Architecture Diagrams

Traffic Segmentation Options in AWS Direct Connect

Traffic Segmentation Options in AWS Direct Connect: Architecture Diagrams

Table of Contents

Traffic Segmentation Options in AWS Direct Connect

Publication date: **November 9, 2023 (Diagram history)**

This architecture shows traffic segmentation options in AWS Direct Connect, include using several private VIFs in the VPC or Direct Connect gateway, using AWS Transit Gateway Connect attachments over transit VIF, and using AWS Site-to-Site VPN connections over transit VIF.

Segment Your Traffic over AWS Direct Connect by Using Several Private VIFs in the VPC or Direct Connect Gateway Diagram

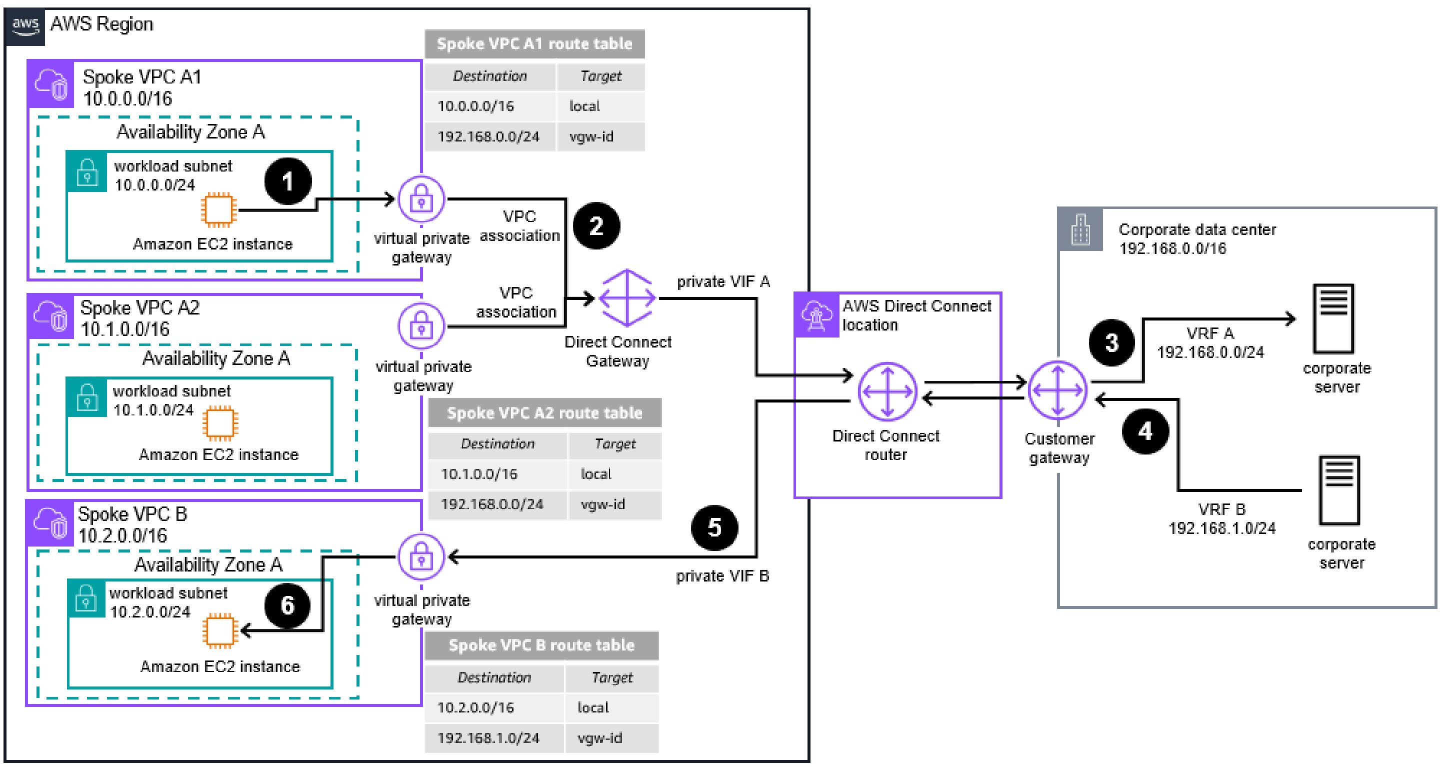

Private virtual interfaces (VIFs) are used to access Amazon Virtual Private Cloud (Amazon VPC) using private IP addresses. Each VIF is configured with an unique virtual local area network (VLAN) tag, which means you can segment your traffic by using one VIF per virtual routing and forwarding (VRF) you have in your data center. Take into account the limits in the number of VIFs, depending the type of Direct Connect connection you have. You can have 50 private or public VIFs with a dedicated connection, and one private or public VIF with a hosted connection. In addition, you can have 30 private VIFs associated to the same Direct Connect gateway. You can check the quotas in the Direct Connect user guide.

