# Cisco SD-WAN

# Cloud First – SDCI Case Study

4Dachs Consulting

August 2023

# Contents

## About the Guide

The designs discussed within this document are presented in the form of a case study for a fictional customer – 4Dachs Consulting – who is taking a cloud-first approach to providing software services to its customers by leveraging benefits of Cisco Software-Defined Cloud Interconnect (SDCI). Although 4Dachs Consulting is a fictional customer, the designs presented within this guide are based on actual customer deployments. The purpose of this document is as follows:

- Present multiple design models for site-to-cloud and cloud-to-cloud connectivity.
- Highlight the benefits of the Cisco SDCI solution for connectivity between different public IaaS/PaaS Cloud Service Providers (CSPs) and from corporate sites to public CSPs in a multi-cloud use case.
- Present some of the considerations that a network engineer will need to focus attention upon when implementing a Cisco SD-WAN design to include public IaaS/PaaS cloud connectivity via SDCI.

### Audience

The audience for this document includes network design engineers, network operations personnel, Cloud Ops, and security operations personnel who wish to implement Cisco SD-WAN networks.

## Background

4Dachs Consulting is a business consulting company operating in the U.S. and Europe that provides software services to its customers. Their approach for providing software services is through web-based Software-as-a-Service (SaaS) offerings. Rather than expending its limited capital budget on constructing data centers and purchasing servers, 4Dachs Consulting instead leverages public IaaS/PaaS CSPs such as Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), and Oracle Cloud Infrastructure (OCI) for the infrastructure and platforms necessary for its SaaS offerings.

4Dachs Consulting's software services are primarily based on newer microservices architectures for application development, where each VM (or container running on a VM) implements a single service of the larger application. Each microservice may contain different pieces of application logic, messaging functionality, database functionality, web-based front-end services, etc. Hence, each call to a VM/container hosting the web-based front-end may initiate multiple calls to other VMs/containers running on different back-end servers, to generate the response.

Initially, 4Dachs Consulting leveraged a single IaaS/PaaS CSP for its software services. 4Dachs Consulting hosted its software services in two separate regions within the IaaS/PaaS CSP – U.S. and Europe, as shown in the following figure.

**Figure 1.** **4Dachs Consulting - Initial Site and Cloud Connectivity Design using Cisco Catalyst SD-WAN**

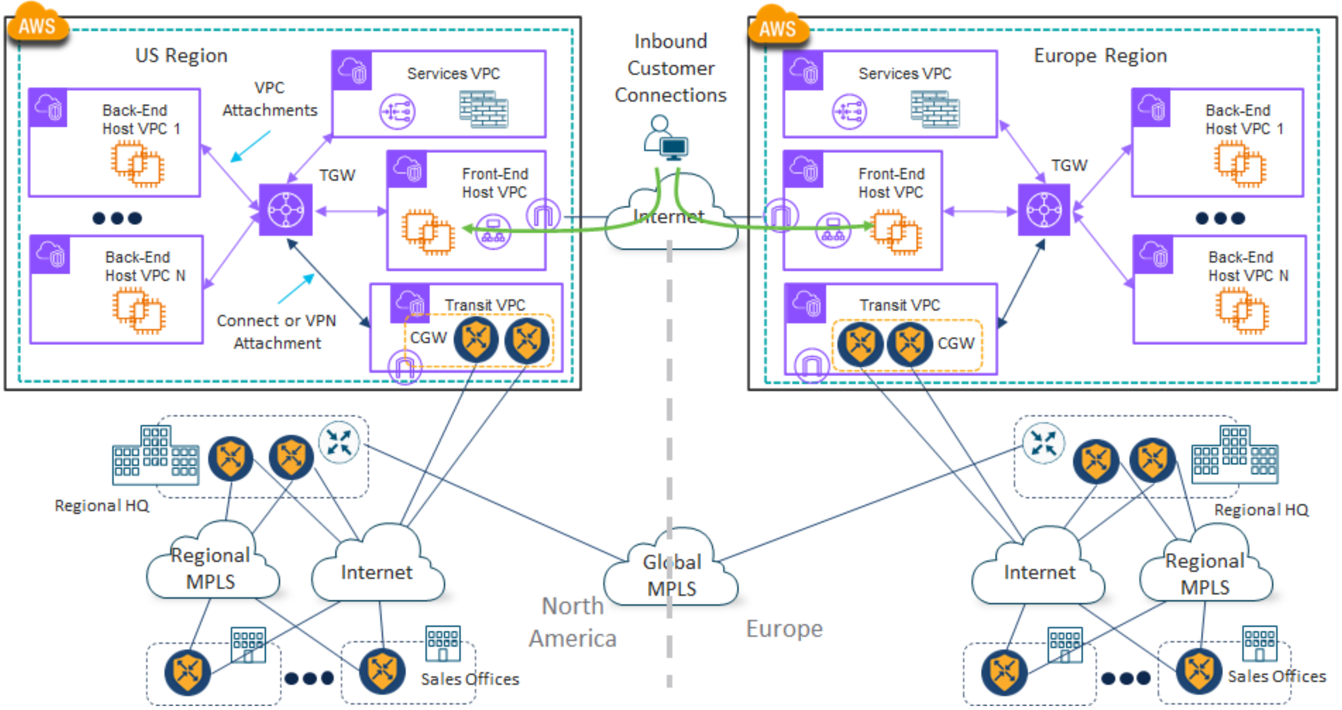

Servers, reachable from the Internet through load balancers, provide customers with access to the web-based front-end of 4Dachs Consulting's software services. Back-end services are supported through additional servers located within other host VPCs/vNets within the same CSP region, with the option of inserting a security appliance between them. An example of the front-end and back-end application data flows with security appliances is shown in the following figure.

**Figure 2.** **Example Front-End and Back-End Data Flows**

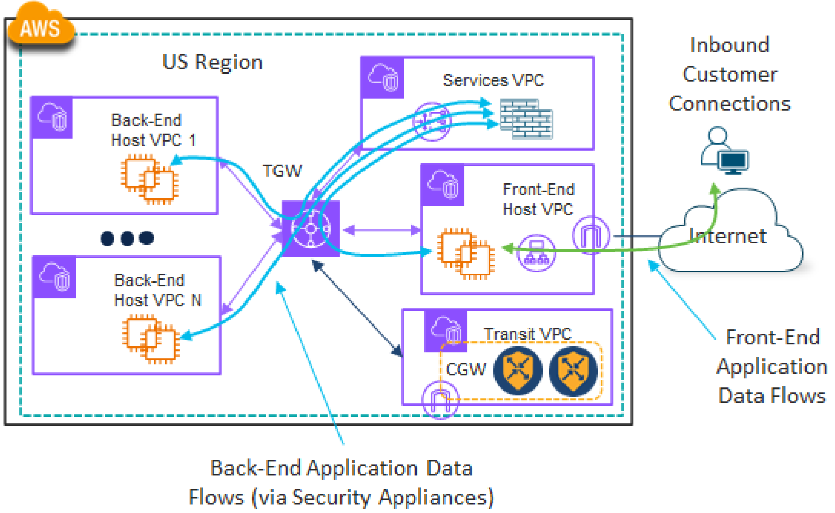
AWS
US Region
Services VPC
Back-End Host VPC 1
TGW
Front-End Host VPC
Back-End Host VPC N
Transit VPC
CGW
Inbound Customer Connections
Internet
Front-End Application Data Flows
Back-End Application Data Flows (via Security Appliances)

4Dachs Consulting leverages the high-speed infrastructure of the IaaS/PaaS CSP within each region for connectivity between host VPCs/vNets that supported the various microservices – which together provide the software services to its customers.  Each region is fully redundant and independent of the other region – meaning there is no requirement for traffic between regions.

| **Technical Note:** |
|---|
| Redundancy through deployment of services in different Availability Zones (AZs) within the IaaS/PaaS CSP is not shown in the figures for simplicity. |

4Dachs Consulting utilizes Cisco Catalyst SD-WAN for regional connectivity between its corporate sites over both a private MPLS service as well as the Internet.  Connectivity between geographic regions (U.S. and Europe) is via a traditional global MPLS backbone.

Connectivity from corporate sites to the public IaaS/PaaS CSP was initially accomplished by extending the Cisco Catalyst SD-WAN fabric into the CSP, via the Internet, within each region.  This was done through a Cloud Gateway (CGW) consisting of a redundant pair of Cisco Catalyst 8000v routers instantiated within the CSP.  One advantage of this design is that traffic is encrypted from the corporate sites all the way into the public IaaS/PaaS CSP network.

