

Connectivity Redundancy Guide

ORACLE WHITE PAPER | SEPTEMBER 2019





Disclaimer

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, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Revision History

The following revisions have been made to this white paper since its initial publication:

Date	Revision
September 5, 2019	Updated Figure 11.
May 30, 2019	Initial publication.



Table of Contents

Overview	4
Design Considerations	4
Redundancy Use Cases	5
Oracle VPN Connect	5
VPN Connect with a Redundant Customer Edge Device	7
VPN Connect Plus FastConnect	9
Redundant FastConnect	11
References	14

Overview

Enterprise customers experience growth in their cloud deployments, and more critical applications are increasingly deployed to the cloud. With this growth, you need to ensure that your cloud infrastructure is available and connected to your on-premises network in a redundant way so that it can support planned maintenance outages and unplanned downtime.

The purpose of this document is to help you check your current deployment for redundancy or upgrade a single connection to Oracle Cloud Infrastructure to a redundant connection. It reviews some use cases and options for connectivity through FastConnect and through VPN Connect (IPSec VPN) over the internet. It assumes that you are familiar with routing protocols and concepts, IPSec VPN technology and configuration, and Oracle Cloud Infrastructure concepts and components.

Design Considerations

When you deploy resources to Oracle Cloud Infrastructure, you might start small, with a single connection to your on-premises network. This single connection could be through FastConnect or through VPN Connect, which is the fastest way to deploy a connection to Oracle Cloud Infrastructure.

To plan for redundancy, consider all the components (hardware, facilities, circuits, and power) between your on-premises network and Oracle Cloud Infrastructure. Also consider diversity, to ensure that facilities are not shared between the paths.

Table 1 shows the components that you need to consider for a redundant solution.

TABLE 1. DESIGN CONSIDERATIONS

Components	Comments
Internet service provider (ISP)	Not all ISPs are the same. Peering relationships from your ISP let your traffic route more efficiently, reducing the latency as it varies over the internet.
Hardware	Enable services with redundant hardware, and ensure that there's no single point of failure anywhere in the path. How will you handle infrastructure maintenance (by Oracle or your own IT department)? Can you tolerate downtime? How much downtime can you tolerate?
Facilities diversity	Do you have redundant power feeds? Do you have diverse telecommunication entry points into your building? Is your equipment in different racks or data centers?



Components	Comments
Oracle FastConnect POP diversity	Do you want to terminate both FastConnect circuits into the same point of presence (POP) or into different locations? Note that POP diversity is available only in the Phoenix, Ashburn, Frankfurt, and London regions.
Circuit provider diversity	Are you planning to use diverse carriers? Are your WAN or internet circuits fully diverse, or do they share a POP? Note that having different carriers doesn't mean that the circuits are fully diverse.

Redundancy Use Cases

This section describes how to create a redundant connection to Oracle Cloud, starting from a single FastConnect or VPN Connect connection, and walking through the following use cases:

- VPN Connect with a redundant customer edge device
- VPN Connect plus FastConnect
- Redundant FastConnect

The use cases provide instructions on how to set up primary and backup connections within the same region only.

Oracle VPN Connect

Oracle VPN Connect is the quickest way to privately connect your on-premises network to Oracle Cloud by using the internet as the transport and encryption to secure your traffic from the internet. When you create a VPN Connect connection, Oracle provides the public IP addresses for two tunnel endpoints (headends) in the same region for redundancy. In this document, the VPN headends are represented as independent components to make it easier to understand the concepts and the endpoints for the tunnels. However, in general, the connection is to the dynamic routing gateway (DRG).

This solution consists of a single edge device in your on-premises network and two VPN headends (default) in a single Oracle Cloud region. Your edge device can be in your company's headquarters, a data center, a colocation facility, or even in another cloud.

Figure 1 illustrates the default setup for a single VPN Connect connection.




Figure 1. VPN Connect for a Single Region with a Single Customer Edge Device

Figure 2 outlines the routing necessary for this solution to work. Security lists must also be updated accordingly to allow the traffic. Security lists are out of scope for this document because their contents depend on the type of applications and traffic that you want to allow through this connection.