1. Traffic initiated from an instance in Spoke VPC A and destined for the corporate data center server (in VRF A) is routed to the virtual private gateway according to the Spoke VPC A route table.
2. The virtual private gateway is associated with the **Direct Connect** gateway, which uses a private VIF (VLAN) to connect to the corporate data center. To achieve segmentation, use several private VIFs. The on-premises router must be configured to use a different VIF (VLAN) connection per VRF configured.
3. The on-premises router forwards the traffic from private VIF A to VRF A and the destination server.
4. Traffic from the corporate data center's VRF B destined for Spoke VPC B is first sent to the customer gateway located in the data center.
5. The customer gateway uses private VIF B – over the Direct Connect link – to send the traffic to the virtual private gateway in VPC B.
6. Traffic is forwarded to the destination according to the Spoke VPC B route table.

Segment Your Traffic over AWS Direct Connect by Using Several Transit VIFs and Direct Connect Gateways Diagram

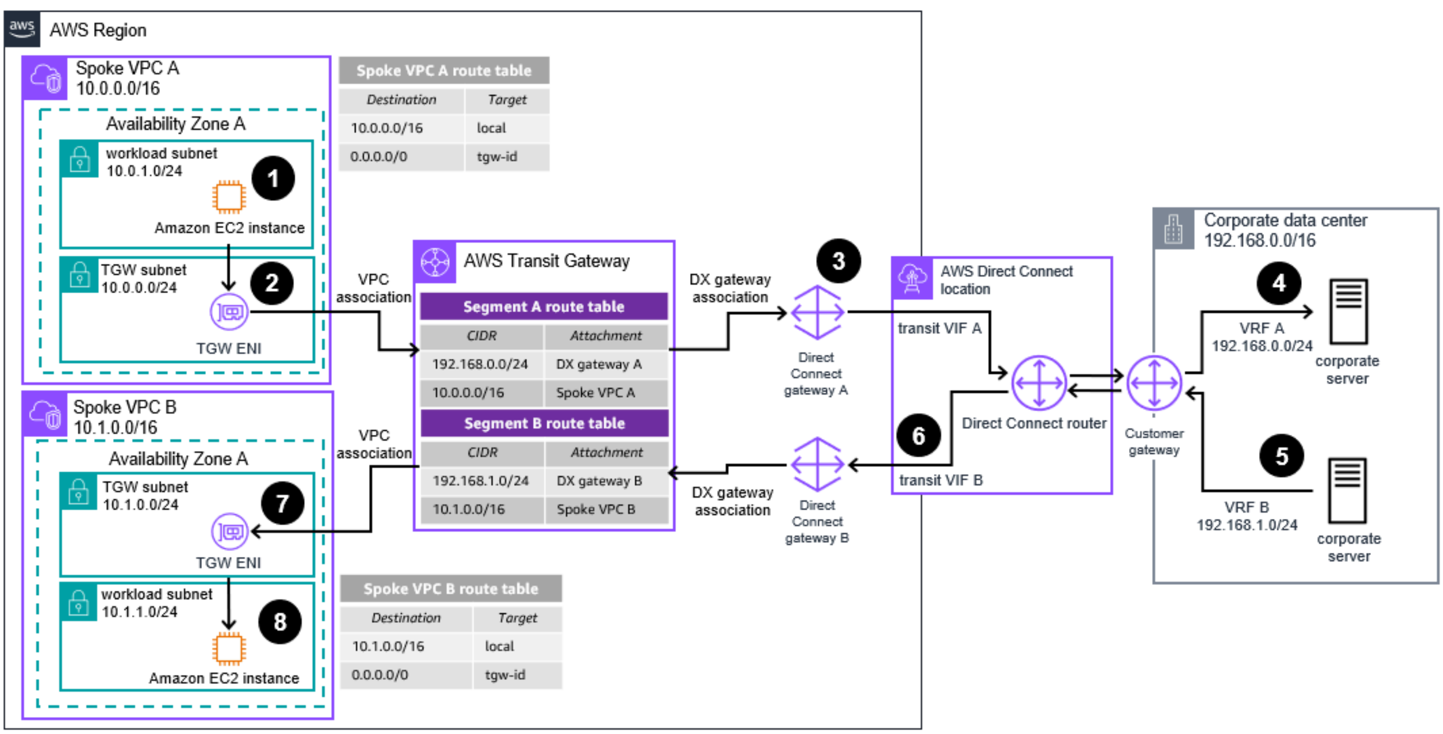

Transit Virtual Interfaces (VIFs) are used to access AWS Transit Gateway by way of an Direct Connect gateway. Each VIF is configured with an unique virtual local area network (VLAN) tag, which means you can segment your traffic by using one VIF per virtual routing and forwarding (VRF) you have in your data center. Take into account the limits in the number of VIFs, depending the type of Direct Connect connection you have. You can have four transit VIFs with a