emma

Redefining multi-cloud connectivity

While other multi-cloud management platforms offer some networking capabilities for managing cloud environments, the emma platform redefines multi-cloud connectivity. By providing a unified networking fabric, emma simplifies the complexities of managing cloud resources across different providers. It facilitates secure connections, implements network policies, and optimizes traffic routing, enabling organizations to build resilient and highly available network architectures.

This seamless connectivity ensures efficient communication and data transfer between cloud environments.

Benefits

By leveraging the emma own network backbone, organizations can benefit in several ways:

Enhanced performance

Designed with high-speed connections and low-latency pathways, allowing data and traffic to move efficiently between cloud providers and data centers. This reduces the time it takes for data to traverse the network, resulting in improved response times and application performance for end-users.

Increased reliability

Incorporated redundancy and fault-tolerant features allow the network to reroute traffic through alternate paths if one connection or pathway experiences issues. This ensures continuous connectivity and minimizes the risk of service disruptions or downtime.

Seamless multi-cloud integration

The network backbone acts as a unified communication channel between different cloud providers and data centers. It provides a consistent and reliable connection for data transfers, enabling teams to seamlessly integrate resources and services across multiple cloud environments without the complexities of managing disparate networks.

Cost optimization

Data transfer is optimized between cloud providers by leveraging direct connections and peering agreements. This reduces data egress costs as data can travel through the backbone at a cost that is three times lower than the standard rates offered by cloud providers for their direct connect services.

Scalability and flexibility

The scalable architecture enables teams to quickly scale resources up or down as per demand. With direct connections between cloud providers, it becomes easier to allocate additional resources or redistribute workloads in response to changing requirements, ensuring optimal performance and resource utilization.

Security and compliance

Robust security features such as encryption, firewall protection, and access controls, safeguard data as it travels through the network, helping organizations to meet compliance requirements and protect sensitive information.

Centralized monitoring and management

A centralized dashboard provides real-time insights into network performance, usage, and potential bottlenecks. This visibility streamlines the management of network configurations and troubleshooting, making it easier to identify and address any issues promptly.

The emma proprietary network backbone helps customers navigate and eliminate the complexities associated with cloud and multi-cloud environments. By embedding networking capabilities into the emma platform, we can provide a unique cloud-agnostic solution to our customers, allowing them to interact with diverse cloud services without any restrictions or dependencies.

How it works?

We manage two principal categories of network traffic

1

East-West Traffic

This refers to traffic that transits within a particular cloud provider’s infrastructure. The East-West traffic can be categorized as:

- **Layer 2 (L2) Connectivity**
  This encompasses network communication occurring within a singular Virtual Private Cloud (VPC) or Virtual Network (VNet). It entails point-to-point data exchange within the same network segment.
- **Layer 3 (L3) Routing**
  This comprises inter-VPC or inter-VNet routing of data traffic facilitated through VPC/VNet peering mechanism, enabling efficient communication across separate network segments.

Traffic from EC2 to EC2 within AWS Cloud

Contact us
info@emma.ms

Visit our site
emma.ms

emma

2

North-South Traffic

This category involves data traffic transiting in and out of our network through customer edge routers, integrated into the VXLAN (Virtual Extensible LAN) fabric. Our networking fabric infrastructure leverages eVPN/VXLAN as its data plane, ensuring efficient encapsulation of Ethernet frames for tunneling across IP networks. Concurrently, our control plane is built upon the Border Gateway Protocol (BGP) that manages both the underlay and overlay networks, providing resilient and scalable network communication.

Our standard network backbone boasts a bandwidth capacity of 10 Gigabits per second (10Gbps), with a scalability feature that allows expansion up to 100 Gigabits per second (100Gbps) and in some cases up to 400 Gigabits per second (400Gbps) . This scalability facilitates dynamic accommodation of escalating data transfer demands, guaranteeing optimum performance under varied load conditions.

