Cloud networking promises to resolve multi-cloud connectivity issues but will disrupt the industry

June 2022

Caroline Chappell and Gorkem Yigit

Over the next decade, we expect that a combination of workplace trends, digital transformation, Industry 4.0 supply chain transformation, edge, AI and 5/6G will drive enterprise use of high numbers of distributed private and public clouds. Today, however, enterprises must navigate an increasingly fragmented connectivity landscape to support app-to-app communication across multiple/hybrid clouds. An emerging category of cloud networking products promise to meet market needs for connectivity and security and be consistently programmable across network domains and multiple clouds, including edge clouds.

Current approaches to multi-cloud connectivity are failing

Cloud technologies are becoming the substrate on which enterprises run their businesses. Companies may implement such technologies themselves on-premises, procure them as a service directly from public cloud providers or indirectly through software-as-a-service (SaaS) providers, but the direction of travel is clear. Enterprise adoption of cloud technologies fuelled public-cloud provider revenue growth of more than 30% in 2021, it is creating multi-billion dollar PaaS and SaaS companies and is driving the need for application connectivity across a cloud landscape that becomes more hybrid and distributed each month.

The latter trend is creating a new breed of multi-cloud connectivity provider, such as Aviatrix, which boasts a healthy valuation of around USD2 billion thanks to its bet that most enterprises are already using multiple clouds across which they need to run business processes. The keys to multi-cloud connectivity are provisioning speed and flexibility. Application developers, particularly those in ‘born-in-the-cloud’ companies, want to spin up inter-cloud networks quickly so they can create cloud compute and storage in the individual clouds that they wish to connect, preferably directly from their continuous integration/continuous deployment (CI/CD) pipelines.

This is not a new requirement, of course. We have seen the wave of demand for multi-cloud connectivity coming for years. Traditional operators have tried to prepare for it by building more automation into their underlay networks but this is proving to be a slow process. Operators are also reselling, or in a handful of cases have built, software-defined overlay networks that provide more automation but which come with drawbacks in a post-pandemic world. SD-WANs depend on best-effort internet, unless an operator expensively integrates them with its own underlay network, and have a branch-based focus in a world moving to hybrid working. Cloud-based security (SASE) solution providers are pitching in to address the issue of securing hybrid workers’ access to cloud services, but they are exacerbating the problem. Now enterprises, which already have to stitch together, or ask a supplier to connect, fragmented underlay and overlay networks from numerous providers to reach multiple IaaS, PaaS and SaaS services, have to factor SASE in, too. Enterprise traffic is tromboning across networks through complex service chains that operators say can take them 6 months or more to set up. This is hardly the frictionless, on-demand cloud experience that their customers are looking for.

The emergence of cloud-based IP routing planes is both an opportunity and threat for vendors

Now that the wave is upon them, network operators, both established and new, need a new approach to routing traffic across a continuum of core and edge clouds. Current solutions for multi-cloud connectivity are merely temporary fixes. A radical rethink of cloud networking is required, which takes advantage of the fact that slowly but surely, cloud technologies are becoming the substrate for networking, too.

The industry has an opportunity to develop a routing control plane for multi-cloud connectivity that is itself cloud-based, built using cloud-native technologies and deployed in a distributed way across an IP network cloud. Such a cloud would span multiple network domains, supported by different types of user-plane hardware in clouds, data centres and wide-area network (WAN) PoPs, using Control/User Plane Separation (CUPS) principles. The cloud-based control plane would naturally implement zero trust because microservices-based Layer 4–7 security functions could be co-located with routing microservices and other functions, for example, path computation and analytics. Such functions could be accelerated by emerging white box-based user-plane hardware that uses Intel Tofino chips, Nvidia GPUs or Broadcom chipsets.

Successful IP network clouds are likely to be built by an ecosystem of technology vendors in a way that mirrors the disaggregation of the radio access network (RAN). However, it will result in even greater disruption than Open RAN because there are many more vested interests in the IP market. Already, a variety of vendors are approaching this opportunity from different directions, including established networking vendors, such as Cisco, F5 and VMware, new entrants, such as Arrcus, DriveNets and Kaloom, new service providers such as Cato Networks and InterCloud and public cloud and data-centre providers. In an as-a-service cloud networking world, the boundaries will blur in future between service providers and vendors, a further threat to traditional network operators (Figure 1).