max combination of 51 private, public, and transit VIFs per dedicated connection. For hosted connections, you can have only one transit VIF. In addition, you can have 30 transit VIFs associated with the same Direct Connect gateway. You can check the quotas in the Direct Connect user guide.

1. Traffic initiated from an instance in Spoke VPC A and destined for the corporate data center server (in VRF A) is routed to **AWS Transit Gateway** according to the Spoke VPC A route table.
2. Transit Gateway is associated with several Direct Connect gateways (DXGW), each one of them using a transit VIF (VLAN) to connect to the corporate data center, achieving traffic segmentation. The on-premises router must be configured to use a different VIF (VLAN) connection per VRF configured.
3. Traffic is forwarded to DXGW A according to the Segment A **Transit Gateway** route table associated with the VPC A attachment.
4. The on-premises router forwards the traffic from the transit VIF A to VRF A and the destination server.
5. Traffic from the corporate data center's VRF B destined for Spoke VPC B is first sent to the customer gateway located in the data center.
6. The customer gateway uses transit VIF B – over the **Direct Connect** link – to send the traffic to **Transit Gateway** Segment B.
7. Traffic is forwarded to VPC B according to the Segment B **Transit Gateway** route table.
8. Traffic is forwarded to the destination according to the VPC B route table.

Use AWS Transit Gateway Connect Attachments and AWS Direct Connect to Extend your On-Premises VRFs over Transit VIFs Diagram