Connectivity from corporate sites to the public IaaS/PaaS CSP is primarily for monitoring and maintenance of the application services deployed within the VMs/containers within the CSP.  No application traffic (traffic between microservices which together provide the application experience to the customer) traverses the SD-WAN connections between the public IaaS/PaaS CSP and the corporate sites.  Hence, the bandwidth and latency requirements for connectivity from the corporate sites to the public IaaS/PaaS CSP are minimal.

4Dachs Consulting is currently not doing any network segmentation within their corporate sites, or out to the cloud.  Hence, they extend only a single Service VPN to the public IaaS/PaaS CSP.

## Business Challenge

Over time 4Dachs Consulting found it necessary to expand to a multi-cloud architecture by utilizing multiple public IaaS/PaaS CSPs to provide a better experience and service to its customers.  This meant that the microservices which constituted a single application could be split across infrastructure located in multiple public IaaS/PaaS CSPs.  This in turn brought up new concerns regarding the application experience for their customers, since the performance of the overall application now depended on traffic flows that potentially could be crossing between different public IaaS/PaaS CSPs.

4Dachs Consulting considered the following methods of providing cloud-to-cloud connectivity between public IaaS/PaaS CSPs:

- Connectivity via a Software-Defined Cloud Interconnect (SDCI) provider, also leveraging private connectivity options of the public IaaS/PaaS CSPs – such as AWS Direct Connect and Azure ExpressRoute.
- Connectivity via the Cisco Catalyst SD-WAN fabric using the Internet.
- Connectivity via a regional MPLS provider, leveraging private connectivity options of the public IaaS/PaaS CSPs – such as AWS Direct Connect and Azure ExpressRoute.

After a thorough analysis of each of the options, 4Dachs Consulting chose to migrate to a model in which connectivity to multiple public IaaS/PaaS CSPs would be provided through an SDCI provider. The decision was based primarily on their business requirements as discussed in the following section.

## Cloud-to-Cloud Connectivity via a Software-Defined Cloud Interconnect (SDCI) Partner

4Dachs Consulting initially considered the option of extending their private network into public IaaS/PaaS CSPs through a traditional colocation provider. This would entail provisioning space within the colocation provider (referred to as a cage), installing and maintaining their own hardware (physical routers and switches or virtual instances running on compute hardware) within the cage, and provisioning physical cross-connects between their cage and the cages of the public IaaS/PaaS CSPs for private connectivity (AWS Direct Connect, Azure ExpressRoute, etc.). However, they chose not to pursue this option based on the perceived cost of provisioning the cage and hardware, and the perceived complexity of maintaining the equipment within the colocation provider. Instead, they chose to look at the “as-a-service” option of extending their SD-WAN network to the public IaaS/PaaS CSPs through an SDCI provider – viewing this as a more cost effective and flexible option.

With an SDCI partner, 4Dachs Consulting realized they still had several design options when leveraging private connections (AWS Direct Connect, Azure ExpressRoute, etc.) for cloud-to-cloud connectivity between public IaaS/PaaS CSPs.

- Option 1: Extending the Cisco Catalyst SD-WAN fabric into the SDCI partner network through Cisco Interconnect Gateways (ICGWs) and leveraging private connectivity (AWS Direct Connect, Azure ExpressRoute, etc.) via Virtual Cross-Connects (VXCs) between the SDCI partner and the public IaaS/PaaS CSP through a Service VPN. In other words, this extends only the Service VPN side of the Cisco Catalyst SD-WAN, and not the fabric itself (WAN transport side) into the public IaaS/PaaS CSP.
- Option 2: Extending the Cisco Catalyst SD-WAN fabric into the SDCI partner network through Cisco ICGWs and leveraging private connectivity via VXCs between the SDCI partner and the public IaaS/PaaS CSP through a WAN transport. In other words, this extends the Cisco Catalyst SD-WAN fabric itself into the public IaaS/PaaS CSP via dedicated connectivity to the CGWs.

Each of these options is discussed in detail in the following sections.

For their SDCI partner, 4Dachs Consulting chose Megaport. This is primarily because Megaport supported the Cisco Catalyst 8000v operating in SD-WAN mode as a Cisco Interconnect Gateway (ICGW) – otherwise known as a Megaport Virtual Edge (MVE) – as of Cisco Catalyst SD-WAN software release 17.9/20.9. This was the long-term software release 4Dachs Consulting was using within their Cisco Catalyst SD-WAN deployment, when evaluating SDCI partners.

### Option 1: Extending the SD-WAN Fabric into the SDCI Partner Network Only

With SDCI Option 1, 4Dachs Consulting would extend their Cisco Catalyst SD-WAN fabric into the SDCI partner network via the Internet, through a pair of Catalyst 8000v SD-WAN routers functioning as Cisco Interconnect Gateways (ICGWs) for redundancy. They would then leverage the peering relationships between the SDCI partner and the various public IaaS/PaaS CSPs to provide private connectivity via AWS Direct Connect, Azure ExpressRoute, etc., from the ICGWs through Virtual Cross Connects (VXCs). This could be automated through the Cisco Cloud onRamp for Multi-Cloud Interconnect workflow for ease of deployment.

**Technical Note:**

As of Cisco Catalyst SD-WAN software release 17.9/20.9, the Cisco Cloud onRamp for Multi-Cloud Interconnect workflow supports connectivity to AWS, Azure, and GCP.

An example of this design in shown in the following figure.

**Figure 3. Cloud-to-Cloud Connectivity via SDCI Provider - Option 1**

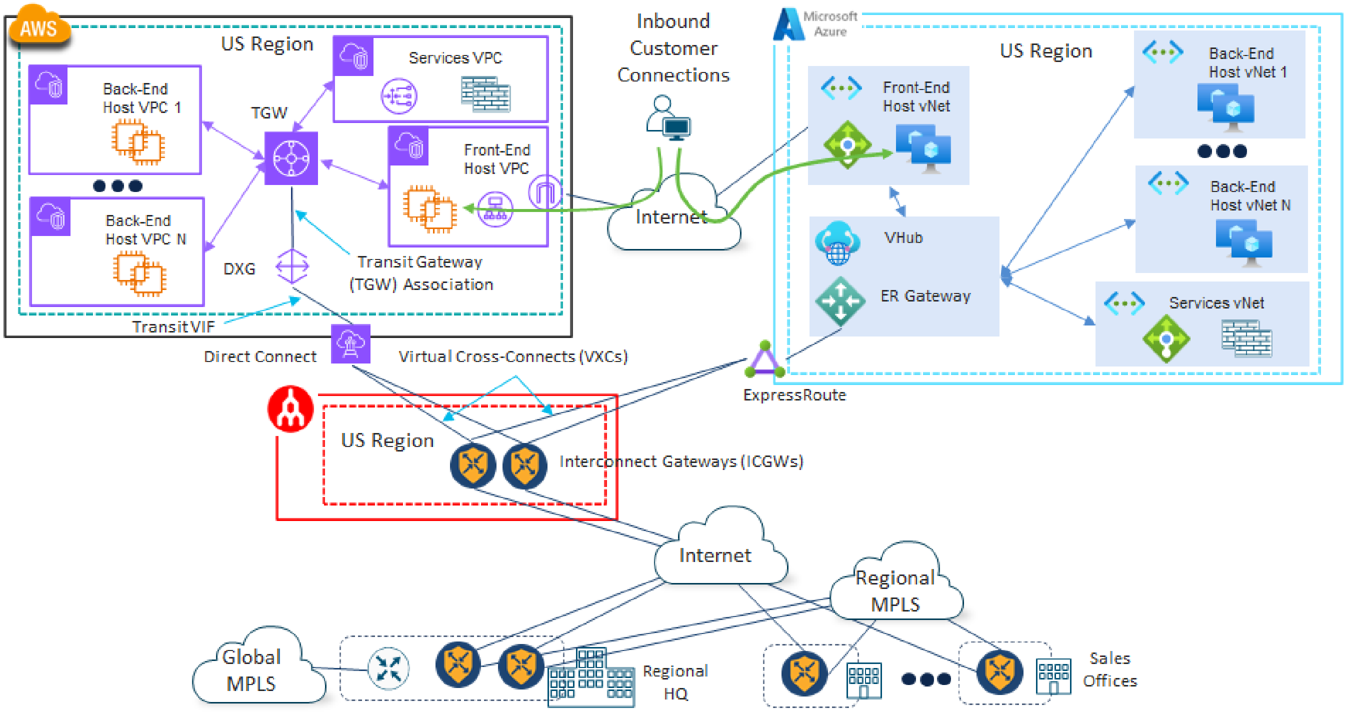

4Dachs Consulting would no longer need the existing CGWs within the public IaaS/PaaS CSPs with this design. Therefore, they could be removed, resulting in some cost savings due to the underlying instance charges and licensing for the Catalyst 8000v virtual routers. However, 4Dachs Consulting would be relying completely on the built-in redundancy within the private connectivity to each of the IaaS/PaaS CSPs for high availability.

