Reference:

FAQs

Category:

Networking and Content Delivery

AWS Transit Gateway

Complete book:

Click Here

Created by:

Ashish Prajapati

What?

- AWS Transit Gateway connects your VPCs and on-premises networks through a central hub. AWS Transit Gateway supports dynamic and static routing between attached Amazon VPCs and VPNs.

Why?

- It acts as a cloud router to simplify your network architecture and puts an end to complex peering relationships.
- With inter-Region peering, everything attached to an AWS Transit Gateway is shared across AWS Regions. This includes VPCs, DNS, Microsoft Active Directory, and IPS/IDS.

When?

- You want to deploy new applications across VPCs without updating massive route tables to create peering relationships.
- You want to host multicast applications such as video conferencing, media, or teleconferencing without redesigning your application or tweaking your on-premises network.

Where?

- AWS Transit Gateway is a regional resource and enables you to attach VPCs and VPN connections (within or across AWS accounts) in the same region and route traffic between them.
- You can peer two transit gateways hosted in the same AWS region or across regions, and route traffic between them.

Who?

- Your transit gateway automatically comes with a default route table. You can segment your network by creating multiple route tables in an AWS Transit Gateway and associate Amazon VPCs and VPNs to them.
- A transit gateway scales elastically based on the volume of network traffic.

How?

- Routing through a transit gateway operates at layer 3, where the packets are sent to a specific next-hop attachment, based on their destination IP addresses. When a packet comes from one attachment, it is routed to another attachment using the route that matches the destination IP address.

How much?

- Charges are determined by two factors:
- AWS Transit Gateway hourly charge - calculated per AWS Transit Gateway attachment.
- Transit Gateway data processing charge – calculated per GB of data processed.