firewall subnet routes the traffic to the destination instances in the App subnet by using the implicit VCN CIDR local route.

On-Premises Instances Access OCI Services Through a Service Gateway

After it's connected to OCI private cloud networks through FastConnect virtual circuits or IPsec VPN tunnels, an on-premises network can access OCI services through the private OCI network as well. Typically, this access is achieved by means of a transit routing design in which traffic between the on-premises network and OCI services is routed through a transit-hub VCN. The on-premises network and the transit-hub VCN are attached to the same DRG, and the transit-hub VCN has an SGW for private access to the OCI services.

The diagram in Figure 42 shows an example in which the on-premises network 10.254.0.0/16 is connected to a DRG and then uses the SGW in the transit-hub VCN for private access to the OCI Object Storage service.

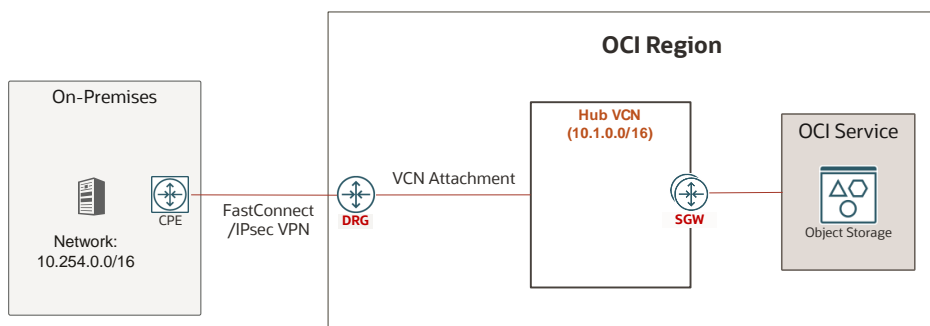


Figure 42. Typical Transit Routing Design for On-Premises Networks to Access OCI Services Through an SGW

Figure 43 shows the required route advertisement or configuration and the routing lookup process for the traffic from the on-premises network to OCI services.

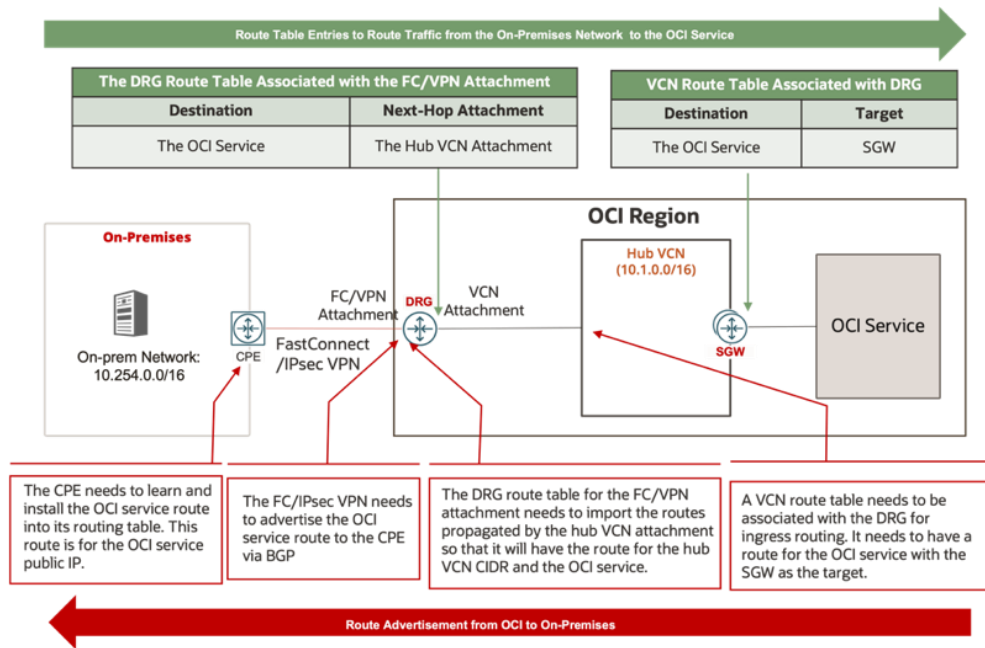


Figure 43. Route Advertisement and Routing Process for Traffic from On-Premises Networks to OCI Services

Figure 44 shows the required route advertisement or configuration and the routing lookup process for the traffic from OCI services to the on-premises network.