### AWS Connectivity Details

4Dachs Consulting was already using Transit Gateways (TGWs) within AWS in each region and wanted to continue with the same design. Hence, within the SDCI provider network, each Virtual Cross-Connect (VXC) would need to be provisioned as a **Connection VIF Type - Transit** from each of the ICGWs to AWS to support connection to a Direct Connect Gateway (DXG). The Direct Connect Gateway (DXG) would then be associated with the Transit Gateway (TGW).

Within AWS, **Hosted Connections** would need to be provisioned. **Hosted Connections** support varying speeds (50 Mbps, 100 Mbps, 200 Mbps, 300 Mbps, 400 Mbps, 500 Mbps, 1 Gbps, 2 Gbps, 5 Gbps, and 10 Gbps) up to 10 Gbps – which is more than enough for 4Dachs Consulting's cloud-to-cloud requirements. 4Dachs Consulting would choose to begin with 1 Gbps of bandwidth on each VXC, since the latency between the public IaaS/PaaS CSPs was the main concern, not the amount of bandwidth. Connectivity through an SDCI partner would also allow 4Dachs Consulting the ability to change the bandwidth as necessary, up to the maximum of 10 Gbps, rather than being locked into a fixed bandwidth as with a regional MPLS carrier.

**Technical Note:**

Note that with an AWS **Hosted Connection**, a new **Hosted Connection** must be provisioned when changing bandwidth. Hence, 4Dachs Consulting would need to take into consideration any potential temporary disruption of service to its customers when increasing or decreasing bandwidth between the SDCI provider and the public IaaS/PaaS CSP.

Additionally, note that although a **Hosted Connection** supports varying speeds from 50 Mbps through 10 Gbps, the **Connection VIF Type** (which can be **Private**, **Public**, or **Transit**) may further constrain the bandwidth options which can be selected. For example, with a **Connection VIF Type – Transit** configuration the available bandwidth choices within the Cisco Cloud OnRamp for Multi-Cloud Interconnect workflow are 1 Gbps, 2 Gbps, 5 Gbps, and 10 Gbps.

### Azure Connectivity Details

Within Azure, 4Dachs Consulting wanted to continue using the vWAN architecture, with a vHub in each region. Hence, within the SDCI provider network, each VXC would be associated with an ExpressRoute circuit that terminates on an ExpressRoute Gateway (ER Gateway) within a vHub in the Azure region.

**Technical Note:**

When using the Cisco Cloud onRamp for Multi-Cloud Interconnect workflow, connectivity to Microsoft Azure via ExpressRoute requires a pair of Interconnect Gateways (ICGWs) within the Megaport fabric for redundancy. Primary and secondary ExpressRoute connections are provisioned on an ExpressRoute circuit, each terminating on a separate ICGW. This is the only configuration supported within Cisco Cloud onRamp for Multi-Cloud Interconnect.

This is different from connectivity to AWS via Direct Connect, which is provisioned in a non-redundant manner. A single Direct Connect connection is provisioned on a Direct Connect circuit to a single Interconnect Gateway (ICGW) within the Megaport fabric. If desired, a second Direct Connect connection can be provisioned on the Direct Connect circuit to a second ICGW for redundancy.

For this case study, redundant Direct Connect and ExpressRoute connections to a pair of ICGWs/MVEs were provisioned for both AWS and Azure connectivity.

ExpressRoute circuits can connect through a service provider or directly to Azure – as specified through the **Port Type** (**Provider** or **Direct**) setting. ExpressRoute connections support varying speeds depending upon the **Port Type** setting. Connections provisioned through an SDCI provider such as Megaport are of **Port Type – Provider** and support speeds (50 Mbps, 100 Mbps, 200 Mbps, 500 Mbps, 1 Gbps, 2 Gbps, 5 Gbps, and 10 Gbps) up to 10 Gbps. Again, this is more than enough for 4Dachs Consulting's cloud-to-cloud requirements.

4Dachs Consulting would choose to begin with 1 Gbps of bandwidth on each VXC, again since the latency between the public IaaS/PaaS CSPs was the main concern, not the amount of bandwidth. Since each regional site (U.S. and Europe) is essentially independent of each other, a **Standard** ExpressRoute circuit is all that would be required.

### Additional Considerations and Benefits

BGP routing would need to be configured between the public IaaS/PaaS CSPs to the ICGWs within the SDCI provider. This allows the private IP addressing prefixes within the VPCs and vNets to be exchanged between the public IaaS/PaaS CSPs for connectivity to be established.

**Technical Note:**

BGP peering between the public IaaS/PaaS CSP and the Interconnect Gateways (ICGWs) within the SDCI provider is automated when using the Cisco Cloud onRamp for Multi-Cloud Interconnect workflow.

4Dachs Consulting realized they would have to adhere to the maximum number of prefixes that could be sent over the private connections to each of the public IaaS/PaaS CSPs. However, they did not believe that this limitation would be an issue for their network.

Although there were no regulatory requirements that applied to 4Dachs Consulting, another area of consideration was that cloud-to-cloud traffic is un-encrypted at the network level with this design. However, at the application level, the data transfers between the back-end services were either REST-based, which made use of HTTPS as the transport, or had secure transport options. Hence if application-level encryption of the back-end services was maintained, network level encryption was not a concern for 4Dachs Consulting.

For 4Dachs Consulting, the primary business agility benefits of this design over using a regional MPLS carrier were as follows:

- The ability to rapidly provision the SDCI service for cloud connectivity via VXCs, versus the longer lead times for provisioning private connectivity to public IaaS/PaaS CSPs via a regional MPLS carrier. This allowed 4Dachs Consulting to expand into new public IaaS/PaaS CSPs, or remove existing CSPs as needed, if the CSP had a peering relationship with the SDCI partner.
- The ability to expand or contract the bandwidth to each public IaaS/PaaS CSP as needed to meet current business requirements, versus being locked into a long-term contract at a fixed bandwidth as with a regional MPLS carrier.

An example of the front-end and back-end data flows is shown in the following figure.