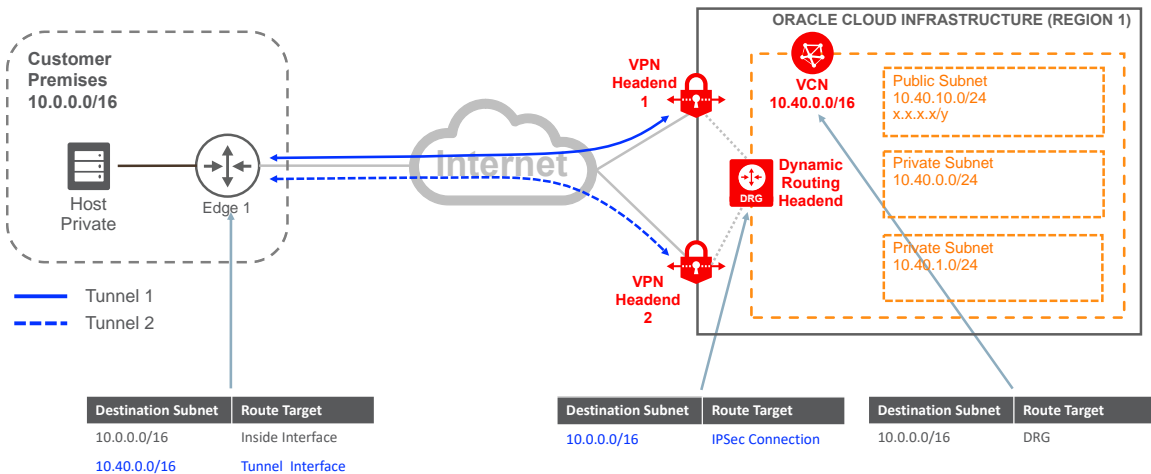


Figure 2. Routing for VPN Connect for a Single Region with a Single Customer Edge Device

As shown in Figure 2, redundancy is provided only on the Oracle end. The customer edge device is the single point of failure in this solution. To overcome this issue, you have two options, as described in the next two use cases.

VPN Connect with a Redundant Customer Edge Device

The VPN Connect solution depicted in Figures 1 and 2 has a single point of failure: the customer's edge device. To overcome this issue, deploy a second edge device, in the same location as the primary device, in a different data center, or in another cloud. If the second device is in the same location as the primary one, verify that they connect to different internet providers, LAN switches, and power units. Ensure that your edge devices don't share a common point of failure.

For simplicity, Figure 3 shows the two customer edge devices deployed at the same location with two carriers connecting to the internet. As stated earlier, VPN Connect automatically provides two termination VPN headends per connection. Oracle headends have diverse connections to the internet and are located in diverse data centers within the same region. As shown in Figure 3, each edge device has two tunnels, represented by the blue (Tunnel) and red (Backup Tunnel) lines.




Figure 3. VPN Connect for a Single Region with Redundant Customer Edge Devices

Four tunnels provide redundancy, but they can complicate your routing deployment. Now you must configure your routing over four tunnels and prioritize routes for each tunnel. In Figure 3, on the customer end, traffic fails back from Edge 1 to Edge 2 only if Edge 1 fails or if its internet circuit is down.

Note: Oracle has several diverse, redundant headends per region. Figure 3 shows only two of them.

If redundancy is maintained, you can choose *not* to create the second tunnel per connection. You could simplify the solution by reducing the number of tunnels from four to two. Oracle still provides connectivity information for a second headend if you want to configure one in the future. This design still provides redundancy and diversity because each edge device builds a tunnel to a different Oracle VPN headend, as shown in Figure 4. This simplified version provides an active/passive solution that you can control more effectively through routing.




Figure 4. VPN Connect for a Single Region with Redundant Customer Edge Devices Simplified

With fully diverse paths in the solution, the next step is to ensure that routing is set correctly on both sides of the connection to define the primary and backup paths. In Figure 5, the blue tunnel (Tunnel 1) is the primary path, and the red tunnel (Backup Tunnel 2) is the backup. To influence the routing, we recommend advertising more-specific routes over the primary path and less-specific routes over the backup path. With this approach, traffic is symmetric in both directions. If the primary path fails, a less specific route is available through the backup path. When the primary path is restored, traffic fails back to the primary path because it advertises a more specific route.

Because this use case uses static routing, it's important to set the routing to withdraw the route from the route table when the path is not available. Otherwise, traffic won't fail to the backup path. VPN Connect also supports Border Gateway Protocol (BGP), so you can manipulate the routes accordingly.