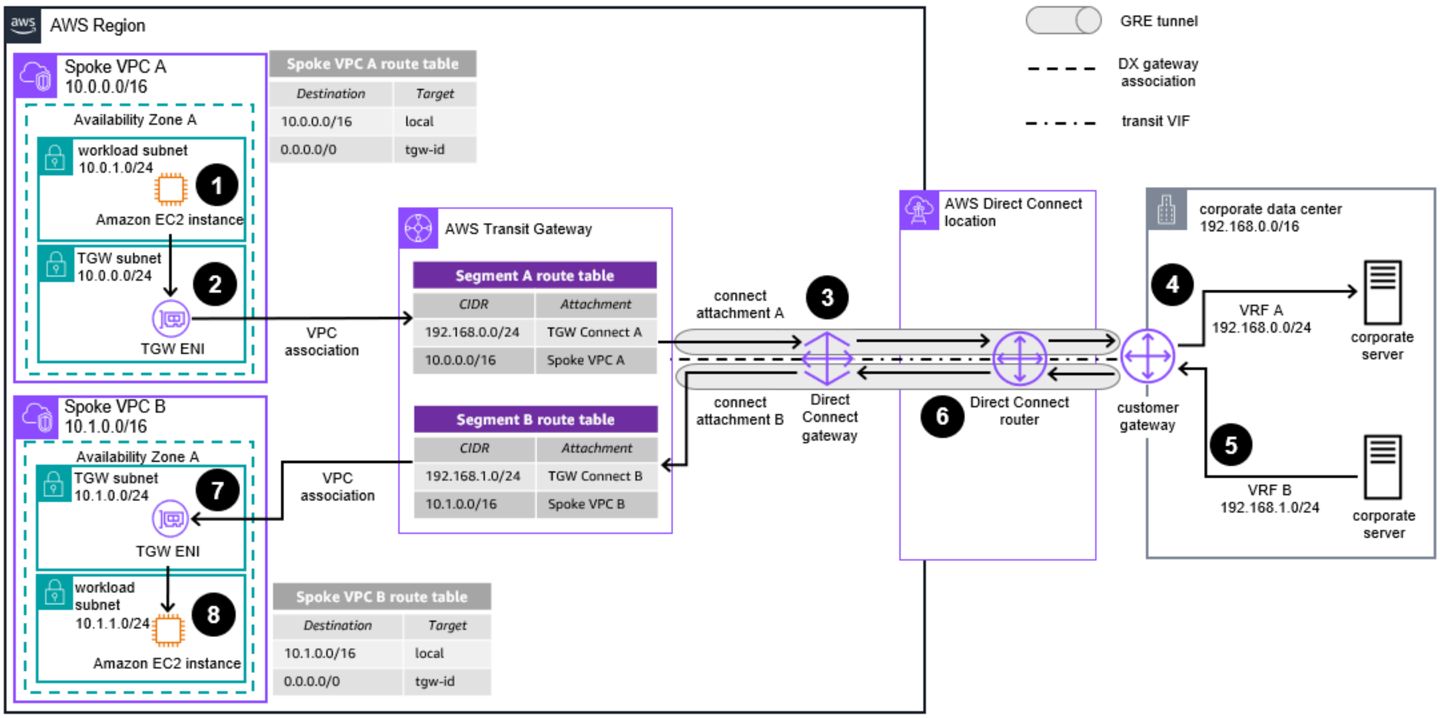

With a hosted connection, you get one transit VIF; with a dedicated connection you get four. When you have more VRFs than transit VIFs, you can use several Transit Gateway Connect attachments (GRE tunnels) over a transit VIF to segment the traffic, creating one connect attachment per VRF you have in your data center. For more information about this use case, refer to Using AWS Transit Gateway Connect to extend VRFs and increase IP prefix advertisement.

1. Traffic initiated from an instance in Spoke VPC A and destined for the corporate data center server (in VRF A) is routed to the **Transit Gateway** elastic network interface (TGW ENI) according to the Spoke VPC A route table.
2. Traffic is forwarded to **Transit Gateway**. Traffic is then routed to the corporate data center through the **Transit Gateway** Connect attachment A according to the Segment A route table.
3. The **Transit Gateway** Connect attachment uses the **Direct Connect** connection as transport and connects **Transit Gateway** to the corporate data center router using Generic Routing Encapsulation (GRE) tunneling and Border Gateway Protocol (BGP). The on-premises router must be configured to use a different connect attachment per VRF configured.

4. The on-premises router forwards the traffic from Connect attachment A to VRF A and the destination server.

5. Traffic from the corporate data center's VRF B destined for Spoke VPC B is first sent to the customer gateway located in the data center.

6. The customer gateway uses the GRE tunnel of the **Transit Gateway** Connect attachment B – over the **Direct Connect** link – to send the traffic to **Transit Gateway**.

7. Traffic is forwarded to the Spoke VPC B attachment according to the **Transit Gateway** Segment B route table.

8. The TGW ENI of Spoke VPC B forwards the traffic to the destination.

Use AWS Site-to-Site VPN Attachments (Public VIFs) and AWS Direct Connect to Extend Your On-Premises VRFs Diagram