**Figure 4. Example New Front-End and Back-End Data Flows**

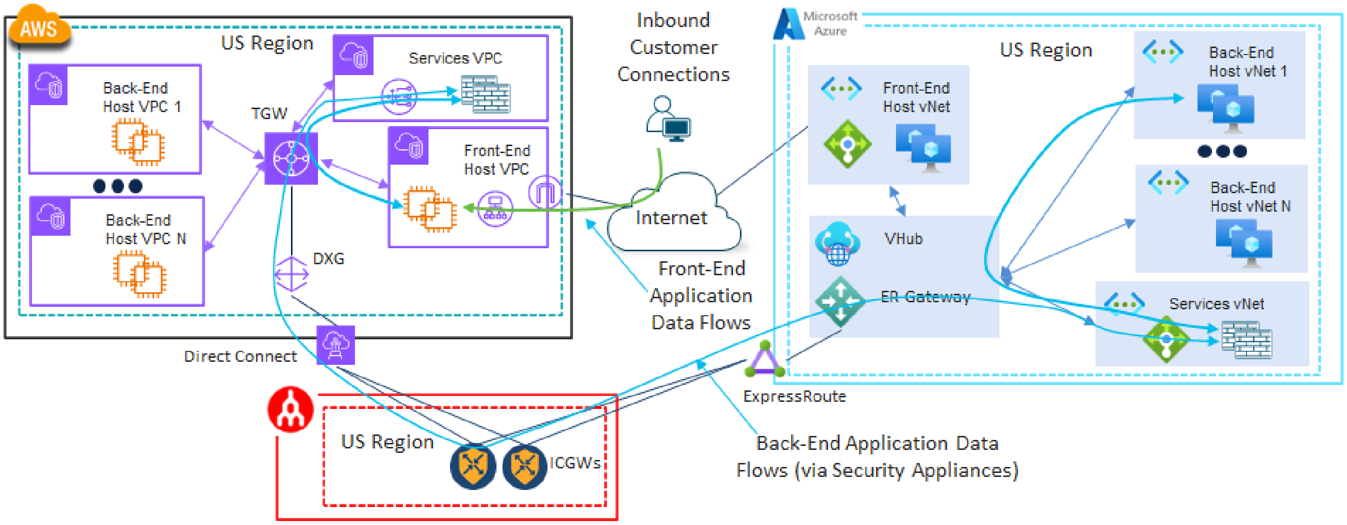

Although not the primary focus for 4Dachs Consulting currently, another benefit of this design is that it optimizes the connectivity from the corporate sites to the public IaaS/PaaS CSPs over traditional Internet connectivity which 4Dachs Consulting was using. Instead of taking multiple hops over the Internet from corporate sites to reach the public IaaS/PaaS CSP, traffic from corporate sites only needed to traverse the local “last-mile” ISP to reach the regional SDCI facility. 4Dachs Consulting felt that this positioned them well for any future considerations they may have regarding replacing the middle-mile site-to-site connectivity – which currently consisted of regional and global MPLS providers – with site-to-site SDCI connectivity.

Finally, although satisfied with this design, 4Dachs Consulting also decided to look at the second option for leveraging private connections (AWS Direct Connect, Azure ExpressRoute, etc.) for cloud-to-cloud connectivity between public IaaS/PaaS CSPs via an SDCI partner, to see if it provided any additional benefits.

## Option 2: Extending the SD-WAN Fabric into the SDCI Partner and Public IaaS/PaaS CSP Networks

With SDCI Option 2, 4Dachs Consulting would also extend their Cisco Catalyst SD-WAN fabric into the SDCI partner network via the Internet, through a pair of Catalyst 8000v SD-WAN routers functioning as Interconnect Gateways (ICGWs) for redundancy. However, they would also maintain the existing CGWs (Cisco Catalyst 8000v router pairs) within each of the public IaaS/PaaS CSPs. Instead of extending a Service VPN from the ICGWs through the private connectivity (AWS Direct Connect and Azure ExpressRoute), 4Dachs Consulting would be extending the WAN Transport VPN (VPN 0) to the CGWs within each public IaaS/PaaS CSP. This could be automated through a combination of the Cisco Cloud onRamp for Multi-Cloud and the Cisco Cloud onRamp for Multi-Cloud Interconnect workflows, for ease of deployment.

An example of this design in shown in the following figure.

**Figure 5. Cloud-to-Cloud Connectivity via SDCI Provider – Option 2**

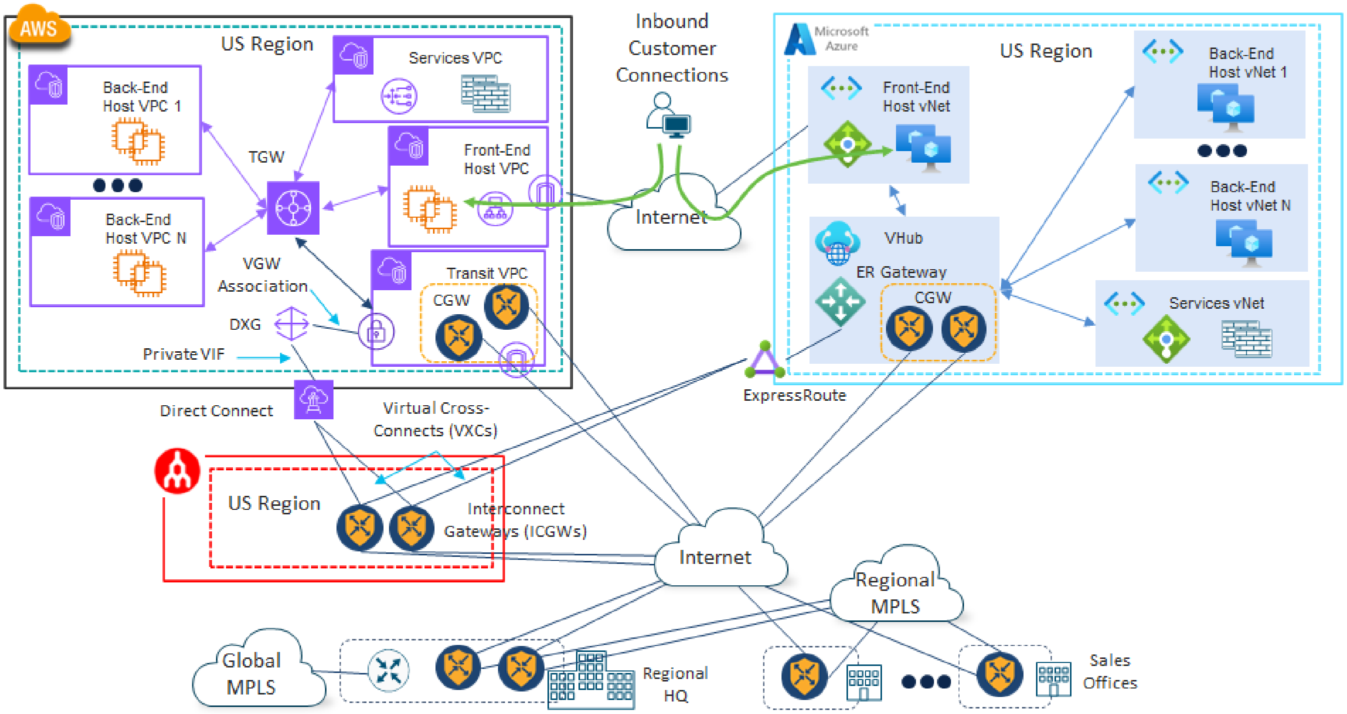

4Dachs Consulting would again leverage the peering relationships between the SDCI partner and the various public IaaS/PaaS CSPs to provide private connectivity via AWS Direct Connect, Azure ExpressRoute, etc., from the ICGWs through Virtual Cross-Connects (VXCs).

Since 4Dachs Consulting would be keeping existing CGWs within the public IaaS/PaaS CSPs with this design there would be no cost savings from removing them as with the previous option. However, 4Dachs Consulting could now leverage the Internet as a backup path (although the latency may not be within their requirements for customer satisfaction) and therefore, not rely completely on the built-in redundancy within the private connectivity to each of the IaaS/PaaS CSPs for high availability.

### AWS Connectivity Details

As with SDCI Option 1, 4Dachs Consulting was already using Transit Gateways (TGWs) within AWS in each region and wanted to continue with the same design. However, in this case within the SDCI provider network, each Virtual Cross-Connect (VXC) would need to be provisioned as a **Connection VIF Type - Private** from each of the ICGWs to AWS to support connection to a Direct Connect Gateway (DXG). The Direct Connect Gateway (DXG) would then be associated with a VPN Gateway (VGW) within the Transit VPC. Within AWS, **Hosted Connections** would again need to be provisioned with a bandwidth of 1 Gbps.

### Azure Connectivity Details

As with SDCI Option 1, within Azure, 4Dachs Consulting wanted to continue using the vWAN architecture, with a vHub in each region. Hence, within the SDCI provider network, each VXC would be associated with an ExpressRoute circuit that terminates on an ExpressRoute Gateway (ER Gateway) within a vHub in the Azure region. 4Dachs Consulting would begin with 1 Gbps of bandwidth on each VXC. Again, since each regional site (U.S. and Europe) is essentially independent of each other, a **Standard** ExpressRoute circuit is all that would be required.

### Additional Considerations and Benefits

In this design, since the Cisco Catalyst SD-WAN fabric is extended into the public IaaS/PaaS CSPs via the CGW, Cisco Overlay Management Protocol (OMP) automatically routes prefixes between the CSPs within secure SD-WAN tunnels. Hence any concerns regarding the maximum number of prefixes that could be sent over the private connections to each of the public IaaS/PaaS CSPs are alleviated.

