Demystifying the Shape of Traffic in the Cloud

How Cloud Monitoring Differs from Traditional On-Prem Solutions

Introduction

Traditional networks have clear-cut edges and obvious directionality, whereas cloud spaces are made up of many distinct private clouds each with separate deployed resources and unique routing.

Distributed Sensors

AWS flow log sensors are located on network interfaces inside of VPCs, not on the edge of the AWS cloud.

Cloud Flow directionality is relative to the location of the flow log collector.

AWS global services are always external to VPCs and never host cloud sensors.

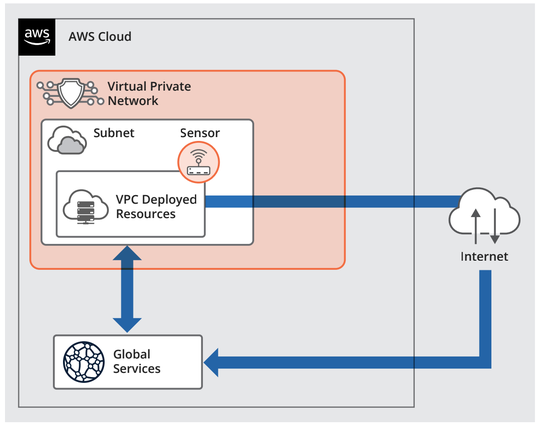

Collection Options

A sensor's Resource ID (location) affects visibility and directionality labeling.

Depending on the resource where the sensor in deployed and the traffic that passes through it, certain fields are more important to analysts.

- pkt-srcaddr vs. srcaddr
- pkt-dstaddr vs. dstaddr

Cloud traffic monitoring requires rethinking typical assumptions of flow labeling, challenging notions of **traffic directionality** and **network scope**.