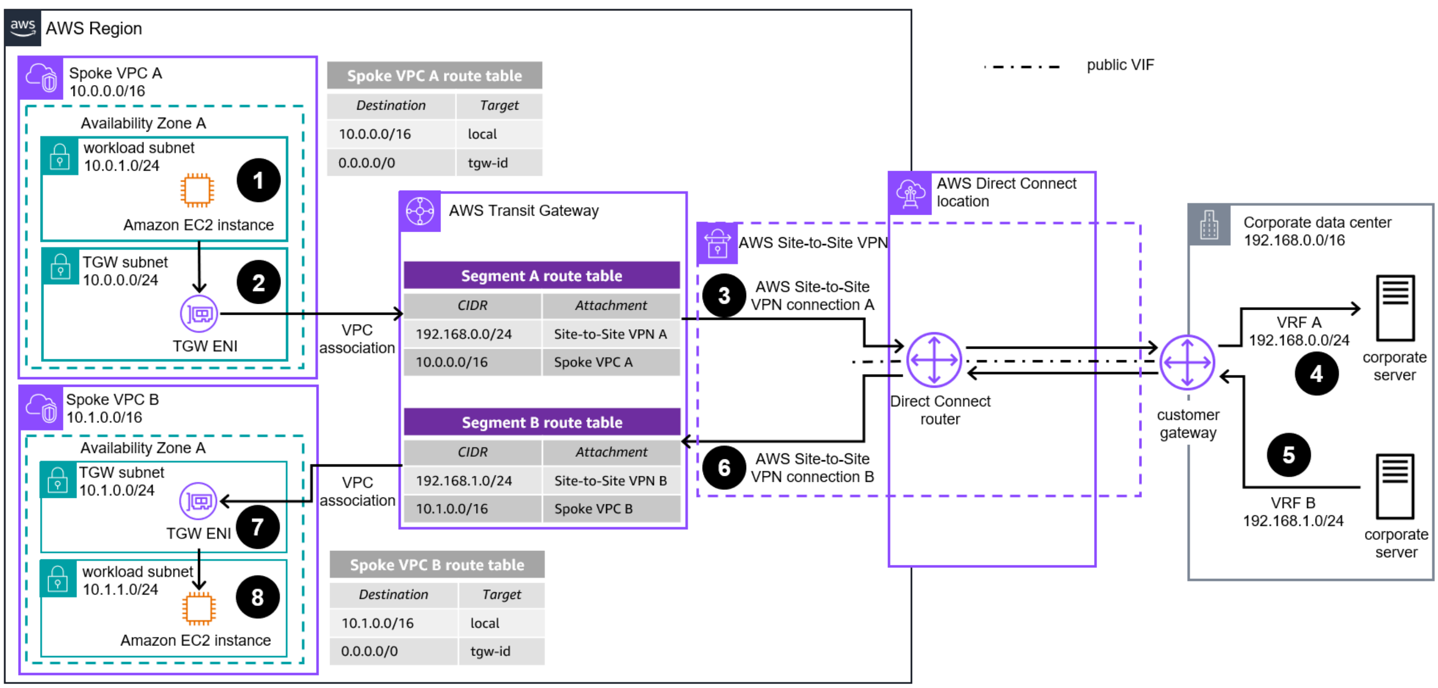

If your customer gateway does not support Generic Routing Encapsulation (GRE), you can still segment your traffic by creating one AWS Site-to-Site VPN connection per each VRF you have in your data center. You can use a public VIF as underlaying transport.

1. Traffic initiated from an instance in Spoke VPC A and destined for the corporate data center server (in VRF A) is routed to the **Transit Gateway** elastic network interface (TGW ENI) according to the Spoke VPC A route table.

2. Traffic is forwarded to **Transit Gateway**. Traffic is then routed to the corporate data center through **Site-to-Site VPN** connection A according to the Segment A route table.

3. The **Site-to-Site VPN** connection is created using a public VIF as transport (VLAN). The on-premises router must be configured to use a different VPN connection per VRF configured.

4. The on-premises router forwards the traffic from **Site-to-Site VPN** connection A to the VRF A and the destination server.

5. Traffic from the corporate data center's VRF B destined for Spoke VPC B is first sent to the customer gateway located in the data center.

6. The customer gateway chooses **Site-to-Site VPN** connection B to send the traffic to the **Transit Gateway** over the **Direct Connect** link.

7. Traffic is forwarded to the Spoke VPC B attachment according to the **Transit Gateway** Segment B route table.

8. The TGW ENI of Spoke VPC B forwards the traffic to the destination.

Use AWS Site-to-Site VPN Private IP VPN Attachments and AWS Direct Connect to Extend Your On-Premises VRFs Over Transit VIFs Diagram