Within AWS, BGP peering would be established between the TGW and the Cisco Catalyst 8000vs within the CGW. Hence, the maximum number of prefixes that could be sent to a TGW applied. Likewise, within Azure, BGP peering would be established between the vHub and the Cisco Catalyst 8000vs within the CGW. Hence, the maximum number of prefixes that could be sent to a vHub applied. However, 4Dachs Consulting did not believe that these limitations would be an issue for their network.

| **Technical Note:** |
|---|
| Configuring of the BGP routing between the Catalyst 8000vs within the Cloud Gateways (CGWs) and the constructs within the public IaaS/PaaS CSP (Azure vHub, AWS TGW, etc.) can be automated as part of the Intent Management section of the Cisco Cloud onRamp for Multi-Cloud workflow if desired, or manually and/or custom automation - depending upon whether the insertion of a security appliance is needed between the host VPCs/vNets. |

Another benefit of this design is that cloud-to-cloud traffic between public IaaS/PaaS CSPs would be encrypted automatically at the network level. However, again there were no regulatory requirements that applied to

4Dachs Consulting requiring network-level encryption. If application-level encryption of the back-end services was maintained via secure protocols, network-level encryption was not a concern to 4Dachs Consulting.

To 4Dachs Consulting, the potential benefits of this design over the previous SDCI design options were as follows:

- It provides a secondary path via the Internet between the CGWs within the public IaaS/PaaS CSP, for additional high availability. However, the private connectivity constructs (AWS Direct Connect and Azure ExpressRoute) within each public IaaS/PaaS CSP were already designed for high availability. The provisioning of a pair of ICGWs, each with a single VXC within the SDCI provider network further increases high availability. Additionally, 4Dachs Consulting had previously identified a concern that cloud-to-cloud connectivity between different public IaaS/PaaS CSPs via the Internet would result in high latency which affected the quality of experience of their customer-facing applications (See **Appendix D: Alternative Solutions**). Hence the backup Internet path may only be marginally useful for them.
- Network-level encryption of cloud-to-cloud traffic between different public IaaS/PaaS CSPs. As mentioned previously, there were no regulatory requirements that applied to 4Dachs Consulting for network-level encryption. Further, 4Dachs Consulting felt that this could be accomplished for their deployment at the application level via secure back-end protocols.
- Since this design option extended the SD-WAN fabric into the IaaS/PaaS CSP, 4Dachs Consulting could extend multiple service VPNs between different CSPs – depending upon the CSP. However, with the Azure vWAN/vHub architecture the CGWs are mapped to a single default route table within the vHub, limiting the extension of the Cisco Catalyst SD-WAN into Azure to a single Service VPN. With AWS, multiple Service VPNs could be extended to the TGW via different route tables. However, as mentioned previously, 4Dachs Consulting's current requirements were for extending only a single service VPN between different public IaaS/PaaS CSPs.

4Dachs Consulting spent time evaluating whether the benefits identified above provided any clear advantage over the previous SDCI design option. They concluded that the previous design option would fully meet their requirements and decided to pursue implementing SDCI design Option 1. However, if design requirements were to change, they would keep SDCI design Option 2 in mind.

## Summary

This case study presented the use case of a fictional customer, 4Dachs Consulting, whose primary business challenge was providing cloud-to-cloud connectivity between different public IaaS/PaaS CSPs with SLA guarantees of bandwidth, low latency, and low packet loss. Such guarantees were necessary to support the back-end application data flows necessary for 4Dachs Consulting to provide software services to their customers through SaaS offerings as they expanded into multiple IaaS/PaaS CSPs.

4Dachs Consulting evaluated multiple options for cloud-to-cloud connectivity between different public IaaS/PaaS CSPs, including using an SDCI partner, the Internet, and regional MPLS providers.

Although use of the Internet for connectivity between different public IaaS/PaaS CSPs was already possible through their existing Cisco Catalyst SD-WAN deployment, consisting of Cloud Gateways (CGWs) deployed within each CSP, latency over the Internet was a concern. The quality of experience of their SaaS offerings depended on back-end flows that could traverse different CSPs. Higher latency would result in a less responsive experience for their customers.

The use of regional MPLS providers between different public IaaS/PaaS CSPs, again leveraging their existing Cisco Catalyst SD-WAN deployment consisting of Cloud Gateways (CGWs) deployed within each CSP, could provide the necessary SLA guarantees.  However, 4Dachs Consulting had concerns regarding business agility – specifically the ability to respond to requirements for increased or decreased bandwidth or add and delete different IaaS/PaaS CSPs as necessary – due to longer-term contracts negotiated with regional MPLS providers.

Finally, the use of an SDCI provider between different public IaaS/PaaS CSPs, was considered.  4Dachs Consulting considered both options – leveraging Interconnect Gateways (ICGWs) within the SDCI provider and extending the Cisco Catalyst SD-WAN fabric into the CSPs via the existing CGWs; and leveraging ICGWs within the SDCI provider without extending the Cisco Catalyst SD-WAN fabric into the CSPs.  Based upon the need to extend only a single Service VPN into the CSPs, and no requirements for network-level encryption of the data flows between CSPs, 4Dachs Consulting chose to eliminate the existing CGWs within each of the IaaS/PaaS CSPs, and move to a design in which a Service VPN was extended into each CSP via private connectivity (AWS Direct Connect and Azure ExpressRoute) through Virtual Cross Connects (VXCs) from the ICGWs within the SDCI provider.  This met the SLA guarantees of low latency and low data loss between different IaaS/PaaS CSPs and provided them the business agility to adjust bandwidth as needed or add and delete different IaaS/PaaS CSPs as necessary to ensure the quality of experience of the SaaS software service they provided to their customers.

## Appendix A:  Changes from Previous Versions

This guide is a new guide.  There are no previous versions.

## Appendix B: Software Version

This guide is based upon Cisco SD-WAN software version 17.9/20.9.

# Appendix C:  Alternative Solutions

The following sections discuss alternative solutions which 4Dachs Consulting considered for public IaaS/PaaS CSP connectivity.  Although each presents a valid method of connecting to multiple public IaaS/PaaS CSPs using the Cisco Catalyst SD-WAN fabric, 4Dachs Consulting ultimately chose not to pursue the alternatives, based on their specific business requirements.

## Cloud-to-Cloud Connectivity via the Cisco Catalyst SD-WAN Fabric Using the Internet

For 4Dachs Consulting, the most obvious way to connect different public IaaS/PaaS CSPs was via the Internet. 4Dachs Consulting thought about leveraging the native IPsec VPN connectivity functionality of each public IaaS/PaaS CSP to establish secure connectivity.  However, this design was quickly dismissed, since 4Dachs Consulting already had the Cisco Catalyst SD-WAN fabric extended into each public IaaS/PaaS CSP through Cloud Gateways (CGWs).

Although, connectivity was primarily for site-to-cloud traffic for monitoring and maintenance of the application services deployed within the VMs/containers within the CSP, the Cisco Catalyst SD-WAN fabric could easily be leveraged for secure cloud-to-cloud connectivity between different public IaaS/PaaS CSPs via the Internet as well.  This is shown in the following figure for the U.S geographical area.

**Figure 6.  Cloud-to-Cloud Connectivity Between IaaS/PaaS CSPs Utilizing the Cisco Catalyst SD-WAN Fabric with the Internet**