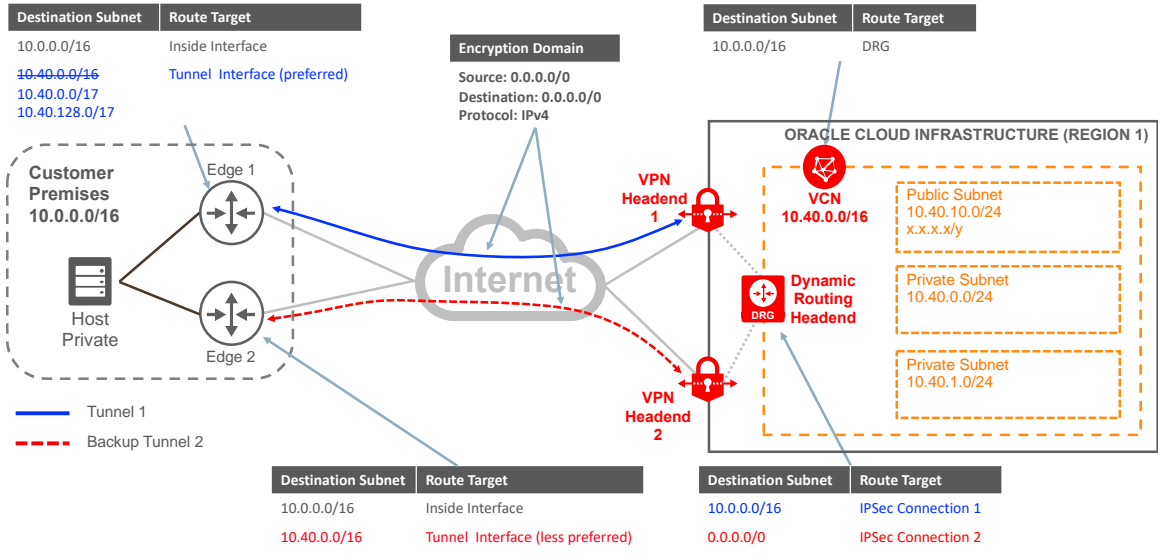


Figure 5. Routing for VPN Connect for a Single Region with Redundant Customer Edge Devices

Figure 5 shows the routing for each of the components, and the color assigned to the route highlights which path it belongs to. In your on-premises network, influence the primary (blue) path based on your internal routing protocol because both edge devices advertise the same VCN subnet. From the Oracle end, advertise your on-premises network for the primary (blue) path and advertise the default route or a less-specific route through the backup (red) path. In Figure 5, strikethrough text refers to the changes that you need to perform on your existing configuration. Figure 5 reflects the advertisement of more specific subnets from your on-premises network over the primary path. You can achieve the same results by manipulating the metrics of your interior routing protocol to prefer the primary path.

Routing is important for the solution to work, and it's independent of the *encryption domain* configuration in the tunnels. With routing, you can decide what traffic is sent to the tunnel interface, whereas the encryption domain determines what traffic is encrypted. In Figure 5, the encryption domain (middle of the diagram) is the same for both tunnels on both sides, but routing is handled at each end of the connection to a primary path and backup path for redundancy. This solution maintains a single encryption domain per tunnel even though the routing uses specific subnets.

Note: The encryption domain defines the “interesting traffic” that is encrypted within the tunnel. Don't create multiple encryption domains to accommodate the various subnets in the Oracle Cloud Infrastructure virtual cloud network (VCN) or your on-premises network. Instead, summarize the subnets into a single *supernet* (several subnets combined or summarized into one network with a single CIDR prefix). For example, if your VCN network is 10.40.0.0/17 and 10.40.128.0/17 and your on-premises network is 10.0.0.0/18, 10.0.64.0/18, 10.0.128.0/18, and 10.0.192.0/18, you could use any-to-any or 10.0.0/16 to 10.40.0.0/16 to create a single encryption domain.