Figure 1: Potential architecture for an IP network cloud

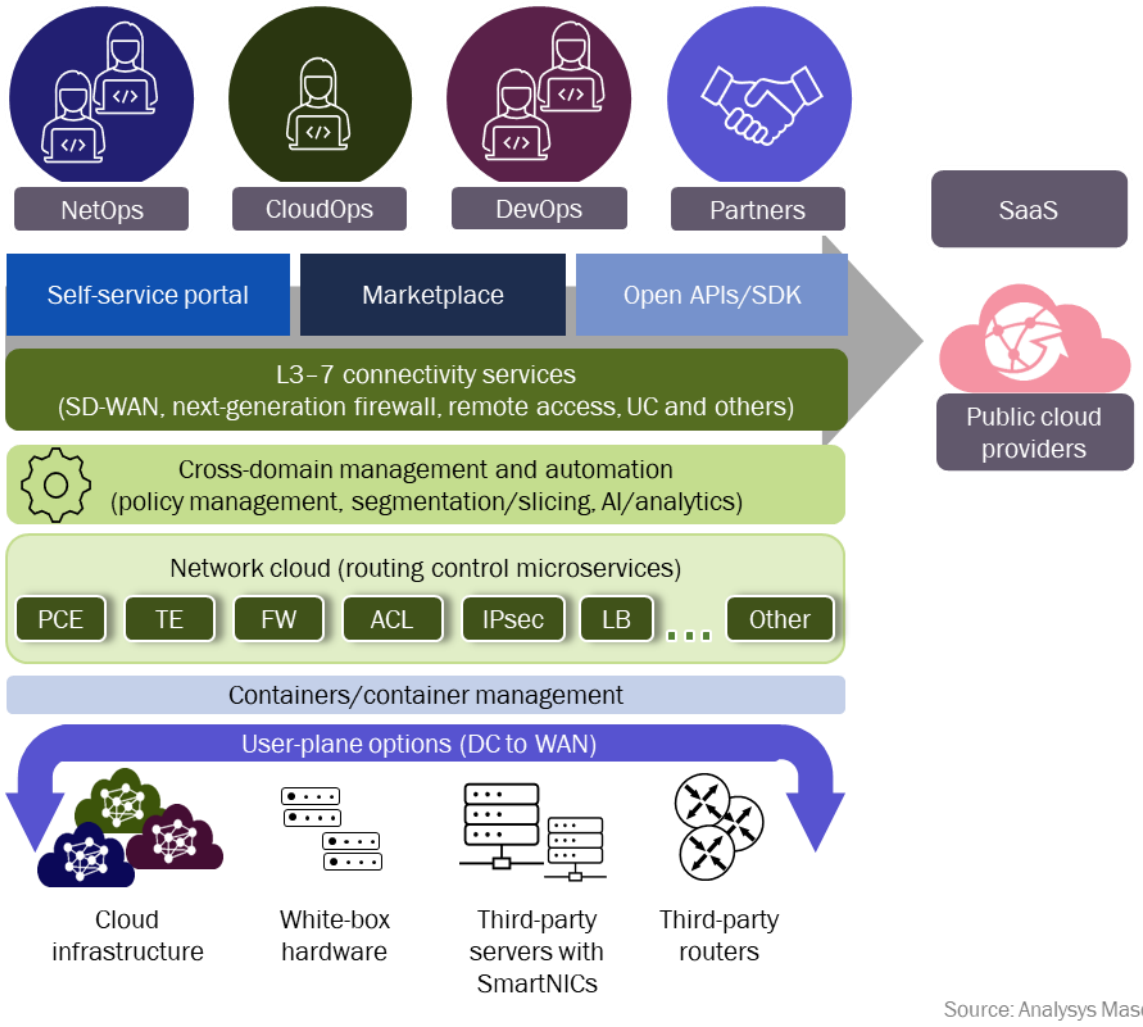

Source: Analysys Mason

Analysys Mason is at the forefront of the cloud networking trend

Our recent report, *SDN in the age of multi-cloud connectivity: the case for a converged programmable network* examines the drivers for software-defined networking (SDN) products to support multi-cloud connectivity. It evaluates the contenders for an IP routing control plane that converges L2/3 routing to L7 application firewall functions, which are currently realised differently in diverse network domains, and makes the case for a single, consistent, cloud-native-based and programmable network that can span multiple cloud environments.

In addition, we are launching a new *Cloud Networking* research programme within the *Cloud* practice in 3Q 2022, which will track vendor and challenger service provider progress as they address multi-cloud connectivity challenges. We expect this new SDN development to transform the connectivity landscape over the next decade.