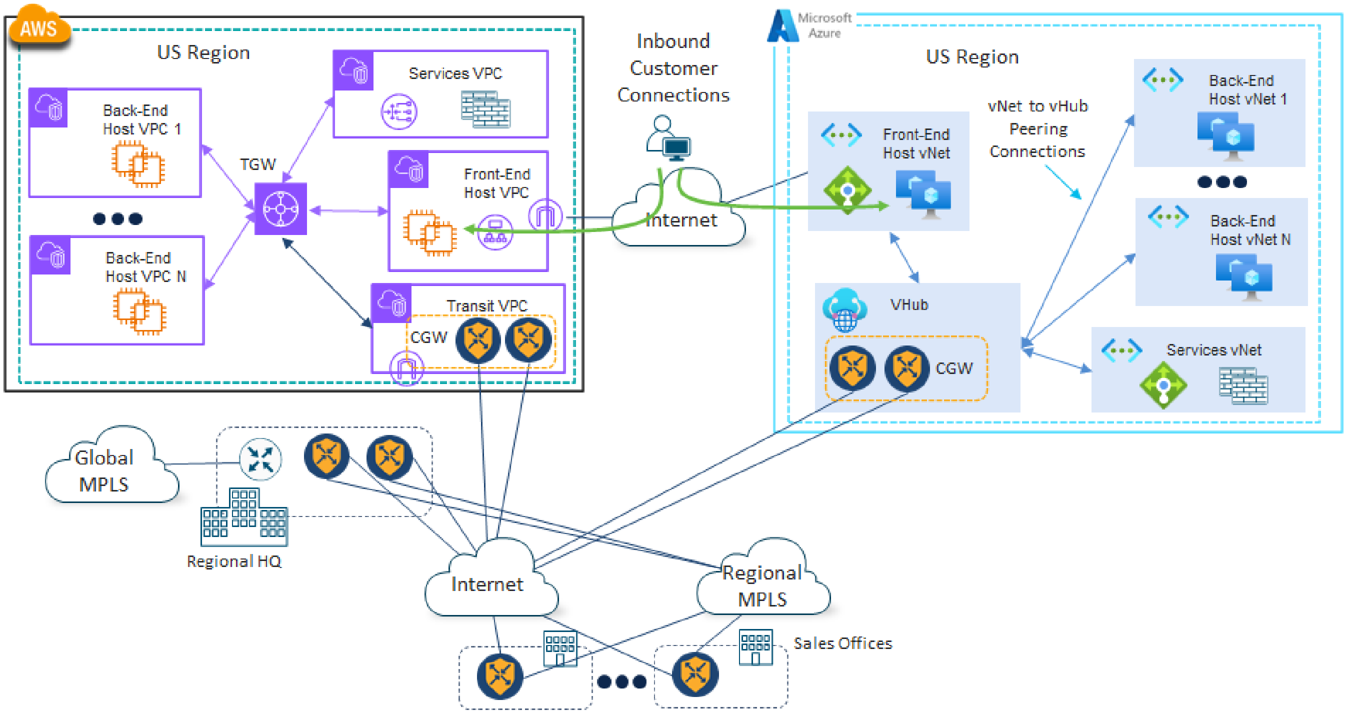

The primary advantage of this design is that 4Dachs Consulting would have to do minimal work to enable the cloud-to-cloud connectivity.  By default, the Cisco Catalyst SD-WAN fabric forms secure full-mesh connectivity between sites which have Catalyst SD-WAN Edge devices – minimizing the network latency between sites and guaranteeing encryption all the way into each public IaaS/PaaS CSP's network.  Alternatively, the Cisco Catalyst SD-WAN fabric provides the flexibility such that 4Dachs Consulting could implement centralized control policy to ensure all site-to-cloud and cloud-to-cloud traffic was sent through a regional corporate site. However, this would increase the latency for cloud-to-cloud traffic, which was still the primary concern for 4Dachs Consulting since the Internet was being used as the transport.

**Technical Note:**

The Azure vWAN architecture maps the Service VPN side of the Cisco Catalyst 8000v routers (instantiated as NVAs within the vHub) to the default route table of the vHub. Hence only a single Service VPN can be extended across the Cisco Catalyst SD-WAN fabric to each vHub. For more information on this design please see the following guide:

https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/Cisco_Cloud_onRamp_for_Multi-Cloud_Azure_Version2.html

Note that multiple vHubs can be implemented within an Azure region – each with a different Service VPN mapped to the default route table of the vHub.

4Dachs Consulting realized that there would be no SLA guaranteeing bandwidth, latency, or packet loss of the cloud-to-cloud traffic with this design. Increased latency between the public IaaS/PaaS CSPs via the Internet would adversely affect the user experience of their customers. Their microservice application architecture required data flows between multiple back-end services – possibly located in different public IaaS/PaaS CSPs – to generate a response which could be sent back to the customer via the web-based front-end. Another concern for 4Dachs Consulting was data transfer charges since traffic between different public IaaS/PaaS CSPs was exiting the region via the Internet. Hence, outbound data transfer charges from the CSP regions applied for the cloud-to-cloud traffic.

Primarily due to not having any SLA guaranteeing bandwidth, latency, and packet loss, when using the Internet for cloud-to-cloud connectivity, 4Dachs Consulting decided to explore further design options which leveraged their regional MPLS carriers.

## Cloud-to-Cloud Connectivity via a Regional MPLS Carrier

With this design, 4Dachs Consulting would leverage the private connectivity options of the public IaaS/PaaS CSPs – such as AWS Direct Connect and Azure ExpressRoute – which would be provided through regional MPLS carriers. This assumed the regional MPLS carriers had the necessary peering relationships with the public IaaS/PaaS CSPs (AWS, Azure, etc.), to provide such services.

4Dachs Consulting realized they had several potential design options when adding private connectivity to public IaaS/PaaS CSPs via MPLS carriers, some of which depended on the offerings of the MPLS providers themselves. For example, AWS supports Direct Connect connections of type Transit Virtual Interface (Transit VIF) to a Direct Connect Gateway (DXG), which then forms an association with the Transit Gateway (TGW) within the region, as shown in the following figure.

**Figure 7. AWS via Direct Connect using a Transit Gateway (Dependent Upon Private WAN Offering)**

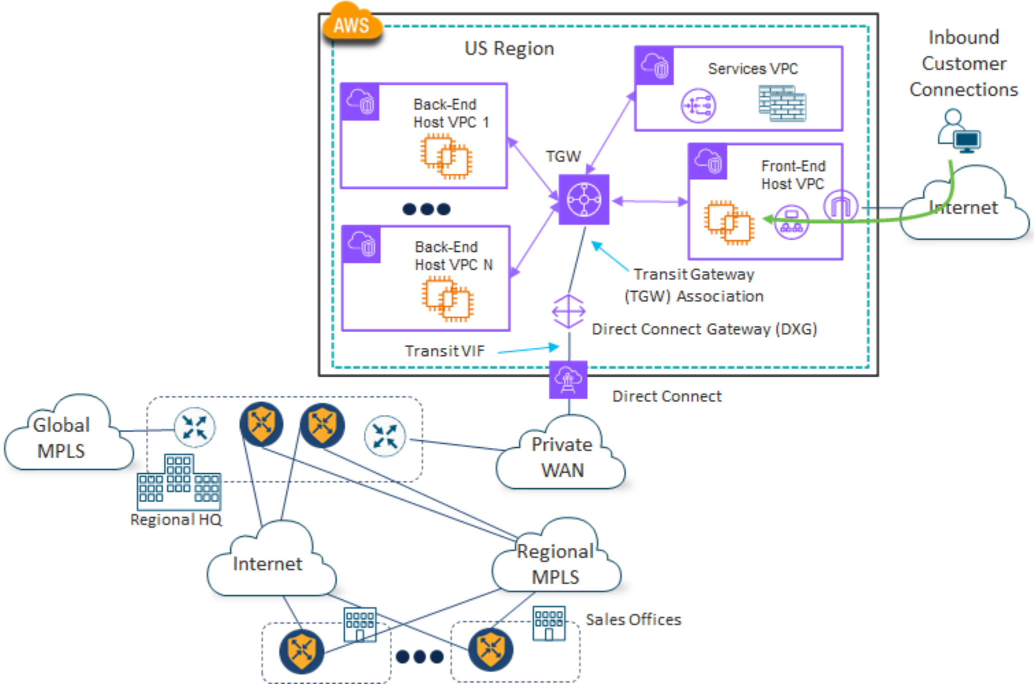

Since 4Dachs Consulting was already using TGWs in each AWS region (U.S. and Europe), this would be the preferred method of bringing in MPLS connectivity via a Direct Connect circuit. However, this design is somewhat dependent upon the private WAN offering. Depending on the MPLS provider, they may only support Private VIF connections, requiring a separate VIF to be provisioned across the Direct Connect circuit for each VPC connected via a VPN Gateway (VGW) with or without a Direct Connect Gateway (DXG). An example of this is shown on the left side of the figure below.

