AWS Direct Connect with AWS Site-To-Site VPN as a failover

1 April 2022 - 7 min. read

Advanced Networking *AWS Direct Connect*

Introduction

Today, many solutions require approaches that implement a joint use of public cloud providers and their own on-prem resources. In this series of articles, we provide an overview of useful AWS services to build a hybrid network that seamlessly extends workloads from local installations to the public cloud according to business needs.

In some scenarios, hybrid cloud networks rely on **AWS Direct Connect (DX)**. This service provides another option rather than the public internet to connect to AWS by delivering data through a **private network connection between** the **on-premise facility** and **the AWS cloud**.

If you like to know more about Direct Connect, please check our previous blog post. It explains what it is and how to choose it regarding some specific needs.

In some cases, this connection alone is not enough. It is always better to guarantee a **fallback connection** as the backup of DX. There are several options, but implementing it with an **AWS Site-To-Site VPN** is a real cost-effective solution that can be exploited to reduce costs or, in the meantime, wait for the setup of a second DX.

Please bear in mind that this solution does not ensure an SLA, and it is impossible to obtain it with a connection over the public internet.

High-Level Architecture

This article will progress on what has already been put in place in the previous article (Hybrid Cloud Networking: centralized NAT Gateway through AWS Transit Gateway). We presented how to set up a Transit Gateway and share network appliances between AWS Accounts. So, please, it might be helpful to read it first to fully understand what will be explained here.

Let's assume that a fictitious company currently wants to connect its own on-premises facilities with the AWS Accounts directly to a DX connection by using a VPN Site-To-Site as a backup.

Before proceeding, it is required that a Direct Connect has been already requested and ordered. Plus, the whole AWS networking has been centralized through Transit Gateway.

As shown below, the connection between the local facility and AWS through the Direct Connect private connection can be configured by directly attaching the DX to the AWS Transit Gateway.

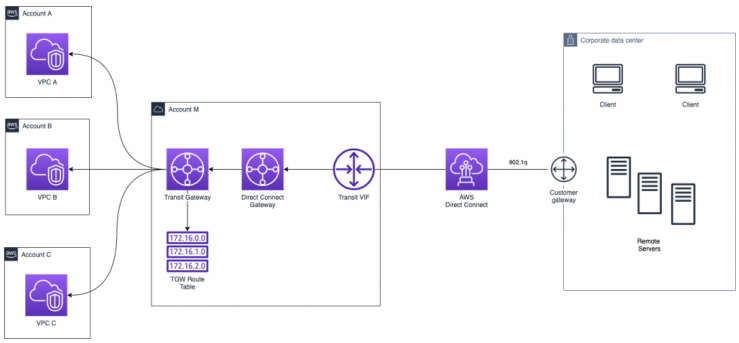

Let's dive in:

Direct Connect offers a location over a standard Ethernet fiber-optic cable. One end is connected to the on-prem router, and the other to the AWS DX router. A user can associate it directly with the Transit Gateway through the **Transit Virtual Interface (VIF)**. This specific interface is different from the private/public interface because it is linked directly with the Transit Gateway. This enables connectivity to all VPCs that share an attachment.

Be careful! Please note that it is possible to attach only 1 transit virtual interface for DX dedicated connection. This limit cannot be increased.

After the setup has been done, the company wants to use the **site-to-site VPN** as a backup. It consists of an encrypted link, called VPN tunnel, that connects the on-prem site with AWS. Each VPN connection includes two VPN tunnels which can simultaneously be used for *high availability*.

In our case, the two endpoints will be the Customer Gateway (CGW) for the on-prem side and the Transit Gateway for the AWS one.

Please note that VPN has a couple of limits that must be taken into account in this design:

- Maximum 50 Site-to-Site VPN connections per AWS Region.
- Maximum bandwidth per VPN tunnel, up to 1.25 Gbps.

The modified scenario will look like this:

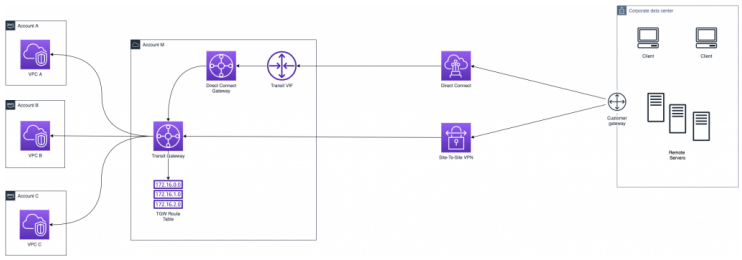

With this improved configuration, the traffic will flow through Direct Connect and, when a **failure happens** on the **DX**, the traffic will be switched to the failover **VPN Channel**. The purpose of using DX as the **primary active path** is that it guarantees:

1. High speed up to 100 Gbps (not in every Region)
2. Bypassing public internet and so network congestion and unpredictability.