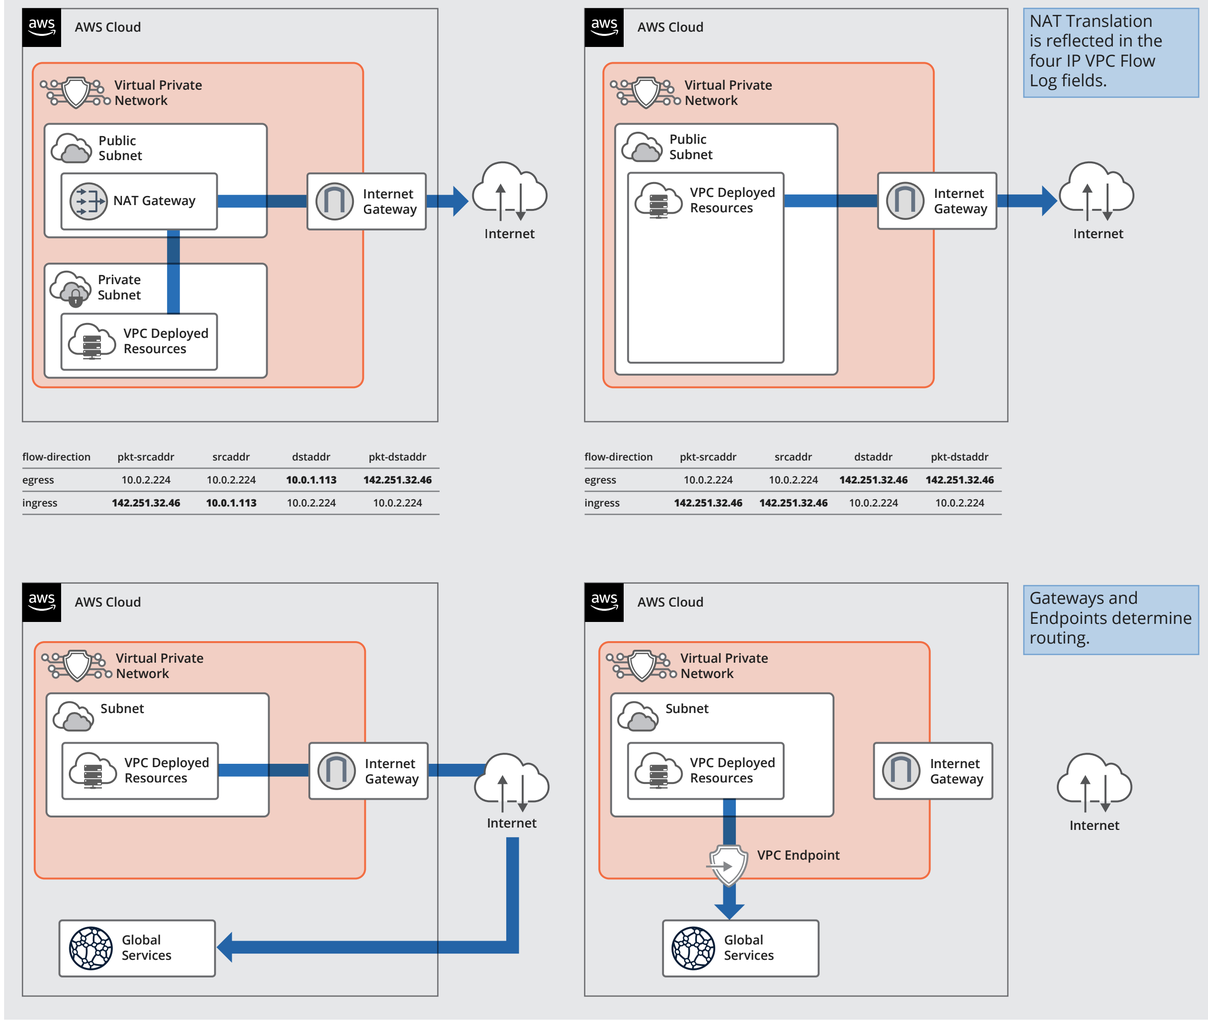

Gateways typically touch the most data, but can leave artifacts in the flows, like reporting duplicate flows and confusing directionality.

Architecture Determines the Shape of Traffic

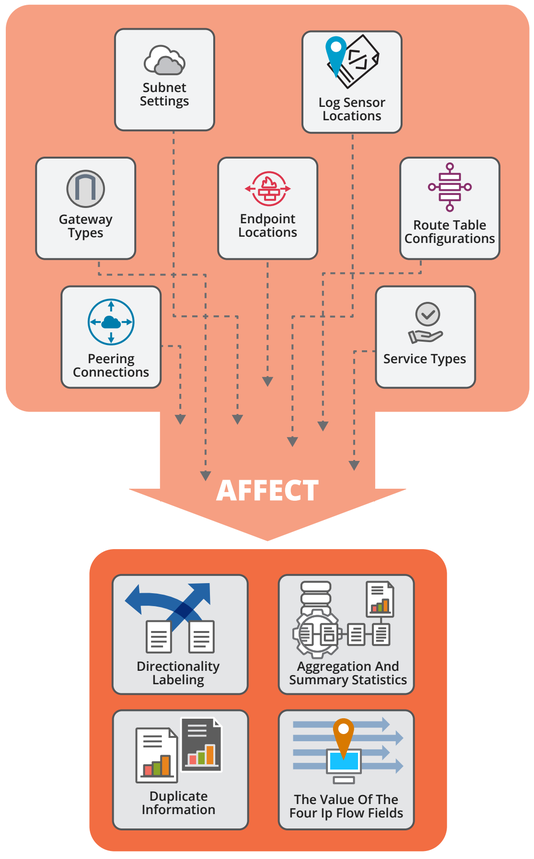

Understanding how architecture elements and settings influence the shape of traffic gives analysts context crucial to deciphering flows.

Exploring Configurations

At CERT, we're exploring different cloud configurations and architecture designs to better understand and document how deployment settings impact traffic routing and monitoring throughout the cloud.

Carnegie Mellon University
Software Engineering Institute

Thomas Benoit | tgbenoit@sei.cmu.edu