VPN Connect Plus FastConnect

You might need to upgrade your connection and deploy a FastConnect solution to Oracle Cloud. FastConnect is a solution that lets you connect to Oracle Cloud through a private connection. FastConnect provides better performance and higher bandwidth than VPN Connect. Oracle offers the following types of connections through FastConnect:

TABLE 2. FASTCONNECT OPTIONS

FastConnect	Description
Oracle provider	This connectivity option is suitable if you plan to use or are already using network connectivity services from any Oracle FastConnect partner. Depending on your partner, you might have to order redundant cloud connectivity services from the Oracle FastConnect partner. For a list of data center locations, see Oracle FastConnect Partners .

FastConnect	Description
Third-party provider	This connectivity option is suitable if you have existing relationships with certain network carriers, or if your on-premises or remote data center location is not served by any of Oracle's FastConnect partners.
Colocation	This connectivity option is suitable if you already have presence at an Oracle FastConnect location or want to establish a colocation presence at one. You can order two such connections into a data center for redundancy.

For more information about each of these options, see the [FastConnect documentation](#).

In Figure 6, the private cloud represents the provider or a cross-connect if you are colocated at a FastConnect location. In this solution, you continue to use VPN Connect, but you use it as the backup path rather than the primary path. Over FastConnect, you use BGP from the DRG to Edge 2 to exchange routes.

Don't deploy FastConnect and VPN Connect from the same edge device in your on-premises network (which would create a single point of failure). Deploy the services using separate edge devices, as shown in Figure 6.

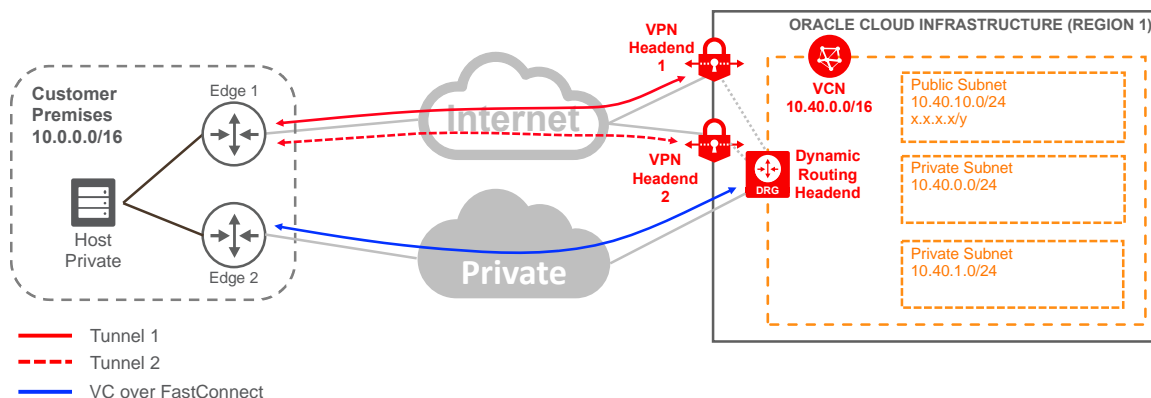


Figure 6. FastConnect Plus a Single VPN Connect Connection

For routing, follow the same approach as the previous solution, in which you advertise more-specific routes through the primary path (VC over FastConnect) and less-specific routes through the backup path (VPN Connect). The DRG learns your on-premises network subnets through BGP over FastConnect, and over VPN Connect it can use static routing or BGP. Within your network, manipulate your routing to prefer routes learned through FastConnect over routes learned through VPN Connect. For example, use AS prepend or local preference.

Figure 7 shows the routing configuration on both ends of the connection.