**Figure 8. AWS via Direct Connect via MPLS (Possible Provider Options)**

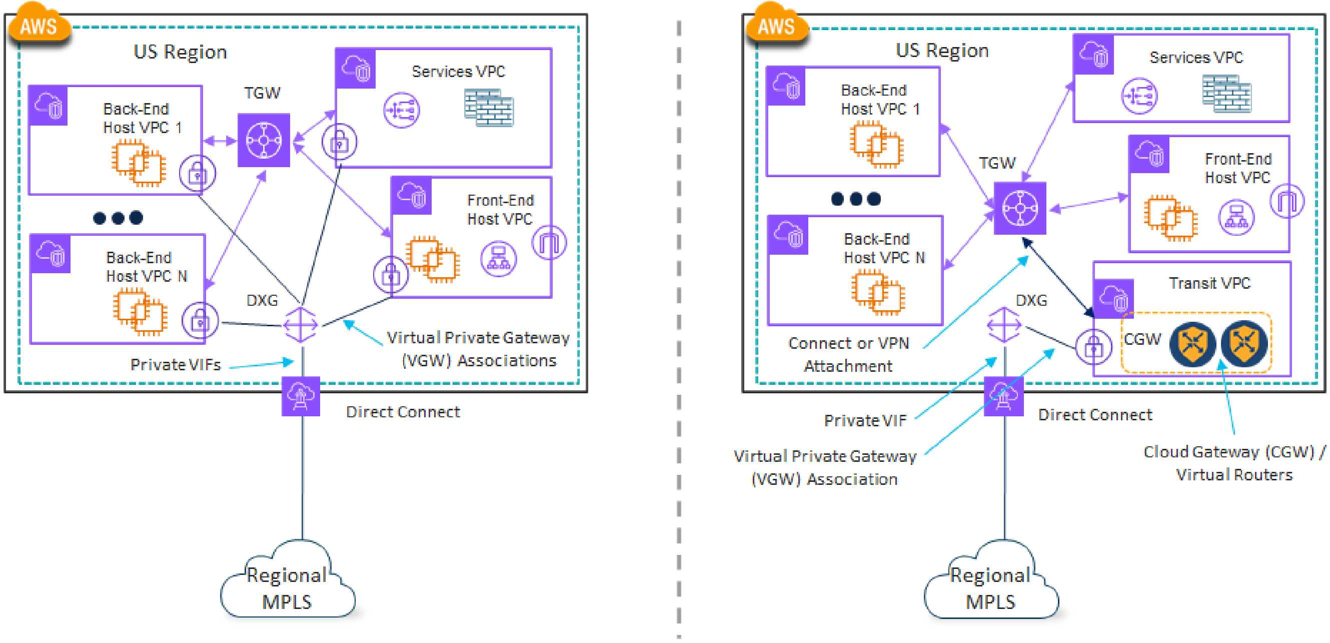

When looking at the design on the left side of the figure above, a further complication is that VPC-to-VPC connectivity via the Direct Connect circuit is generally not supported unless the traffic is backhauled to a corporate site. Therefore, the TGW may be needed for direct VPC-to-VPC connectivity.

Alternatively, the MPLS provider may support a Private VIF connection to a VPN Gateway (VGW) within a Transit VPC, again with or without the Direct Connect Gateway (DXG). This is shown on the right side of the figure above. This design relies on third party virtual devices such as the Cisco Catalyst 8000v routers operating as a Cloud Gateway (CGW) within the Transit VPC to provide dynamic routing for reachability between the Host VPCs behind the Transit VPC and the corporate sites. The design on the right in the figure above is somewhat simpler in that VPC-to-VPC connectivity within the region is via the TGW and all VPC connectivity outside the region is via the Transit VPC.

| Technical Note: |
|---|
| Actual limitations regarding what designs are supported and not supported when provisioning private connectivity (AWS Direct Connect, Azure ExpressRoute, etc.) to public IaaS/PaaS CSPs through an MPLS provider depend on both the MPLS provider and the CSP – and are not covered in this guide. The limitations brought up in the example above are for illustration purposes only. The reader is encouraged to thoroughly investigate supported designs by both the public IaaS/PaaS CSP and the MPLS provider before making any design decisions. |

Assuming the MPLS provider supports the ability to bring in the Direct Connect connection through a Transit VIF to a DXG, which then forms an association with the TGW within the region, the most basic design for 4Dachs Consulting would be to leave the Cisco Catalyst SD-WAN deployment as is, as shown in the following figure.

**Figure 9. Cloud-to-Cloud Connectivity via Regional MPLS Carrier in Parallel with the Cisco Cloud Gateways (CGWs)**

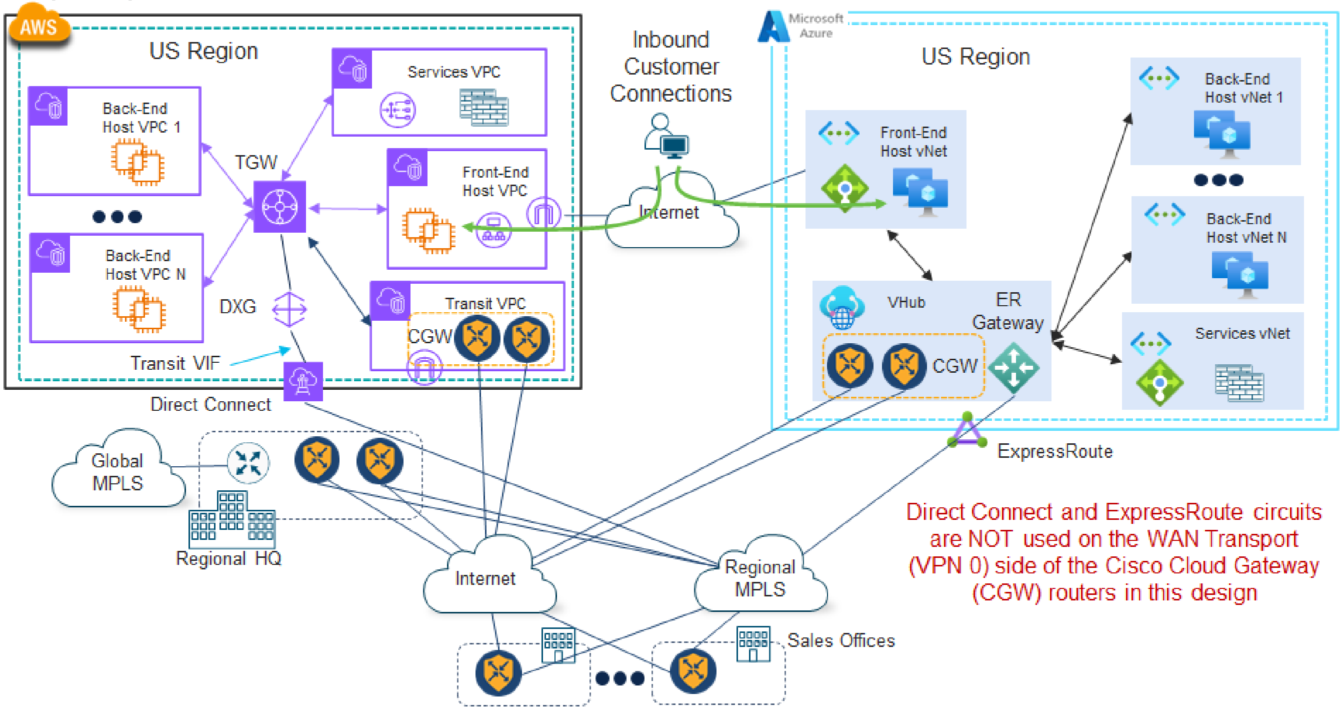

Within AWS, a Direct Connect circuit connects through a Transit VIF to a DXG, which then forms an association with the TGW within the region.  Within Azure, an ExpressRoute circuit, provided through the regional MPLS provider, terminates on an ExpressRoute Gateway (ER Gateway) within the regional vHub.  The ER Gateway is associated with the default route table of the vHub.

With this design, the AWS Direct Connect and the Azure ExpressRoute circuits are NOT used as WAN transports (VPN 0) for the Cisco Cloud Gateway (CGW) within the AWS Transit VPC or the Azure vHub.  Essentially the AWS Direct Connect and ExpressRoute circuits sit in parallel to the Cisco Catalyst SD-WAN deployment. 4Dachs Consulting realized that this could complicate the routing within the overall deployment, since the subnets within the public IaaS/PaaS CSPs are seen by the Service VPN side of the Cisco Catalyst 8000v routers and transported over the SD-WAN fabric, as well as being directly transported across the regional MPLS provider.  4Dachs Consulting realized that complexity of routing could lead to mistakes in configuration that could result in outages that affect their customers.