The entire routing behavior and the failover mechanism is managed by the BGP protocol, and it's a constraint for the customer gateway router to support it. The router will send various keepalive packets to check if the DX path is active, and if a failure is detected, the path is switched to the fallback one (VPN). In fact, BGP dynamic routing is directly controlled by the Customer Gateway. The CGW, whether software or a physical appliance, must be configured appropriately by the network administrator to exchange routing information among the various routers. It works by managing a table of IPs (called prefixes) that provides information on the reachability of the different networks we're observing. Then, BGP is responsible for exchanging IP blocks

advertisement (IP prefixes) to the various systems.
But if BGP is advertising prefixes, how are the **priorities** of the routes managed by AWS? Well, to understand it better, let's shift the focus on the Transit Gateway and resume them in a table:

Order	Description
1	**Local routes** to the VPC
2	The **most specific prefix**. *i.e., 10.0.1.0/24 is preferred rather than 10.0.0.0/16*
3	**Static route entries** preferred **over dynamic ones**
4	**Dynamic Routes** over **BGP**
5	Routes learned via Static VPN
6	**BGP routes from the VPN** - Based on shorter AS path

Hence, to have our infrastructure in place, it is necessary to **advertise the same prefix** for both the **Direct Connect (VIF)** and the **VPN**, because if the CGW is advertising the same routes toward the AWS VPC, the Direct Connect path is always preferred, regardless of AS path prepending.

Please, by considering the table's point 3, be careful when configuring the VPN Site-To-Site. Inside the VPN wizard form, a user must specify **"Dynamic"** as **Routing Option,** otherwise all the reasoning would fall.

Sometimes, a scenario of asymmetric routing might occur with this kind of configuration. **Asymmetric routing** means that a package traverses from a source to a destination in one path and takes a different path when it returns to the source. *i.e., traffic from AWS to on-premise might flow through the Direct Connect link, but for the opposite, the traffic might traverse the VPN tunnel instead.*

To solve this issue, please consider this AWS troubleshooting guide.

Hybrid Cloud

Now that the overall networking has been set up, we will share what is achievable in exploiting these technologies:

- **Predictable latency and throughput:** With AWS DX, it is possible to choose the right bandwidth. Moreover, the Transit Virtual Interface (VIF) supports jumbo frames (8500 bytes), improving throughput.

So, let's assume that you are building an HPC (High-Performance Computing) solution with EC2 instances that will operate on large datasets that must be transferred between the on-prem facility and the AWS Cloud.

To speed up the data transfer, it is possible to leverage jumbo frames to send data because more data is transferred in fewer packets.

Moreover, they are also supported natively by the infrastructure, including the EC2 instances! (only the current generation, though).

- **Add encryption layer in DX**: What if there is a need for strict security compliances and to provide encryption in the Direct Connect link? It is possible to combine a VPN over Direct Connect.

If there's the plan to use the VPN Site-to-Site managed by AWS, it will cap the bandwidth because of the limit of 1.25 Gbps for the VPN Tunnel. Instead, it is feasible to set up a dedicated IPSec VPN and take advantage of the Direct Connect 100 Gbps Dedicated Connection using a high-end router. Slow down there because it will be very costly... Only top tier appliances support such high bandwidth and they're not so affordable.

To set up this approach:

1. A Direct Connect Public Virtual Interface must be created first.
2. Then, a VPN should be configured as depicted in the image below.
3. A BGP connection is established between the AWS Direct Connect and the Customer Gateway on the public VIF.
4. A BGP session or a static router will be established between the AWS Transit Gateway and the router on the IPSec VPN tunnel.

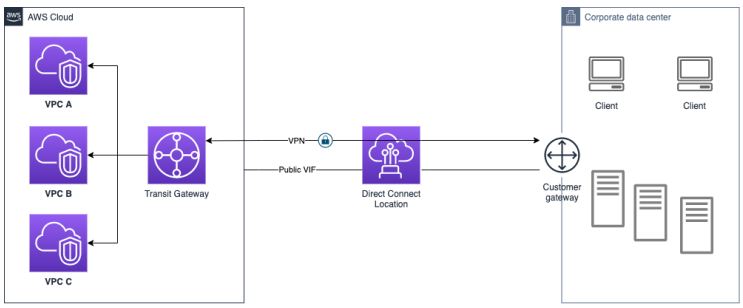

This infrastructure will simplify management and minimize the cost of IPSec VPN connections to multiple Amazon VPCs exploiting the Transit Gateway and, a private dedicated connection is ensured over an internet-based VPN.

- **Support for Inter-Region Connectivity:** Transit Gateway supports inter-Region peering. This means that Direct Connect Gateways attached (also VPN Site-to-Site) to a Transit Gateway, hosted in a specific AWS Region, can exchange traffic with other regions' resources. The whole traffic will traverse the **AWS Global Network**, and it is not exposed on the public internet. The AWS Global network connects all the various Availability Zones with high-bandwidth, low-latency networking over fully redundant, dedicated metro fiber.

Considerations

In this article, we have overviewed how it is possible to add a VPN site-to-site as a fallback method for a Direct Connect connection that links on-premise facilities to the AWS Cloud and how you can benefit from a fully redundant hybrid networking infrastructure.

We remark that a VPN is not designed to provide the same level of bandwidth available to most Direct Connect connections, as it relies on the public internet, so the performances might be unpredictable and indeterministic. On the other hand, it assures a reliable solution when the budget is limited and or when it is impossible to have a second DX provisioned. It is necessary to have at least two DX in order to obtain a high resiliency level of 99.9%.

Have you ever considered this option for the hybrid cloud? Or to one of the proposed solutions? Let us know what you think about this!

Gabriele Zaccagno

Technical Chatterbox @ beSharp. I deal with Cloud Migrations and Data Analytics. I spend most of my time in front of a computer while listening to the whole Radiohead discography (B-Sides included).