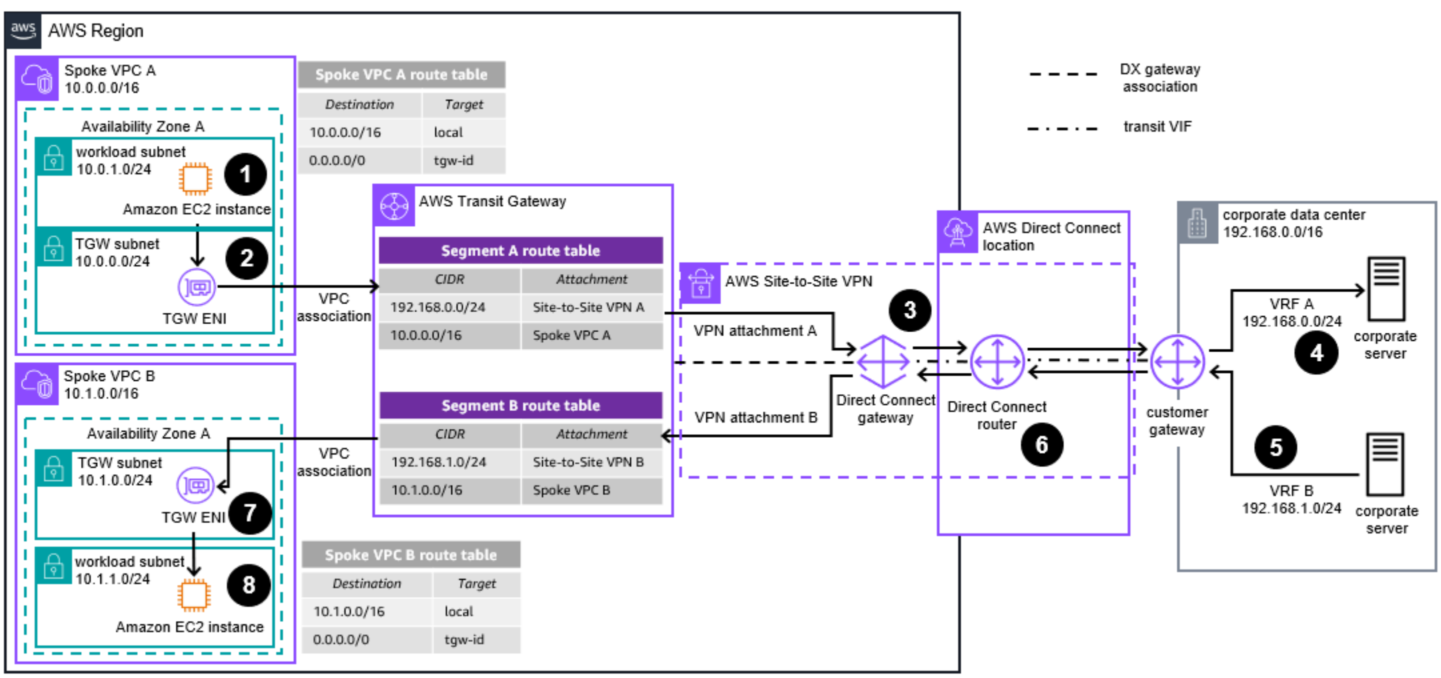

If your customer gateway does not support GRE (Generic Routing Encapsulation), you can still segment your traffic by creating one AWS Site-to-Site VPN connection (with two tunnels) per each VRF you have in your data center. With the Site-to-Site VPN Private IP VPN feature, you can create the VPNs using private IPs with the AWS Direct Connect gateway and transit VIF as underlying transport.

1. Traffic initiated from an instance in Spoke VPC A and destined for the corporate data center server (in VRF A) is routed to the **AWS Transit Gateway** elastic network interface (TGW ENI) according to the Spoke VPC A route table.
2. Traffic is forwarded to **Transit Gateway**. Traffic is then routed to the corporate data center through **Transit Gateway** VPN attachment A according to the Segment A route table.
3. The **Transit Gateway Site-to-Site VPN** attachment uses the **Direct Connect** connection as transport, and connects **Transit Gateway** to the corporate data center router using one IPSec VPN connection (with two tunnels). The on-premises router must be configured to use a different VPN attachment per VRF configured.
4. The on-premises router forwards the traffic from VPN attachment A to VRF A and the destination server.
5. Traffic from the corporate data center's VRF B destined for Spoke VPC B is first sent to the customer gateway located in the data center.
6. The customer gateway uses **Site-to-Site VPN** connection B – over the **Direct Connect** link – to send the traffic to **Transit Gateway**.
7. Traffic is forwarded to the Spoke VPC B attachment according to the **Transit Gateway** Segment B route table.
8. The TGW ENI of Spoke VPC B forwards the traffic to the destination.

Further reading

For additional information, refer to

- AWS Architecture Icons
- AWS Architecture Center
- AWS Well-Architected
- Direct Connect user guide

Contributors

Contributors to this reference architecture diagram include:

- Pablo Sánchez Carmona, Specialist Solutions Architect, Amazon Web Services
- Thaddeus Worsnop, Senior Solutions Architect, Amazon Web Services

Diagram history

To be notified about updates to this reference architecture diagram, subscribe to the RSS feed.

Change	Description	Date
Diagram updated	Added *Segment Your Traffic over AWS Direct Connect by Using Several Transit VIFs and Direct Connect Gateways*	November 9, 2023
Initial publication	Reference architecture diagram first published.	June 30, 2022

> **Note**
>
> To subscribe to RSS updates, you must have an RSS plugin enabled for the browser you are using.