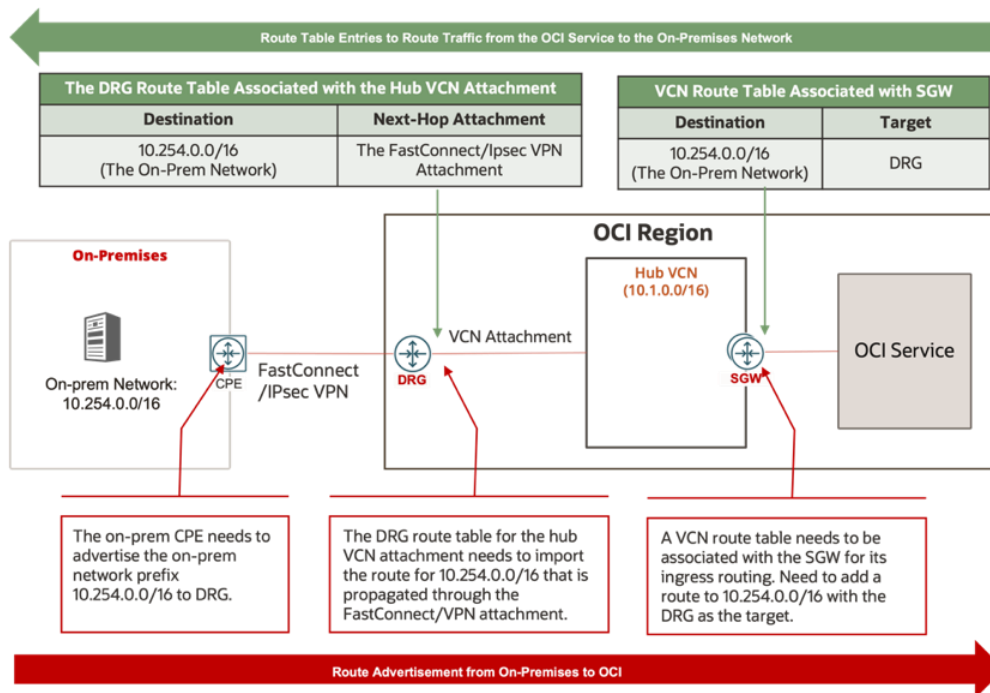


Figure 44. Route Advertisement and Routing Process for Traffic from OCI Services to the On-Premises Network

Conclusion

Oracle Cloud Infrastructure (OCI) offers feature-rich routing capabilities in its virtual network solution. This technical paper provides a comprehensive knowledge base of OCI virtual network routing with a good set of real-world network design examples. It's a good reference for anyone who wants to learn about routing in OCI virtual networks or needs to design, operate, or troubleshoot OCI virtual networks.

To learn more about OCI networking, review the following resources:

- [OCI Networking service documentation](#)
- Oracle and OCI blog posts:
 - [OCI Networking Best Practices - Part One - OCI Network Design, VCN, and Subnets](#)
 - [OCI Networking Best Practices - Part Two - OCI Network Security](#)
 - [OCI Networking Best Practices - Part Three - OCI Network Connectivity](#)
 - [Introducing global connectivity and enhanced cloud networking with the dynamic routing gateway](#)
 - [Announcing OCI intra-VCN routing and VCN gateway ingress routing enhancements](#)
 - [Defense in Depth, Layering using OCI Network Firewall](#)

Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2023, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0120