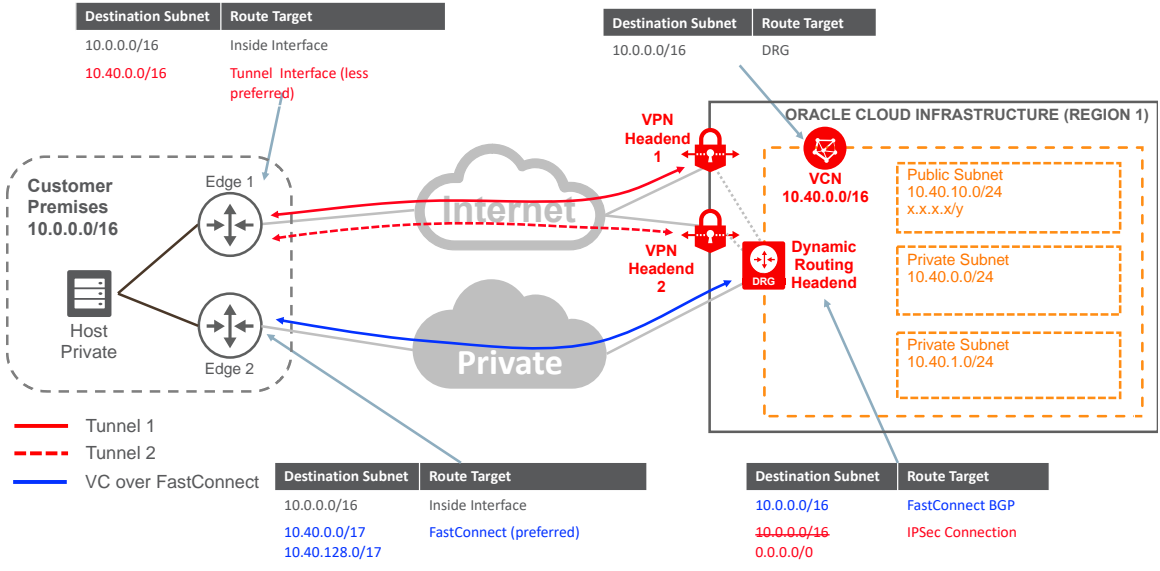


Figure 7. Routing for FastConnect Plus a Single VPN Connect Connection

Redundant FastConnect

As stated in the previous section, Oracle offers three types of FastConnect: with an Oracle provider, with a third-party provider, and colocation with Oracle. For simplicity, the private cloud shown in Figure 8 represents any of the FastConnect options. If there is a specific requirement for any of the FastConnect types, it will be called out.




Figure 8. Single FastConnect

You might have started your deployment to the cloud with a single FastConnect connection and a private virtual circuit (VC), as shown in Figure 8. If you deployed FastConnect with an Oracle provider, you connected to the provider network at some POP (point-of-presence), or the provider was already your backbone provider. If you deployed FastConnect with a third-party provider, you

requested a circuit from the provider's network to your on-premises network. If you were colocated with Oracle, you requested a cross-connect from the facility provider to connect to Oracle.

To build redundancy with FastConnect, pay special attention to the physical connectivity to ensure that it's redundant and diverse. As you work with partners and carriers, ensure that they understand the physical connectivity of your existing connection so that they can provide the diversity you require.

Within some FastConnect locations (Phoenix, Ashburn, Frankfurt, and London), Oracle provides two physical locations where partners can connect to Oracle. Within each location, Oracle provides redundant edge devices, as shown in Figure 9.

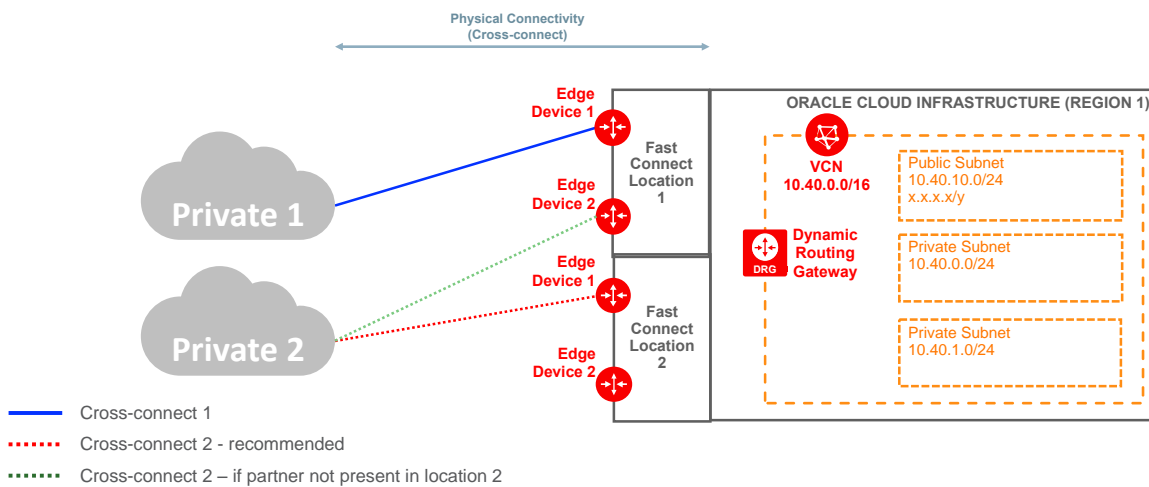


Figure 9. Physical Overview of an Oracle FastConnect Region

Following our example, you already have the cross-connect 1 (blue), which terminates at Location 1, Edge Device 1. When you deploy your second FastConnect connection, we recommend terminating it at Location 2, as depicted by the red line. It could be that an Oracle provider, the third-party provider, or you are colocated only at Location 1. In that case, you don't have location diversity, but you do have hardware diversity because the second FastConnect connection would terminate at Edge Device 2 in Location 1, as shown by the green line in Figure 9.