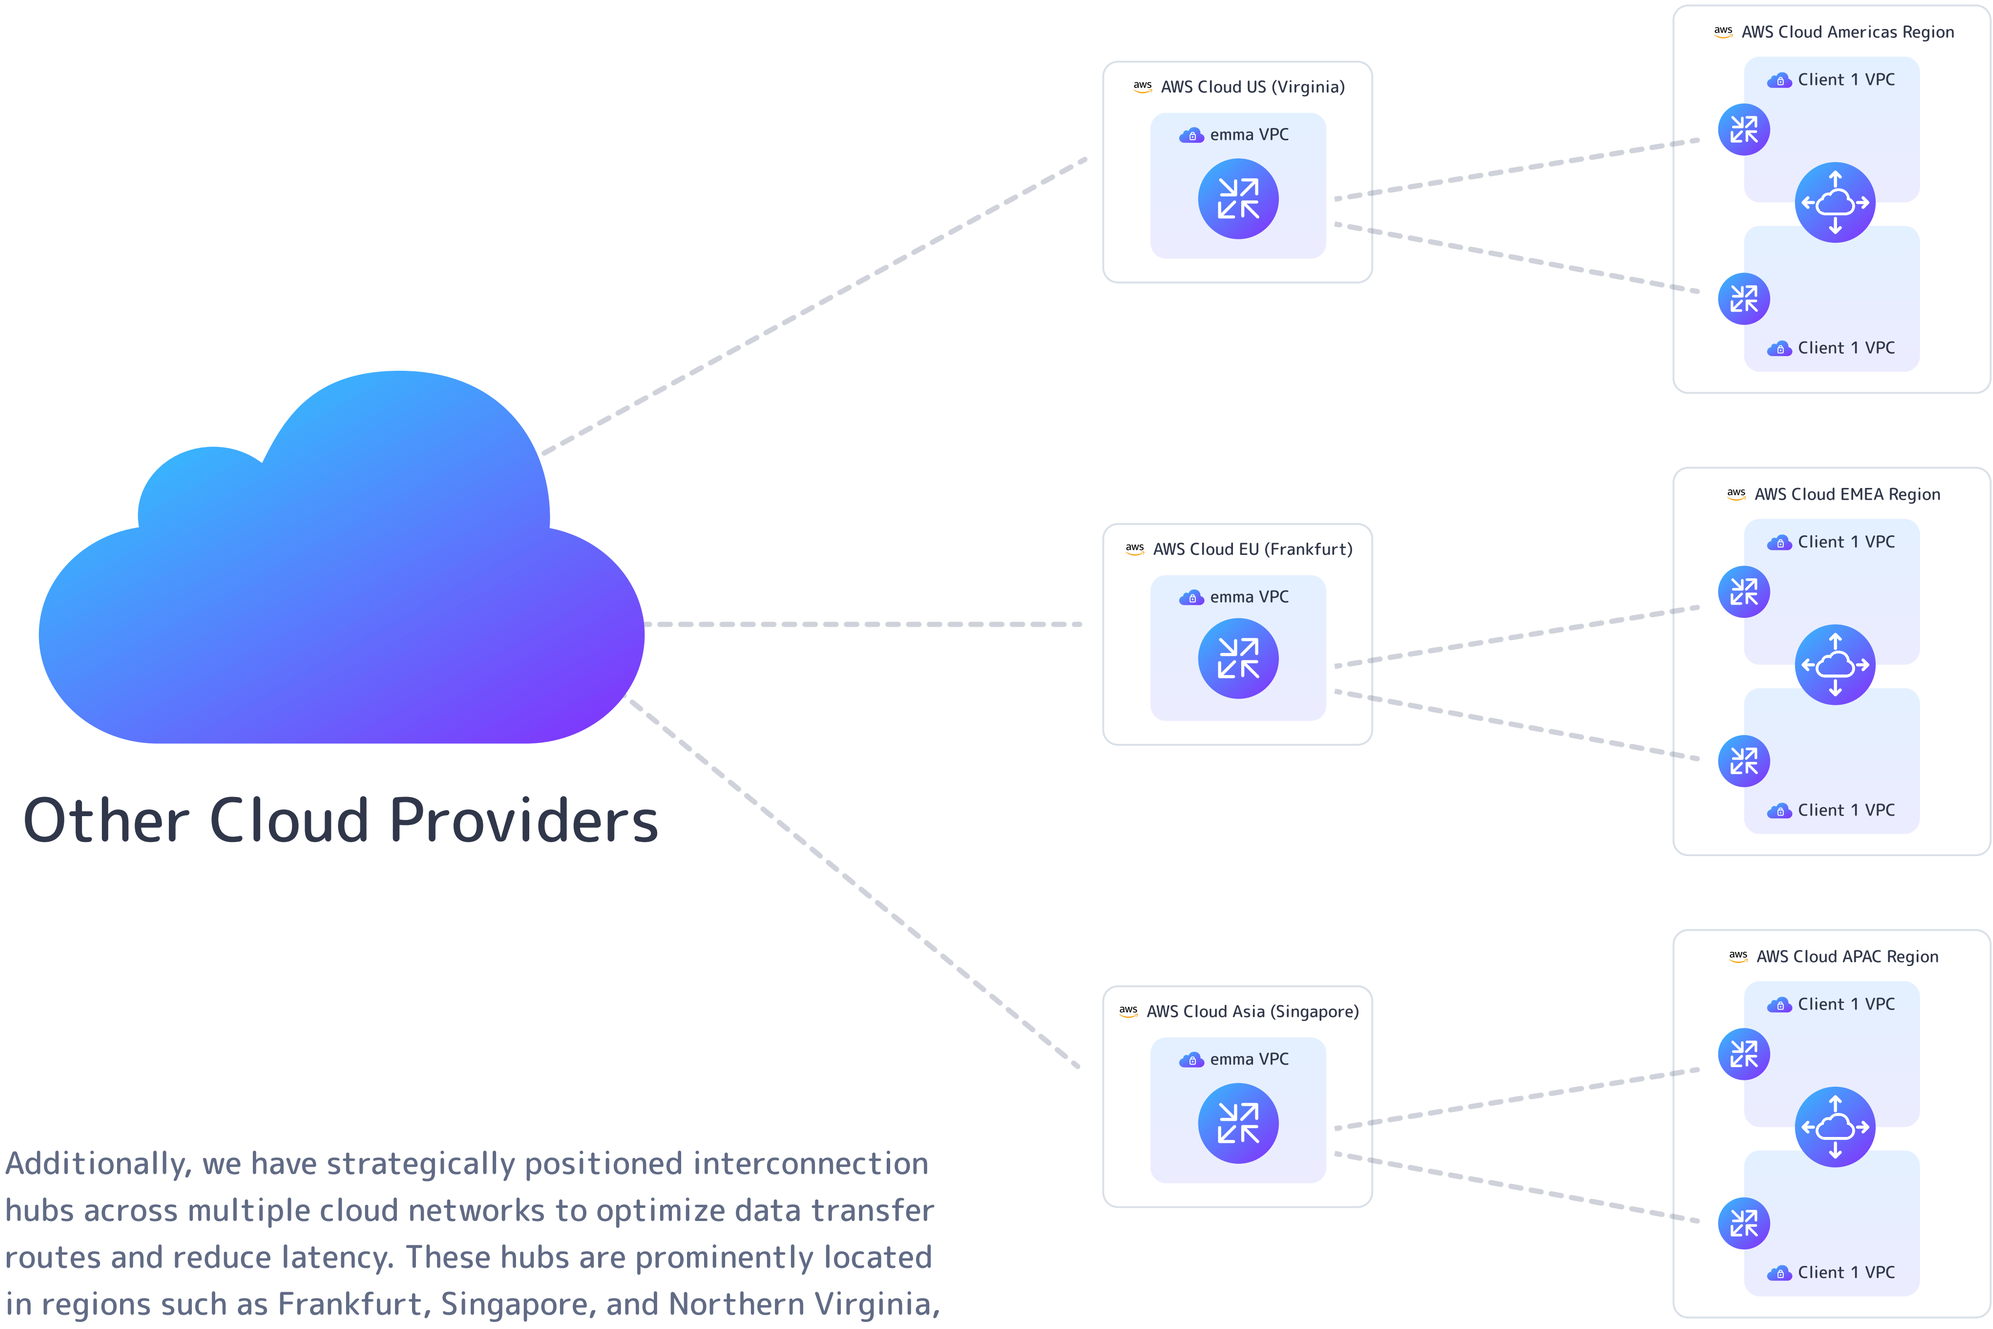

Additionally, we have strategically positioned interconnection hubs across multiple cloud networks to optimize data transfer routes and reduce latency. These hubs are prominently located in regions such as Frankfurt, Singapore, and Northern Virginia, serving as pivotal points for inter-network data exchange.

The following table provides a detailed overview of the actual latency times between various service providers and regions within the emma network backbone, showcasing the efficiency and speed of our global networking solution.

		To AWS			To Azure			To GCP		
		Asia 10.6.1.198	EU 10.7.1.243	US 10.8.1.94	Asia 100.64.92.4	EU 100.64.93.4	US 100.64.88.4	Asia 198.18.30.3	EU 198.18.31.3	US 198.18.32.3
From AWS	Asia 3.112.171.214		201 ms	98 ms	146 ms	242 ms	344 ms	70 ms	225 ms	280 ms
	EU 34.249.163.203	200 ms		117 ms	255 ms	42 ms	180 ms	180 ms	25 ms	117 ms
	US 34.212.138.78	98 ms	117 ms		366 ms	172 ms	134 ms	290 ms	155 ms	66 ms
From Azure	Asia 20.210.102.52	153 ms	255 ms	368 ms		216 ms	97 ms	79 ms	233 ms	305 ms
	EU 20.77.76.91	249 ms	44 ms	174 ms	217 ms		135 ms	175 ms	19 ms	111 ms
	US 20.3.188.118	367 ms	182 ms	132 ms	97 ms	135 ms		295 ms	158 ms	68 ms
From GCP	Asia 35.213.158.179	70 ms	179 ms	290 ms	79 ms	175 ms	293 ms		295 ms	205 ms
	EU 35.207.148.77	232 ms	25 ms	156 ms	233 ms	19 ms	158 ms	295 ms		87 ms
	US 35.212.39.198	304 ms	116 ms	67 ms	305 ms	111 ms	69 ms	207 ms	87 ms	