Note that for 4Dachs Consulting, the regional MPLS carrier transports were currently used primarily to provide a secondary SD-WAN transport (Internet and regional MPLS) between corporate sites, to leverage Application Aware Routing (AAR) for the best path between corporate sites within each of the geographic regions (U.S. and Europe).  Eliminating the regional MPLS as a transport between corporate sites – to potentially simplify the routing – was not an option.

4Dachs Consulting also had to consider that the traffic between the public IaaS/PaaS CSPs would be unencrypted as it traversed the regional MPLS provider with this design.

Hoping to simplify the overall routing of the deployment, 4Dachs Consulting decided to investigate a design in which the AWS Direct Connect and Azure ExpressRoute circuits could also be used as WAN transports (VPN 0) for the Cisco Catalyst 8000v instances functioning as Cloud Gateways (CGWs) within the AWS Transit VPC and the Azure vHub.  This is shown in the following figure.

Figure 10. Cloud-to-Cloud Connectivity via Regional MPLS Carrier as Cisco Catalyst SD-WAN Transport

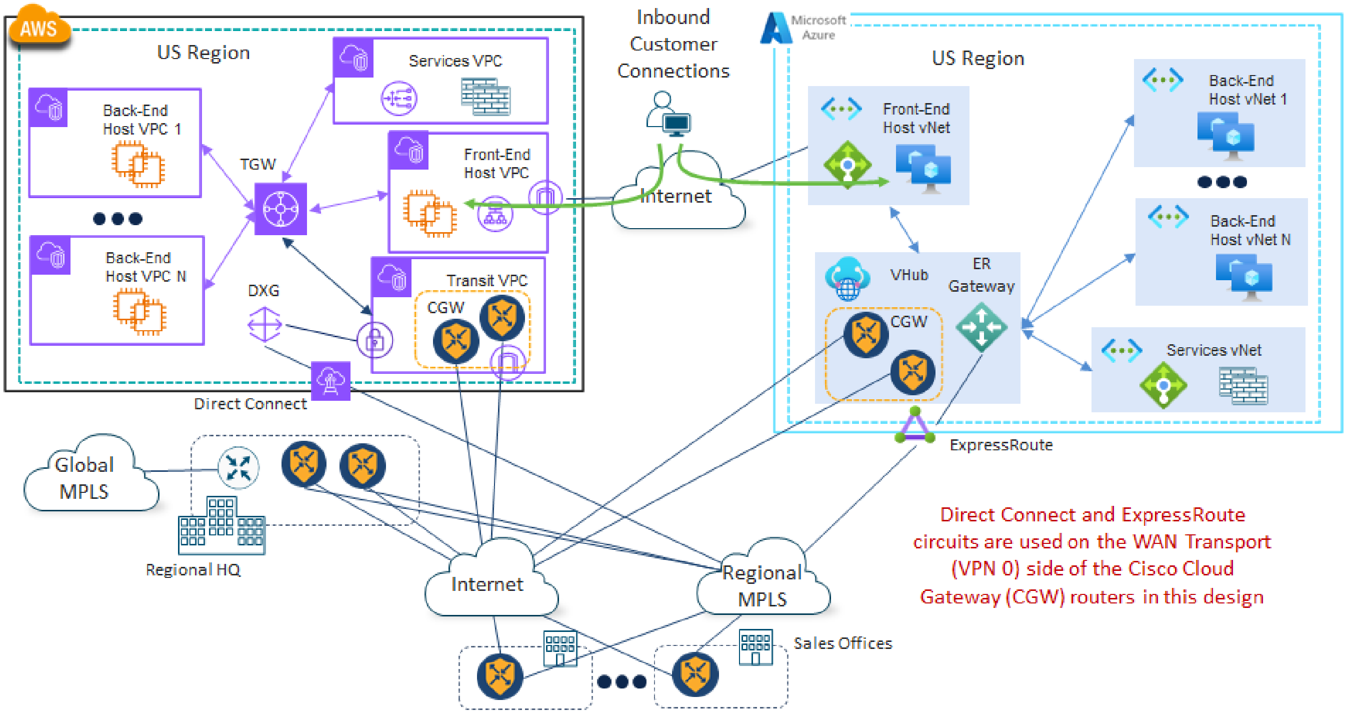

In the figure above, the AWS Direct Connect circuit connects through a Private Virtual Interface (VIF). The Private VIF can either connect directly, or via a Direct Connect Gateway (DXG), to the VPN Gateway (VGW) within the Transit VPC – instead of the individual host VPCs. Within Azure, the ExpressRoute circuit, provided through the regional MPLS provider, again terminates on an ExpressRoute Gateway (ER Gateway) within the regional vHub.

With this design the AWS Direct Connect and Azure ExpressRoute circuits are configured to be used as WAN transports (VPN 0) of the Cisco Cloud Gateway (Catalyst 8000v instances) – along with the existing Internet connectivity.

**Technical Note:**

The **Alternative Azure Designs** section of the **Extending the Cisco SD-WAN Fabric into Azure with Cisco Cloud onRamp for Multi-Cloud** deployment guide, located at the following URL, discusses how Azure ExpressRoute circuits can be accommodated as Cisco Catalyst SD-WAN transports using Cisco Cloud onRamp for Multi-Cloud. A similar design can be implemented for AWS.

https://www.cisco.com/c/en/us/td/docs/solutions/CVD/SDWAN/Cisco_Cloud_onRamp_for_Multi-Cloud_Azure_Version2.html#_Toc103241865

This design provides the ability to leverage the regional MPLS carrier for cloud-to-cloud traffic between different public IaaS/PaaS CSPs via the Cisco Catalyst SD-WAN fabric. Hence, 4Dachs Consulting could leverage the SLA guarantees of latency, bandwidth, and packet loss of the regional MPLS carrier, to better ensure the end-user experience of their web-based software services to their customers – as they transitioned to a multi-cloud design. Also, cloud-to-cloud traffic between different public IaaS/PaaS CSPs would be encrypted.

However, leveraging the regional MPLS carriers had some downsides that 4Dachs Consulting also had to consider.  First, their MPLS contracts were typically negotiated long-term with a fixed amount of bandwidth. Although there were some advantages, in that the cost was also fixed over the term of the contract, 4Dachs Consulting also felt that this could decrease their business agility.  Should they require additional bandwidth due to increased customer demand, new or updated software requirements of their customer-facing applications, or the deployment of additional customer-facing applications – they would have limited ability to respond to the requirements.  On the other hand, should the business environment change such that they did not require as much bandwidth, they would be locked into paying for the bandwidth – at a premium price due to the SLA guarantees of the MPLS provider – although the bandwidth would be under-utilized.

Primarily due to their concerns around losing business agility with an MPLS provider, 4Dachs Consulting decided instead to pursue design options which leveraged Software-Defined Cloud Interconnect (SDCI) partners.

## Appendix D: Glossary

| | |
|---|---|
| **AAR** | Application Aware Routing |
| **AWS** | Amazon Web Services |
| **AZ** | Availability Zone |
| **BGP** | Border Gateway Protocol |
| **CSP** | Cloud Service Provider |
| **CGW** | Cloud Gateway |
| **DXG** | Direct Connect Gateway |
| **ER Gateway** | Express Route Gateway |
| **GCP** | Google Cloud Platform |
| **IaaS** | Infrastructure-as-a-Service |
| **ICGW** | Interconnect Gateway |
| **MVE** | Megaport Virtual Edge |
| **NVA** | Network Virtual Appliance |
| **OCI** | Oracle Cloud Infrastructure |
| **OMP** | Overlay Management Protocol |
| **PaaS** | Platform-as-a-Service |
| **REST** | Representational State Transfer |
| **SDCI** | Software-Defined Cloud Interconnect |
| **SLA** | Service Level Agreement |
| **TGW** | Transit Gateway |
| **VGW** | VPN Gateway |
| **vHub** | Virtual Hub |
| **VIF** | Virtual Interface |
| **VM** | Virtual Machine |
| **vNet** | Virtual Network |
| **VPC** | Virtual Private Cloud |
| **vWAN** | Virtual WAN |
| **VXC** | Virtual Cross-Connect |
| **WAN** | Wide Area Network |

## Feedback

For comments and suggestions about this guide and related guides, join the discussion on Cisco Community at https://cs.co/en-cvds.

**Americas Headquarters**
Cisco Systems, Inc.
San Jose, CA

**Asia Pacific Headquarters**
Cisco Systems (USA) Pte. Ltd.
Singapore

**Europe Headquarters**
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.