Note: Check with your Oracle partner to determine their redundancy with Oracle.

After you have worked out the physical redundancy with the Oracle provider, third party-provider, or your colocation provider, you will have two redundant FastConnect connections into Oracle Cloud, as shown in Figure 10.

Now you need to decide which connection is primary and which is backup. You can select based on the performance or capacity of the connections. In this document, Private 1 cloud is the primary and Private 2 is the backup. Over the second FastConnect connection, you also need to create a virtual circuit (VC) and establish BGP with your second edge device.




Figure 10. Redundant FastConnect

Manipulate the routing accordingly at both ends of the connection (DRG and customer edge devices) to route the traffic through the correct path. You can achieve this by doing AS prepend or advertising more and less specific routes, or using another technique.

With FastConnect, BGP is used to exchange routes between your on-premises network and Oracle.

- If you use a Layer 2 provider, the virtual circuit's BGP session is between your edge device and Oracle.
- If you use a Layer 3 provider, the virtual circuit's BGP session is between the provider and Oracle.

We recommend advertising more-specific subnets through the primary path and less-specific subnets over the backup path, as shown in Figure 11.

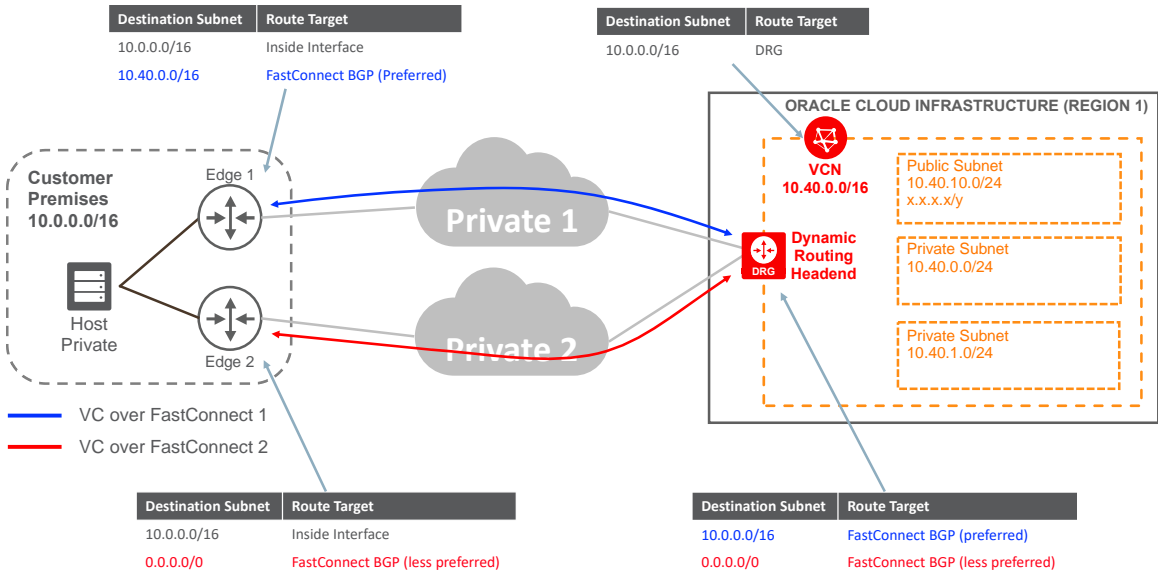


Figure 11. Redundant FastConnect—Routing

References

- [VPN Connect documentation](#)
- [Oracle Cloud Infrastructure Networking documentation](#)
- [FastConnect documentation](#)







Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2019, 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. 0919

Connectivity Redundancy Guide
September 2019
Author: Javier Ramirez

